



AI IN L&D: AROUND THE CORE

By **Donald H Taylor** and **Eglė Vinauskaitė**

October 2026

Sponsored by





THE L&D SURVEY SERIES

AI in L&D

2026: AROUND THE CORE

© Donald H Taylor Services Limited 2026, Eglė Vinauskaitė 2026. All rights reserved.

Cover and layout by Wassaa.

INTRODUCTION

This is our fourth annual report on L&D's use of AI, and the sixth in the L&D Research Series since the public launch of ChatGPT nearly four years ago. In that time, L&D's use of AI has moved from novelty to normality. Two years ago, about 40% of respondents said their organisation was using AI in L&D. Last year the figure passed half, a shift we called the Implementation Inflexion. This year it is over two thirds.

With AI now a regular part of the work of L&D, it is worth asking why we continue authoring these reports. When everyone uses AI, what is there left to talk about? And that is exactly why we continue – to understand **how** exactly people are using AI. Has it changed the way they work, and has it changed how the organisation works with them?

The answers to these questions are what make up this report. They are one part of a much bigger question about what AI means for L&D. While there is important work to be done in imagining what the function could become, there is also value in understanding the practice taking shape now. Teams are starting from very different places. What is established practice for one may still be a useful next step for another; what the industry has already learned can also become a foundation for those pushing further.

In search of answers, we changed how we approached the survey. In previous years, respondents chose their uses of AI from a multiple-choice list, and the results barely moved. This year, we went back over four years of respondents' own descriptions of their usage to analyse the trends within it. The result gives this report its title. L&D has not stood still. It has built around its core. Content production and learning design remain the most likely ways that L&D will use AI. What has changed is both the detail of how these are carried out, and what is built around them. Expectations of AI have shifted too, and so have the obstacles standing in the way.

In **Deep in the detail**, we look beyond the numbers at what practitioners are doing with AI today, category by category. In **What about agents?**, we test the year's most talked-about term against what respondents actually describe. The case studies then show the work in context: real people, in real organisations, tackling real problems with AI. Finally, **The challenge of change** sets this year's findings against the models we have built since 2024.

Our thanks, as always, to our sponsors, without whom this research would not happen: Sana, Docebo, OpenSesame and Speexx. Our thanks also to the 531 practitioners in 55 countries who told us what they are doing. As ever, we look forward to your comments on the report.

- Donald H Taylor and Eglė Vinauskaitė



ABOUT THE AUTHORS

Donald H Taylor

donald.taylor@donaldhtaylor.co.uk

Donald is the founder of the L&D Research Series, and has chaired the Learning Technologies Conference in London since 2000, most recently co-chairing it in 2026. He writes and speaks worldwide about his research and preparing learning and development for the future. His annual L&D Global Sentiment Survey, started in 2014, provides a unique perspective on L&D trends from some 4,000 people in over 100 countries. The AI in L&D series of six reports, co-authored with Eglė Vinauskaitė, have been widely cited as providing an invaluable guide to understanding the practical use of AI by learning and development practitioners.

He currently advises several EdTech companies and from 2010 to 2021 chaired the Learning and Performance Institute. The author of **Learning Technologies in the Workplace**, Donald is a graduate of Oxford University and the recipient of an honorary doctorate from London's Middlesex University.

View LinkedIn profile:



Eglė Vinauskaitė

egle@nodes.works

Named one of HR's Most Influential Thinkers in 2025, Egle Vinauskaitė is the award-winning director of Nodes, an advisory helping organisations build the capability infrastructure for AI-powered work. With a background spanning business, behavioural science and adult development, she brings both analytical rigour and practitioner depth to the organisational and human dimensions of how AI reshapes work. Over the past eight years she has been trusted by leading companies across industries – from technology and professional services to finance and pharma – to navigate capability building, AI strategy in L&D, and the complexity of AI transformation.

Her work has earned industry recognition, including the Rising Star Award from the Learning and Performance Institute. Egle holds an M.Ed. from Harvard University and is known for translating complex shifts in AI into clear, practical insight for leaders.


View LinkedIn profile:





CONTENTS

ABOUT THE REPORT	7
ABOUT THE SURVEY	9
FROM NOVELTY TO NORMALITY	11
HOW L&D IS USING AI	15
BARRIERS TO USING AI	24
HOW PEOPLE FEEL	31
DEEP IN THE DETAIL	39
WHAT ABOUT AGENTS?	62
CASE STUDIES	67
CASE STUDY 1 – IATC, RIYADH	68
CASE STUDY 2 – EY	72
CASE STUDY 3 – RANDSTAD	75
CASE STUDY 4 – KPMG	78
THE CHALLENGE OF CHANGE	83
CONCLUSION	86
RESOURCES	88
SPONSORS	89



ABOUT THE REPORT

THE REPORT STRUCTURE

This report has two parts. The first examines the results of a global survey of over 500 L&D practitioners from 55 countries. This is split into quantitative and qualitative analysis.

We begin with quantitative analysis – exploring what the numbers in the survey tell us about the changing use of AI by L&D practitioners.

The second part is an in-depth exploration of the same topic, drawing on their free-text survey answers and a collection of case studies. This qualitative analysis adds depth to the breadth of the survey results.

Throughout the report, where permission has been given, quotes are attributed. If a quote is not attributed, permission has been given to use the quote anonymously.

You will see some vendors mentioned in the quotes. We do not benefit from this, financially or otherwise. Also, the mention of a vendor is not necessarily a recommendation. The goal of this report is to help the L&D community imagine what can be done with AI and how to do it. After careful consideration, we decided that being specific about the tools used would make the case studies more practical and help our colleagues understand the resources involved in technology rollouts of this kind.

CAVEATS AROUND THE SURVEY DATA

When considering the survey part of this report, please also consider these caveats around the data set and the interpretation of results.

Respondents are likely to be more tech-savvy than most

Most respondents choose to contribute after seeing an invitation on social media or email. They are, therefore, a self-selecting group. Because they are contacted – and answer – electronically, respondents are certainly users of technology, and probably more likely to feel positively about technology than the general population. This method of canvassing votes excludes many, particularly those who are not active on social media.

Year-on-year comparisons may be unsound

Numbers of respondents have grown year-on-year. There is very little overlap across the years with people returning to the survey. This could lead to variations between surveys arising from changes in the make-up of the surveyed population, rather than changes in the views of the wider L&D practitioner population.

Where respondents worked




Figure 1: Where respondents worked

ABOUT THE SURVEY

WHAT DID WE ASK? WHO ANSWERED?

We ran a nine-question survey between 25 June and 6 September 2026. The survey repeated some questions from previous surveys, allowing us to draw comparisons over time.

The first two questions were optional and asked about personal details and where people worked.

They were followed by five questions asking about:

- Q3** their progress using AI in L&D (multiple choice, obligatory)
- Q4** changes in their L&D department in the previous 12 months (multiple choice, optional)
- Q5** their current use of AI in L&D (free text, optional)
- Q6** what benefits they expected to see from using AI in L&D (multiple choice, obligatory)
- Q7** barriers experienced in using AI (multiple choice, obligatory).

The final two questions allowed respondents to comment further and to say whether they were happy to be quoted in this report.

The survey was publicised via a combination of social media (largely LinkedIn) and email newsletters. The 531 respondents were self-selecting. As such, they are unlikely to be representative of the overall population of L&D practitioners. Also, because they volunteered to contribute, they are more likely than most to have strong opinions about AI and to have experienced working with it.

All but three respondents answered the question “Where do you work?”, with the majority working on an L&D team in the workplace (**Figure 1**).

The survey tool used, SurveyMonkey, provides IP addresses. These showed respondents were split across the United Kingdom (26%); the rest of Europe (36%); North and Central America (15%); South America (5%) and the South-west Pacific (7%) – this being largely made up of responses from Australia. The remaining respondents were distributed across other regions.

AI becomes business as usual for L&D




Figure 2: AI becomes business as usual for L&D

FROM NOVELTY TO NORMALITY

L&D NOW USES AI REGULARLY

Question 3 asked respondents “How would you describe your progress in using Artificial Intelligence in workplace L&D?” Six multiple-choice answers were provided. The wording was identical to the surveys of the previous two years, and very close to that of 2023:


1. Our organisation has no intention of using AI as part of L&D
2. Our organisation has not begun exploring the role of AI in L&D
3. Our organisation has experimented with AI tools in L&D, but has not implemented anything
4. Our organisation is currently piloting the use of AI in L&D
5. Our organisation is using AI in some parts of our work in L&D
6. Our organisation is using AI extensively as an integral part of our work in L&D

The six answers can be grouped into three phases:

- 1, 2 – Pre-use
- 3, 4 – Pre-implementation, experimenting or piloting
- 5, 6 – Using AI in some way

As **Figure 2** shows, last year over half of respondents said they were using AI in some way. This year that share rose to over two thirds.

Last year, we called this shift the ‘Implementation Inflexion’. With over 50% of respondents reporting they used AI, it was clear that a wave of enthusiasm was pushing AI use forward. This year, it seems safe to say the wave of enthusiasm has been replaced with something else. In four years, AI has made its way from novelty to normality.



SETTING THE SCENE

The richest source of data in our survey has been our regular question, asked since the inception of the survey: “If you are using AI in L&D in any way, please share what you’re doing.” Although optional, this question was answered meaningfully by 90% of respondents, generating a total of more than 15,000 words. The average response was just under 32 words long.

The answers from this year’s survey provide a snapshot of how L&D is using AI today. When combined with the answers from previous years, the total dataset of over 50,000 words gives an invaluable insight into how the use of AI has altered since the public launch of ChatGPT.

We analysed these answers in two ways. In **Deep in the detail**, we take a qualitative look at what people are doing with AI today. In this section, we use a quantitative approach to examine how patterns of AI use have shifted over time. We do this in a simple, methodical way: by counting words. This apparently primitive research technique serves two vital functions. At the beginning of the research process, it is a simple way of identifying things worth investigating – if the frequency of a word increases over time, there must be a reason for it. In addition, it acts as a counterbalance to confirmation bias. Any ideas that occur to us while reading the responses can be checked against the word count to see whether the two can be reconciled.

The analysis starts with a check of the occurrence of 93 word stems in each response. A word stem is the part of a word which enables us to identify it and its derivatives. Counting occurrences of ‘analy’ includes uses of ‘analysis’, ‘analyse’ and so on. The stem ‘skill’ includes ‘skill’, ‘reskill’, ‘skilling’, ‘unskilled’ and other variations.

We count the number of responses a word stem appears in and then rank the word stems accordingly: which appeared in the greatest number of responses, and which in the fewest.

Figure 3 shows the 10 word stems that occurred in most comments in 2026 and their occurrence in previous years. The enduring presence in the top four positions of 'content', 'generat' and 'creat' shows content generation remains the task that AI is most commonly used for.

	2023	2024	2025	2026
content	1	1	1	1
generat	2	2	2	2
creat	4	3	4	3
learn	3	4	3	4
analy	28	8	5	5
data	37	13	9	6
develop	14	7	6	7
design	27	17	7	8
research	52	19	10	9
skill	8	10	12	10

Figure 3: Rankings of word stems, 2023–2026

Below the top four positions, however, note the rise of 'analy', 'data' and 'research', reflecting how L&D increasingly sees AI as an analytical tool with value beyond content creation. And some words have broadened in use. Initially, 'design' meant designing content. It is now used as much to describe designing learning experiences and interventions at a larger scale.

The rankings change most from 2023 to 2025. From 2025 to 2026, the changes are relatively minor. This is partly because moving up any ranking becomes harder as you approach the top. For example, it takes many more mentions to move from #10 to #5 than from #15 to #10.

Of all the word stems falling in this year's survey, none has fallen further than ChatGPT, which has fallen 71 places from its 2023 rank to #84 in this year's rankings. Along with this, the other 10 major fallers include 'chatbot' (down 32 places), 'generative' (down 40) as well as 'script' (down 39), 'outline' (down 32) and 'voice' (down 28).

These falls illustrate two trends. First, they do not mean people are not using ChatGPT. We know anecdotally that they are. However, as AI moves into general use, it becomes a generic, accepted tool, and people feel less need to mention the precise tool they are using. This is seen with Claude, a large language model (LLM), which ranks only #69 this year, despite its apparent popularity. NotebookLM fell 60 places in just one year – from #87 last year to #147 this year.

Similarly, the fall in the frequency of use of 'script' and 'voice' does not mean AI is not being used for audio work. It is. Almost certainly people find that use unremarkable compared with other things they are doing. The same almost certainly applies to 'outline'. LLMs summarise text far quicker than humans, and using them to create course outlines is probably now business as usual.

In last year's report we said these word counts showed "not just a wider use of AI, but an increased sophistication in its use". It also did not reveal anything dramatically new. 2026 is very much a continuation of previous years' trends and anyone expecting a shift in L&D activity will be disappointed. There are, however, more nuanced shifts, and the rest of the report dives deeper into them.



HOW L&D IS USING AI

In previous years, we have explored respondents' use of AI in two ways. First, through the open question explored above in **Setting the scene**. Second, through a multiple-choice question offering a choice of uses of AI such as **Researching subject matter** and **Administrative tasks**. This second approach, however, produced a chart that changed remarkably little over three years, while our qualitative research suggested that AI use itself was becoming more nuanced. We therefore conducted an analysis of four years of free text responses to capture and examine how AI use has evolved over time.

Across the four years, 1,473 usable responses, comprising more than 50,000 words, were included in the analysis. Responses that described only anticipated future uses of AI, or reasons why respondents could not use AI, were excluded because the analysis focuses specifically on current use. The remaining responses yielded 5,289 coded excerpts describing uses of AI.

Developing and applying the categories was a painstaking process involving several stages of analysis and validation, which both authors carried out independently before comparing their results. Starting with the same existing human-coded data, each author used AI to independently derive categories from the free text responses and develop a separate codebook. These AI-derived categories were then systematically reconciled with the human coding framework, examining areas of agreement as well as differences in boundaries, granularity and structure. The two resulting analyses were sufficiently similar that we were able to agree readily on which to adopt for the final analysis. The coding was also subjected to multiple rounds of human and AI validation, including deliberately challenging boundary cases and independent repeat classification. In a validation sample weighted towards ambiguous cases, agreement between the AI classifier and researcher was 85% at the category-family level; the classifier's family-level self-consistency was 94%.

This retrospective approach to developing the categories differs importantly from the multiple-choice question used in previous years. Those categories represented our best assessment, at the time the survey was designed, of the ways AI might be used in L&D. Here, the process works in the opposite direction: with four years of responses now available, we could identify the patterns and distinctions that have actually emerged in practice, and then apply that framework retrospectively across all four years. This allows us to examine change using categories grounded in what respondents themselves reported doing, including distinctions that would have been difficult to anticipate when the original survey was designed.




TEN CATEGORIES OF USE

The ten categories of AI use in L&D are:

1. **Content development:** the creation of learning materials themselves
2. **Learning design:** shaping the learning experience, its structure, objectives and activities
3. **Administration and communications:** supporting the L&D department's routine work
4. **User-facing delivery:** AI the individual interacts with directly, from practice and feedback to personalised content
5. **Discovery:** using AI to support the discovery phase of learning development work
6. **Skills intelligence:** understanding and mapping the capabilities of people at work
7. **Knowledge and point-of-need support:** organising knowledge and providing it when needed
8. **Evaluation and analytics:** measuring and reporting whether learning activities worked
9. **Operations:** L&D organising its own work, including workflow automation, internal tools and setting learning strategy
10. **Principled non-use:** respondents making the deliberate decision **not** to use AI or use it in limited ways only

Answers to the question “How are you using AI in L&D?” can fall into one or more of these categories, and over time, this number has increased. In 2023, answers fell into an average of 2.1 categories. By 2026 that number had increased to 2.7.

Use of AI in L&D by category 2023–2026



Circle area proportional to share (%). Ties broken by count, then next year's share. Δ = rank change 2023 to 2026.

Source: "AI in L&D 2026: Around the Core" by Donald H Taylor and Eglė Vinauskaitė, nodes.works/reports

Figure 4: Use of AI in L&D by category 2023–2026

Figure 4 shows how the categories of use have shifted over time, represented by the percentage of respondents in each year who described at least one use case in that category. Because respondents typically named several use cases, the categories are not mutually exclusive and the percentages within a year sum to well over 100.

As we have already seen in the word stem analysis, content dominates L&D's use of AI, which is clearly visible in the top two lines in **Figure 4**. It is what almost everyone does first. The change is happening further down the table.

Snaking its way across the chart, one category stands out: **Evaluation and analytics**, which rose from #8 in 2023 to #3 in 2026. There was a notable increase in using AI to analyse data and produce reports and dashboards, and mentions of data itself quadrupled over the same period, from 5% to 20% of responses. These uses are detailed later in the report, in the section **Deep in the detail**.

Operations followed a similar path, from #9 to #6, mainly driven by automation and internal tool building.

And while **Administration and communications** fell in the rankings, its overall share of comments increased from 21.3% to 24.9%. This fall in ranking is unlikely to mean less administration is happening. More likely, routine usage, such as email or document writing, has simply stopped being mentioned, as we saw with the fall in mentions of ChatGPT.

Discovery held its place at fifth, and its overall share barely moved, from 22.3% to 23.5%. The change is in what sits inside it. Needs analysis, including understanding the audience, grew from 7.5% of respondents to 9.8%, while researching the subject matter itself declined from

5.8% to 4.7%. The emphasis has moved from AI helping practitioners understand a topic to AI helping them decide what should be built, and for whom.


Skills intelligence looks like a category in decline, and on the chart it reads that way, dropping from #6 in 2023 to #9 in 2026. The underlying figures tell a flatter story. It has moved between 9.5% and 13.5% of respondents across four years without ever establishing a direction, and this year it recovered to 12.6%, close to where it started. What looks like a fall is mostly other categories rising past it. One promising sign is that building skills frameworks and taxonomies has grown modestly but steadily, making it the only part of the category with a consistent trend. However, for a topic that occupies a substantial share of industry conversation, the practice described by respondents remains notably thin.

Two categories that put employees at the centre of the impact of L&D's activity have shown no positive direction over the four years. **User-facing delivery** and **Knowledge and point-of-need support** both fell in rankings over the years, although there is more going on inside **User-facing delivery** than the ranking suggests. The category peaked in 2025 and has now fallen to its lowest point in four years. The bump in 2025 came largely from interest in AI coaching, which has since receded. Additionally, system-level personalisation, such as adaptive learning and AI-powered content recommendation, has been falling for four years. What held steady is practice and simulation, the only part of the category still growing.

Knowledge and point-of-need support tells a different story. It increased its share of responses marginally over the period, but the rise is not impressive:

from 9.9% to 12.8%. Its modest rise is carried entirely by curation, which grew from 4.0% to 6.4% as practitioners turned toward organising material that already exists. Point-of-need assistants, the part that would signal AI reaching people in the flow of work, peaked at 7.7% in 2025 and has since fallen back to 6.2%. Neither category has yet found the traction the surrounding conversation would suggest.

Taken together, four years of data describe a function that has built around its core rather than moved away from it. Content and learning design have sat at the top in every year, and the growth has come around them: in working out what is needed before anything is designed, in analysing and evaluating what happens afterwards, and in automating how the department itself runs. The spread of uses per respondent widened sharply in 2025 and then held, and no new category has emerged. Meanwhile AI that reaches learners directly peaked in 2025 and is now at its lowest point in the series. For most of the industry, AI is still something L&D applies to its own work.



EXPECTATIONS OF AI

In an obligatory, multiple-choice question, we asked respondents to “Select up to 3 benefits you expect to see from your use of AI in L&D”. We first asked this in 2023, but in 2024 we expanded the list to include data analysis, quality of learning design and point-of-need support, reflecting how the uses of AI in L&D had developed. Because respondents can choose only three, a longer list changes how votes are distributed, so it is impossible to make direct comparisons with 2023. We therefore compare 2024 with 2026, the first and most recent years in which the question was asked in exactly the same form. The two years also span the period in which L&D's use of AI broadened most, as mentioned earlier.

The short answer, as seen in **Figure 5**, is that **Creating learning content faster** at the top of the list has not moved. What sits beneath it has, and those changes reveal something of the mindset of L&D when dealing with this technology.

Four things stand out from these results. First, the greatest perceived benefit of using AI remains the ability to create content faster. It seems that content production still dominates thinking in our field, just as it dominates what L&D does with AI.

Second, the leap in **Use data analysis to support decision-making**, from 31.9% to 44.5%, moving it from third to second. This echoes **Figure 4**, where the **Evaluation and analytics** category of uses rose to 34.1% of answers in 2026.

Third, **Improve the quality of learning design** rose from 31.0% to 37.1%, climbing from fifth to third, in line with the growth in reviewing and quality-checking described earlier.

Fourth, **Improve efficiency/reduce costs within L&D** moved the other way, falling from 41.9% to 36.0% and from second to fourth, the largest fall on the list. Because respondents could choose only three benefits, this reflects priority rather than belief: it does not show that L&D expects AI to save less, only that fewer now place cost-saving in their top three. Notably, it has not been displaced by the closely related **Create learning content faster**, which held steady. The ground it lost appears to have gone instead to decision-making and quality, the two benefits that rose most.

Taken together, L&D's expectations of AI seem to be shifting from what we might

call 'first-order' benefits to 'second-order' ones. The first-order benefits, doing existing work faster and more cheaply, have held steady or fallen. The second-order benefits, making better decisions and designing better learning, are the only ones rising. Expectations have moved in step with practice: what L&D hopes to gain from AI now mirrors what it is actually doing with it, a stable core of faster content with growing attention to analysis and quality around it. And almost everyone now expects something from it, with the share saying they are not yet using AI falling from 7.4% to 1.7%.

Expected benefits of using AI in L&D




Figure 5: Expected benefits of using AI in L&D

The 'break-out' group against other workplace L&D teams, 2026




Figure 6: The 'break-out' group against other workplace L&D teams, 2026

BUILDING BEYOND THE CORE

So far, we have seen that the trends of 2024 to 2025 appear to be continuing into 2026. L&D has largely settled into a pattern in how it works with AI, using it to accelerate and enhance its core work – notably the production of content and the design of learning activities and interventions.

A small group of respondents, however, seems to be using AI subtly differently. They are not ignoring its value in helping them design and create learning content. That core activity dominates responses to our survey and this group is no exception. Rather, they are doing this core activity and then building beyond it.

We found this group through a new question introduced in this year's survey, added to explore how AI is changing L&D itself: "Which of the following has changed in your L&D department in the past 12 months, as a result of AI?" In this part of the report, we explore responses from the 298 members of workplace L&D or HR teams who answered. We have narrowed the field to L&D/HR teams because the

question asks about "your L&D department", which consultants, suppliers and educators may not have. Respondents could tick more than one option, so the percentages in this list sum to over 100%:

- (a) No significant changes. We work now in the ways we always did. (92 responses, 21%)
- (b) We have changed, but not primarily because of AI. (71, 16%)
- (c) The rest of the organisation now approaches us for different things than before. (43, 10%)
- (d) At least one function our team used to do is now routinely handled by AI. (96, 22%)
- (e) At least one existing role has changed in nature (not just the tools used) because of AI. (58, 13%)
- (f) Most of our team now directs and reviews AI work rather than doing the work directly. (60, 14%)
- (g) L&D no longer exists as a centralised, dedicated function in our organisation. (12, 3%)

By themselves, these responses show a profession in flux, with change being driven by AI for most respondents – of the 298 respondents, 58% reported at least one change caused by AI. The data also raise a question: do respondents choosing one of these answers differ in any way from the rest of the workplace L&D/HR teams?

The answer is yes. Those picking (c), **The rest of the organisation now approaches us for different things than before**, form a group with some differences from the mainstream of workplace L&D. We call this the ‘break-out’ group, and the differences are explored in **Figure 6**.

Of the 43 respondents who chose **The rest of the organisation now approaches us for different things than before**, 40 describe their AI use. Their descriptions differ from those of other teams in two areas (**Figure 6**). They are twice as likely as other workplace L&D teams using AI to describe user-facing delivery: offering skills practice, coaching and dynamic feedback for learners. They are also almost three times as likely to describe using AI for operations work, which includes automating tasks and building internal tools. This greater use of AI for operations work could be a signal that the L&D team is more reachable, responsive, or better integrated with the rest of the business.

Notably, the ‘break-out’ group has not cut back on content development either. Of this group, 88% describe doing it, against 80% of other workplace L&D/HR teams. Nor is this simply heavier AI use. A third of this group report using AI extensively, against one in ten of other teams, but most do not.


The pattern, then, is addition rather than substitution. This rather small group (about 14% of the workplace L&D and HR teams who answered the question) does the core content-heavy work of L&D

and, alongside it, automates internal work (45% against 17%) and puts AI in front of learners (38% against 17%). What distinguishes them is not how much AI they use but what they use it for. While the group is small, the differences on operations and on user-facing delivery are too large to be the result of chance.

The differences between the break-out group and the other workplace L&D/HR respondents also extends to their expectations of the benefits of AI, as explored in the previous section. Only 37% of the break-out group chose **Create learning content faster** as an option, against 63% of the other L&D/HR teams. And while 6% of the others chose **Provide extra knowledge testing**, 21% of the break-out group did (again, these differences are statistically significant). Their attitudes towards how AI can help L&D in its role differ significantly from the rest of workplace L&D.

Two caveats should, however, be remembered. First, these respondents tell us the organisation now approaches them for different things, but not why. It is tempting to assume that the organisation comes to them because they are offering something different and better than in the past, but it may be that the organisation has asked L&D for help with something specific, such as adopting AI, and the new work followed. Two of the responses point that way. Second, being approached is not always good news. One respondent reported losing their job after running an AI training programme.

Even allowing for both, the finding stands. Something changes the way an organisation sees its L&D team, and in this data it is not the volume of AI work but its direction – outward, to the business and to learners, rather than inward.



BARRIERS TO USING AI

In question 7, respondents were asked: “What barriers, if any, are you experiencing in your work with AI?” This was an obligatory question and they could choose any number of responses from a list of 13. **Figure 7** shows what share of the people responding annually to the question chose each option. The most obvious difference is the addition of a new option this year – **Ethical and environmental concerns** – added in response to write-in comments on these issues in last year’s survey.

In 2026, the average number of options chosen fell slightly to 2.89, from 3.03 in 2025, having risen from 2.62 in 2024. The fall came despite a new option being added this year. Set it aside, and respondents named fewer barriers in 2026 than in either previous year: 2.59.

Almost every barrier followed the same shape, rising in 2025 and easing in 2026.

Data privacy and security concerns remains the most common by a wide margin, named by 52% of respondents, though down from 59% last year.

Ten of the 13 options fall into three groups:

- ▶ L&D’s own **readiness**: skills, knowing where to start and making a business case.
- ▶ **Friction** of putting AI to work: integration, cost, time and resistance.
- ▶ **Risk**: data privacy, trust in outputs and regulation.

These are shown in **Figure 8**, where we show the share of respondents naming at least one barrier in the group.

Barriers to using AI, 2024–2026




Figure 7: Barriers to using AI, 2024–2026

Group	Options in the group	2024	2025	2026
Readiness	Skills, not knowing where to start, business case	49%	49%	36%
Friction	Integration, cost, time, resistance	60%	68%	66%
Risk	Data privacy, trust in outputs, regulation	69%	77%	74%
Risk+	Above plus ethics	69%	77%	77%

Figure 8: The barriers grouped, 2024–2026

HAS L&D'S READINESS IMPROVED?


It appears so. The barriers which L&D can address by itself held at 49% between 2024 and 2025, then fell to 36% this year. Not knowing where to start has halved, from 19.8% to 9.8%, as has difficulty making a business case, from 11.9% to 5.8%.

Some individuals still struggle to keep up. One respondent wrote simply:

"[the barrier is] My own ignorance of what's possible. It's so hard to keep up on changes in the AI world."

But this comment was unusual in saying barriers belonged in the respondent's own knowledge. Most place it somewhere else (for more, see **How people feel** below).

One caveat: this is self-reported, and fewer people naming a barrier can mean several things. It may have been overcome, or it may simply feel less pressing than the ones that have taken its place. It may also mean that teams have settled on the uses they know work for them, rather than taking on ones that would test them further. That would fit what we saw earlier – a core of uses that has held steady while the work around it grew. But the direction is clear. As the barriers L&D can solve on its own fall away, the ones that remain sit largely or partly outside its control: organisational friction and risk.



PUTTING AI TO WORK IS NOW THE BIGGEST CONSTRAINT

The barriers concerned with making AI work in practice rose from 60% in 2024 to 68% in 2025, and have barely eased since, at 66%. Integration, cost and time are all higher than in 2024; only resistance has fallen. **Integration issues** has grown the most and is now the third most chosen barrier overall. It's also a barrier typically encountered when moving beyond surface-level, single chatbot-based AI use, which fits with the idea of L&D practitioners building around the core of L&D activity.

Respondents describe AI arriving faster than the organisation can absorb it:

“

“Our organisation is very keen to explore AI but doesn't have the structures in place to do this efficiently or effectively [...] approval required from 3 separate teams.”

“Every tool has integrated AI into their offering making it difficult for IT departments to keep up with the need to review and approve all these different solutions.”

“Everything has to be a 6-figure contract with an external vendor, which is pushing us to use shadow IT.”

”

“

“My wish is to have a local AI that might be ethical and trained in a kind of environment (social, economic, environmental) that doesn’t exist yet.”

”



RISK STILL DOMINATES, BUT IS CHANGING SHAPE

Risk remains the largest group, though the established risks have eased this year: data privacy fell from 59.2% to 52.0%, and trust in outputs from 39.4% to 36.4%. Sometimes organisational fear of risk is outside L&D's control:

"Due to security limitations AI in the defence sector is handled with care. Organisationally we are developing our own AI for extensive use whilst "sand boxed" versions of CoPilot are utilised for general admin and large document development."

The comments show risk turning into a cost in time.

"You MUST check everything. If you are spending a lot of time doing this then the claimed benefits of AI saving time are negated."

Three respondents used the word 'negated' in this way. Their complaint is not that AI fails, but that managing the risk cancels the time it saves.

This year's survey added a new option, **Ethical and environmental concerns**, which entered fifth of the 13 options and belongs in the Risk group. It is what keeps the group at 77%: without it, risk would have eased along with almost everything else. Where the established risks ask whether AI is safe for the organisation, this one asks whether AI is right at all:

"Will AI really do what they say without crippling the economy, hurting the environment or destroying the working class?"

PAST THE FIRST HURDLE

Taken together, the barriers describe an industry that has cleared its first hurdle. In 2024, the friction of putting AI to work was named by 11 points more respondents than L&D's own readiness; in 2026, the gap is 30 points. It seems that L&D is less held back by knowing how to start and is now constrained more by the step that comes after.



HOW PEOPLE FEEL

What did respondents really think of AI and its use in 2026?

The numbers in this document report what can be counted. They do not record how L&D feels. For that, the survey leaves spaces for respondents to comment, and in 2026, 130 people used them. The percentages below are of those 130. Thirty of them fell into small groups; we only report here on the largest six clusters of responses.

Sorted by sentiment, the comments split into three main themes. About a third (32%) are positive: 17% straightforwardly enthusiastic, 15% confident but explicitly bounded. It is more difficult to give an exact number for the negative comments because here the sentiment builds on itself and overlaps: 13% are sceptical, 10% afraid for L&D's future, 6% troubled by the impact of AI on thinking. Beyond that, 15% are plain angry – not with AI, but with their own organisations.

While the enthusiasts are brief, the critics use more words and hit harder, painting in some cases a truly dark picture.

One further pattern is particularly surprising. Where respondents are positive, the whole comment is optimistic and optimism never appears as a second, underlying sentiment in other comments. Anxiety, however, often behaves like a drumbeat under their other comments about anything: cost, content, learners, adoption.

General thoughts

The enthusiastic comments often celebrate the innovation and new opportunities AI has presented to L&D.

"After working with AI in L&D for the last 30 months [...] my abilities to create multimedia and develop effective learning have increased in ways I never imagined possible."

– Lawrence Jacobson

"With more than 20 years experience in L&D i was very sniffy about using it, but over the past 6 months i have leant into this and now regularly use CoPilot as a starting point for workshop design [...]"

"Pretty much everything we do contains a first AI layer."

– Alvaro Aldana

Working with limits


The second-largest group is not enthusiasm but qualified confidence. Twenty respondents describe extensive use of AI within the limits set by human judgement.

"I have a strict set of parameters where I won't let Claude do any of the thinking for me."

"Judgement, verification, context and accountability still sit with the practitioner." – Andrew Jacobs

"While we are seeing clear gains in productivity and speed, human expertise remains essential to ensure learning quality, business relevance and instructional effectiveness."

– Jenny Randau Mellander



“

“For me, the biggest value is not that AI gives me a finished answer, but that it helps me iterate much faster.” – Audrey Leduc

“AI can leverage improvements and innovation in the L&D space – it has become a vital tool.”

“I strongly believe that AI helps us as L&D community to mature and grow, and I find it exciting. We need to transform to stay actual because life keeps changing and world keeps changing. L&D cannot stay the same. And that’s good :)”

”



Adoption and the organisation

One of the loudest notes in the comments is frustration with respondents' own organisations, although almost none of it is aimed at the technology. That frustration dominates twenty comments and runs underneath something else in seven more.

The complaints rest on specific details: approvals are spread across teams, insufficient licences have been issued, ownership is complicated. It's clear that sometimes adoption stalls not because of technology or a lack of user enthusiasm, but it simply hits organisational barriers.

"Too much of what we could do sits in the IT teams 'too difficult' box."

"Our L&D team has been tasked with the development of an academy for upskilling the workforce in AI. But at the same time, we don't have permissions to use any AI tools outside Copilot."

"The C-suite reads those same posts and ads and is asking why we aren't doing these things."

"Compliance / Data privacy / Data stewardship / Governance / Workers Council / IT SRA and sustainment etc etc. lots of extra work nobody sees and quite frankly no senior manager cares about."

Doubt, fatigue, and the cost of checking

Some comments are sceptical, and this scepticism has changed since 2025. It is no longer doubt about whether AI works but fatigue at the distance between claims and results.

"We're floating around in the AI bubble – making the best of it, waiting until it pops."

"having same conversations we started 6 years ago or more."


"One would think that with all the talk about AI, there would be more change in our organisation and our work. But not really that much has changed. Not with our customers either."

"I hate it. The generative AI is not helping anyone."

"if you are using it as a fancy spell checker, we might have gone wrong somewhere."

Underneath the mood sits a tangible problem: verification costs time.

"I've learned about the waterbed effect: you push down certain tasks but they pop up somewhere else because you still need to extensively check the output." – Matthias Nauwelaers



Cost

Cost has arrived alongside the fatigue, and it is new. Several comments raised the price of AI, some of them as their main point. In 2025, this issue barely registered.

“biggest challenge: trying to juggle the costs as AI token costs skyrocket.”

“Assumption that AI should ‘save us money’ but actually each AI bot or tool is expensive.”

“The diversity of cost models and drivers associated with the tools and components used to create and operate agents and platforms is getting more and more complex and out of control [...]”

“most organisations that invested and experimented have drawn the focus of the CFO and the spend on tokens appears to have been wasted.”

Jobs

More personal than the organisational frustrations is the fear of what AI might mean for people's own jobs and for L&D itself. It runs through a substantial share of the free text, sometimes as the main concern and sometimes beneath a comment about something else. For some, the consequences have already arrived. For most, they are still anticipated.

“The entire Training Specialist role at my firm was eliminated in lieu of the AI tools the firm put in place. We trained the field how to use them, then they discarded us.”

“L&D is a bit like the string quartet on the Titanic. We're blissfully enhancing our performance and beautiful AI-driven tunes while management is lowering the rails to start pushing us overboard.”

“Until and unless L&D redesigns its value proposition as a strategic business partner, it will go the way of the typing pool.”

“The biggest fear for me, and a lot of others, is that today's efficiency will eventually be tomorrow's problem. Outdating the human component in L&D.”

“AI is an excellent tool that will replace some functions in the L&D environment [...] Unfortunately, leadership are putting a lot of hope in AI [...] This approach will have a negative rebound, but then it will be too late.”

“

“The debate about AI in L&D is still too focused on content generation. Its more significant value may lie in diagnosis, research, analysis, performance support and helping practitioners make better decisions about where learning is, and is not, an appropriate response.”


– Andrew Jacobs

“I have noticed that over reliance within teams have made some staff lazy.”

“AI is pushing L&D to evolve from delivering learning to shaping capability.”

– Bridie Stevens

”



Thinking

Another concern is the cognitive impact of AI. Respondents describe people outsourcing their thinking, accepting outputs without questioning them, and producing work they cannot fully explain. In some teams, these concerns are already changing how AI is used.

"I see it already with my newer team members in their design work – their proposed structures, models activities [sic] are strong but they don't know why."

"People are outsourcing their thinking to AI. People are no longer thinking critically and are delegating everything to AI without questioning its outputs. This should be a major concern for L&D departments."

*"AI is amazing when you use it to enhance thinking – but when you use it to shortcut cognition – you get rubbish."
– Claire O'Toole*

"AI allows untrained people to create bad training faster."

The opportunity

A different set of comments treats AI as a lever for changing what L&D is for. Respondents argue for moving beyond content and learning delivery towards capability, performance and a more strategic role in the business. The language is consistent across all of them, and it is the same argument this report makes in its conclusion.

"This is a unique opportunity for L&D practitioners to shift from being order takers, to being strategic advisors to the business." – Ian Walker

"if the content you were creating and the impact you were having can be replaced by AI – you were never creating effective content in the first place."

A few respondents point out that learners have already moved, suggesting that L&D will need to catch up with them.

*"Learners can learn anything they want through AI [chatbots], including myself. I have far less patience than before to watch a webinar or training program [...] our learners are no longer interested in our content."
– Zoe Qiu*

“

**“Content production is still at the centre,
and learning design itself has barely
moved. What has changed is how
much more of the surrounding work
AI now touches”**

”

DEEP IN THE DETAIL

KEY DIFFERENCES THIS YEAR

While the past 12 months have brought another big jump in what AI can do, the map of L&D work looks remarkably familiar. No major new category of activity has emerged. Content production is still at the centre, and learning design itself has barely moved. What has changed is how much more of the surrounding work AI now touches: working out what is needed, deciding what is worth using, checking whether the output is good enough, understanding what happened afterwards, and increasingly automating parts of the process. While the shape of the work may be familiar, the role AI plays within it is expanding. In practice, that shows up in several ways.

Checking becomes part of the workflow

Reviewing and quality-checking content has grown sharply. Respondents increasingly describe using AI to critique work, check accuracy, and review it against specifications or quality standards. The growth reflects greater attention to whether the output – increasingly AI-produced itself – is accurate, coherent and good enough to use.

Analytics keeps climbing

AI use for analytics and reporting has continued to expand, extending one of the clearest trends across all four years of the research. L&D teams are analysing larger and more varied datasets, interrogating learning metrics, extracting patterns from feedback and combining learning data with wider organisational measures. AI is increasingly being used not only to produce learning content, but to understand what happened as a result.

Automation starts to take shape


Agents and workflow automation are becoming more visible. Teams are beginning to automate recurring processes, connect steps across workflows and build their own tools for specific L&D tasks. Custom building is still far from universal, but it is moving beyond isolated experiments and starting to become part of how some teams work.

Multimedia fades into the background

Some of last year's most visible production uses have fallen sharply. Translation saw the largest decline, while video, audio and image generation were also mentioned less often. This does not necessarily mean these tools are being abandoned. A more plausible reading is that many of these activities have become ordinary enough to stop feeling worth mentioning: generating a voice-over, translating content or creating an image with AI is no longer novel in the way it was even a year ago.

Curation makes a return

More respondents are using AI to work with existing content rather than create something new – finding, filtering, organising and reshaping material into something more useful. Alongside the rise in review and quality checking, this marks a subtle but important change: the emphasis is shifting from creating more content to extracting more value from what is already there.



OVERVIEW OF AI USE

This section looks at the range of AI uses reported in our survey.

We grouped these into nine broad domains:

1. [Discovery](#)
2. [Learning design](#)
3. [Content development](#)
4. [User-facing delivery](#)
5. [Knowledge and point-of-need support](#)
6. [Skills intelligence](#)
7. [Evaluation and analytics](#)
8. [Administration and communications](#)
9. [Operations](#)

The pages that follow set out the wide range of uses we heard about within each of these domains.

“

“We use it heavily within the discovery piece – especially if we are trying to uncover a learning need and have done research, workshops, stakeholder discussions. We use it to pull together the data, theme it up and playbacks.”

– Isabelle Shawyer

”

DISCOVERY

Audience research

- ▶ **Understanding the audience:** Researching learner insights and target-group needs to build a clearer picture of who learning content is being designed for.
- ▶ **Profiling and segmentation:** Creating learner profiles and personas, segmenting audiences, and using AI to help define different learner groups.
- ▶ **Understanding the wider context:** Exploring stakeholder dynamics alongside audience characteristics to provide more context for design decisions.

Subject matter research

- ▶ **Exploring unfamiliar topics:** Gathering information quickly, searching for credible sources, investigating best practice, and using AI as a starting point for deeper research into unfamiliar subjects.
- ▶ **Synthesising information:** Summarising large volumes of material, comparing sources, identifying recurring themes, interrogating assumptions, analysing source material and synthesising literature or trends.
- ▶ **Building subject understanding:** Explaining unfamiliar concepts and industry jargon, researching regulatory requirements and other specialist material, and pulling together relevant information and examples for learning content.
- ▶ **Working with SMEs:** Preparing for and conducting SME interviews, analysing transcripts and SME-provided content, and identifying tasks, supporting knowledge, resources and workflow information. AI is also used to draft specialised content for SME review and validation.

Needs analysis

- ▶ **Bringing together evidence:** Analysing learner and business data, performance data and behaviour, internal sources, discovery discussions and complex client requests to understand learning and performance needs.
- ▶ **Diagnosing the underlying problem:** Identifying capability and performance gaps, analysing root causes and pain points, and examining where learning steps underperform, progress stalls or particular target groups deviate.
- ▶ **Turning evidence into priorities:** Synthesising needs-analysis data to identify common development needs and improvement opportunities, understand what people need to learn, and make recommendations about appropriate learning approaches.
- ▶ **Supporting the discovery process:** Drafting questions for training needs analyses, recording and transcribing discovery meetings, diagnosing stakeholder needs, and using AI during early scoping and performance-consulting work.

LEARNING DESIGN

Learning design strategy

- ▶ **Learning objectives:** Translating broader learning needs into learning and performance objectives for the learning experience.
- ▶ **Ideation:** Brainstorming concepts for programmes, workshops and other learning initiatives, particularly at the beginning of the design process.
- ▶ **Instructional approaches:** Proposing instructional strategies, pedagogical approaches and learning formats for digital, virtual and in-person learning, as well as ongoing performance support.
- ▶ **Initial designs:** Developing course and session outlines, curricula, learning architectures and other early solution concepts that give designers a starting point to build from. Creating initial storyboards and design mock-ups.
- ▶ **Learning journeys:** Organising content into modules and logical flows, sequencing learning to support progression, and strengthening the overall storytelling of the learning journey.
- ▶ **Design review:** Critiquing designs to surface blind spots, hidden assumptions and alternative approaches, while drawing on learning science and best practice to strengthen the proposed design.

Learning activities

- ▶ **Practice:** Creating exercises, assignments, reflection questions, practice activities and capstone projects that allow learners to apply skills and knowledge.
- ▶ **Interactivity:** Developing interactives, simulations, gamification, flipcards, workbooks and custom digital interactions for e-learning, workshops and other learning experiences.
- ▶ **Scenarios and case studies:** Generating realistic scenarios, case studies and role-play materials for workshops, e-learning, practice and assessment.
- ▶ **Workshops:** Generating workshop activities, facilitator guides, trainer notes and agendas, as well as job aids and interactive resources that support practice in the flow of work.
- ▶ **Assessments:** Generating quiz questions and open-ended knowledge checks, and reviewing or refining questions, answers and distractors.

Thought partnership

- ▶ **Exploring alternatives:** Using AI as a thinking or sparring partner to surface possibilities teams may not have considered and to develop multiple approaches before committing to one.
- ▶ **Stress-testing ideas:** Challenging assumptions, testing early concepts and using AI in a facilitator, coach or synthetic focus-group role to expose designers to alternative perspectives.
- ▶ **Refining concepts:** Developing initial ideas through dialogue, reflection and iteration rather than using AI solely to generate a list of possibilities.

"I use the AI in Articulate 360, but mainly copilot to help thought process but also presentation design including data. It helps surface ideas we may not [have] considered and overall, time wise, making us more efficient. The perfect assistant."



“We have been using Claude Code to design and develop interactive learning materials in SCORM File format which has opened up our ability to create very custom learning materials that can be hosted in our LMS effectively, we can also capture bespoke data points, and also avoiding the need for content design platforms like Articulate.”

“The place that I have seen issues is when people don’t properly define what good looks like and just ask the AI to generate outputs. I use my expertise to craft constraints based on experts, for example Patti Shank about writing assessment questions, which I can then feed to the system to get more consistent outputs. It takes some time to initially set up but with that guidance created, it can be reused and the outputs are improved.”

“I am using AI, sadly, to build content based on data from the organisation. So much data across the I/O sphere now provides data on pain points as well as poor work choices and with intelligent ID models and intelligent authoring software, agentic AI can now build learning for quick deployment when and as needed.”





“Content creation is completely AI-generated. The “speed to market” and consistency across diverse industries is unhinged and AI has equipped the team with the tools to respond quickly and promote dialog between designers and process owners.”

“Reconstruction, not generation. We use AI to deconstruct existing e-learning and work instructions and rebuild them into short, context-driven learning steps that people apply directly on the work floor. This has significantly shortened the time it takes us to deliver a first working programme. We deliberately do not use AI to generate more courses. More content does not solve a transfer problem.” – Remco van den Hout

“We try to provide detailed prompts, however AI appears not to understand our audience/learners so we always need to review our content with our learners and their context in mind. Matching our brand voice and tone is also a challenge when editing AI output. We still need to add personal insights and real-life examples to our content. We also still need to simplify complex language and remove repetitive phrases from AI content. Fact checking, reviewing statistics and verifying claims made in the content remain human tasks currently.”



CONTENT DEVELOPMENT

Text-based content creation

- ▶ **Writing copy:** Generating course and e-learning content, training materials, documents and supporting resources, sometimes from organisational data, SME conversations or other source material. Several respondents explicitly describe AI as providing the starting point that is then reviewed and developed further by a human.
- ▶ **Learning resources:** Creating facilitator and participant materials, workbooks, job aids, guides, checklists, playbooks, knowledge articles, process documentation, observation checklists and other performance-support assets.
- ▶ **Scripts:** Generating scripts for learning content, including video and voice-over scripts, dialogue for animation and VR, and other scripted learning experiences.

Content transformation

- ▶ **Copy editing:** Proofreading, rewriting and improving existing text for grammar, clarity, concision, flow and consistency, as well as strengthening wording and incorporating feedback into subsequent versions.
- ▶ **Audience adaptation:** Simplifying or tailoring content for different audiences, roles, levels of knowledge, and learning objectives.
- ▶ **Reformatting:** Turning source material into more structured content, reorganising text, changing formats and repackaging existing material – from converting policies, procedures and SME knowledge into microlearning to rebuilding existing e-learning or work instructions into shorter, context-driven learning steps. Also, condensing documents, transcripts and other source material into summaries, follow-up materials, checklists or more digestible learning content.
- ▶ **Translation and localisation:** Translating learning materials into other languages and localising content.

Quality assurance

- ▶ **Quality review:** Reviewing learning content for gaps, repetition, clarity, consistency, cohesion, completeness and overall quality, and providing critique or suggestions for improvement.
- ▶ **Checking against standards:** Reviewing content against the original specification, internal rubrics, format and process standards, brand and style guides, accessibility requirements and inclusion criteria.
- ▶ **Accuracy and validation:** Checking accuracy, comparing content with reference materials, validating claims and assumptions, and pressure-testing or double-checking completed work.
- ▶ **Compliance and risk checking:** Checking areas such as copyright and regulatory or organisational requirements where these form part of the content specification.

“Generation of tools that allow us to work faster, check things, and pressure test new content. Never for actually creating content, but for everything around it.”

Video

- ▶ **Avatar-led video:** Creating training and explainer videos using synthetic avatars and AI-generated voices.
- ▶ **Interactive video:** Using AI-generated video within role plays and branching learning experiences. This is a more specific use and appears less often than straightforward video creation.
- ▶ **Creative video:** Experimenting with more generative formats, including creating video from photos and AI-generated films.
- ▶ **Video production and editing:** Using AI to support video editing and generate images or graphics for use within videos.

Audio

- ▶ **Voice-over and narration:** Generating AI voice-overs and narration for e-learning and video using text-to-speech, which saves time by removing the need to record and edit narration.
- ▶ **Audio content:** Creating podcasts and other audio-based learning content, including turning documents into podcasts and, in one case, even experimenting with generated songs as learning content.
- ▶ **Transcription and subtitling:** Transcribing audio and generating subtitles or captions alongside narration and other media production.

“Voiceover has been disruptive in the sense we used to use professional human voiceover artists and now use almost 100% AI. Clients are asking for more voiceover, especially for accessibility.”

Visual content

- ▶ **Image generation:** Creating images, illustrations, graphics, infographics, backgrounds, flyers and other visual assets for learning materials, videos and presentations, including visual identities for particular courses. AI is especially useful when suitable stock imagery cannot be found, and for manipulating or cleaning up existing illustrations to produce a more precise result.
- ▶ **Presentations:** Developing slide decks and presentations, from generating structures and visual designs to producing complete slide decks. One respondent also described using AI-assisted coding to build HTML-based course slides.

Format conversion

- ▶ **Converting existing content into new formats:** Moving learning materials from paper, Word documents, PDFs and slides into formats such as interactive modules, podcasts, or video.
- ▶ **Diversifying learning materials:** Reworking the same underlying material into visual, written, audio or interactive formats to give learners different ways to engage with it.



“While [ChatGPT and Claude] are helpful for generating presentation structures and ideas, I still prefer to create or significantly refine the visual design myself, as I find AI-generated visuals are not yet at the quality I need.”

“AI is good for visuals & videos within limited use cases. It is also good for phrasing & structure suggestions. However, the more the LLMs get trained on previously AI produced slop, the more the utility for anything more complex is degrading.”

– Maura O'Toole

“Manipulating or creating new illustrations – very handy for avoiding copyright issues and getting precisely the desired image.”

“Some people [are using AI] to create PowerPoints but I don't like these as they lead to formulaic presentations, and encourage didactic teaching.”





“We had added AI activities into program design but recently made the decision to remove some of it as we realised that the LEARNING was in the thinking and struggle to apply the learning, by encouraging use of AI we were in fact removing the deep processing or need for the learner to truly learn and apply the skill.” – Suzanna Horsfall

“The LXP has AI simulations but these are mostly hard-coded and not as personalised as they first appear. We use AI Claude skills to expedite simulation creations for this but they then require a developer to build them in the platform.”

“AI coaching in the eLearning course material with longitudinal, individualised, adaptive feedback using database storage as persistent memory for the AI coaches.” – Lawrence Jacobson



USER-FACING DELIVERY

Conversational bots

- ▶ **Coaching and mentoring:** Providing always-on AI coaches or mentors that learners can converse with for reflection, development and support, including leadership, career, performance consulting and goal-setting contexts. Some respondents describe coaching embedded within onboarding academies, e-learning or the flow of work, rather than as a standalone intervention.
- ▶ **Role-play and simulations:** Giving learners opportunities to practise conversations and interpersonal skills with AI, including customer and client conversations, interview skills and other soft-skills scenarios. Respondents also mention using conversational avatars and AI bots to create more realistic simulated workplace interactions and to enable role-play practice at scale.
- ▶ **Tutoring and study support:** Using AI as a tutor, study buddy or conversational learning companion that learners can question, learn with and receive support from as they work through material.

“AI agents being developed to act as a learning coach supporting and coaching our workforce to more habitual learning, applying learning in workflows, and feeding performance/developmental data to and from our HRIS system feeding into performance reviews with line managers.”
– Kelly Rodrigues


Dynamic feedback

- ▶ **Dynamic feedback:** Providing AI-generated feedback that adapts to user input instead of relying on preset responses, for example by responding to open-ended answers in e-learning, reviewing submitted course assignments, or giving personalised feedback on learner work.

Personalised recommendations

- ▶ **Personalised pathways:** Creating learning paths and content recommendations based either on learner data held within the LMS/LXP – such as role, skills, career goals or capability requirements – or on needs expressed directly by the learner using a conversational AI interface.
- ▶ **Adaptive learning:** Adjusting the learning experience around the individual over time, including skill-based adaptive journeys and personalised content layered over existing content libraries for learning in the flow of work.

“We develop and deploy AI-powered coaching agents that support learners in real work situations by providing guidance, contextual knowledge, and performance support. This helps accelerate onboarding and skill acquisition while reducing reliance on subject matter experts.” – Nadjib Salah



KNOWLEDGE AND POINT-OF-NEED SUPPORT

- ▶ **Content sourcing and curation:** Finding credible internal and external sources, locating resources quickly across large collections, and curating existing material into more coherent libraries, programmes or learning paths. AI is also used to add references, create bibliographies, and tag or map content by topic or skill.
- ▶ **Knowledge base construction and maintenance:** Capturing and structuring organisational knowledge from SMEs, policies, procedures, projects and operational experience, and turning it into searchable repositories, FAQs, tutorials, onboarding resources and learning libraries.

“Advanced information retrieval through agent functionality: through Copilot’s agent-based search capabilities, I can target specific information sources (for example, intranet pages) and analyze content more accurately. For example, this has enabled me to identify content gaps and provide concrete improvement recommendations for [...] intranet pages, something that is difficult to achieve to the same extent in a standard chat experience where source selection cannot be controlled as precisely.”

- ▶ **Knowledge search and retrieval:** Helping people find information, expertise or relevant learning through natural-language search rather than manually navigating multiple resources. Examples include deeper search across content, retrieving validated information from knowledge bases, searching LXP catalogues, and surfacing summaries or suggestions alongside results.
- ▶ **Point-of-need assistants:** Using AI assistants to provide contextual guidance and performance support at the moment of need. For example, answering procedural or L&D questions, helping people use unfamiliar tools, supporting onboarding, or providing access to knowledge directly within environments such as Teams/Slack.



“Capturing tacit knowledge. Experienced colleagues give feedback to new hires inside our system, in free text. We use AI to cluster and structure that feedback into concrete improvements to the programme, so the knowledge in people’s heads gets captured while onboarding is running instead of in a separate documentation project.” – Remco van den Hout

“AI supports the creation and maintenance of onboarding resources, learning libraries, tutorials, FAQs, and searchable knowledge repositories. This helps employees find information more quickly and enables more self-directed learning.”





“We [...] use AI to identify, map, and analyze organizational skills and expertise. Through pilot projects, we document the competencies required to successfully integrate AI into business processes, helping us better understand workforce capabilities and future skill needs.” – Nadjib Salah

“[We] assess capability and proficiency through pre-assessments, checkpoints and summative measures. Identify skills gaps and target the most effective development interventions using evidence-based insight. [This results in] clearer workforce capability data, more targeted investment and measurable proficiency gains.” – James Woodward

“AI is heavily used for employee skills analysis, skills demand analysis, and the management of enablement pathways for employees.” – Franz Obermeier



SKILLS INTELLIGENCE

“Large scale JD [job description] review, to work with AI in how roles will work with AI, then skill mapping into future proofing workforce skills and requirements in the age of AI, then using AI to upgrade our skills and competency frameworks, and use AI to tag relevant learning/content/job aid to help close the gap for people.” – Kelly Rodrigues

- ▶ **Skills architecture:** Creating and updating skills and competency frameworks, taxonomies and matrices; defining competencies and observable behaviours; and mapping skills to frameworks. These structures are then used to support career pathways, training pathways and other workforce development decisions.
- ▶ **Current capability assessment:** Assessing current workforce skills and proficiency, building employee skill profiles, identifying skills and capability gaps, and analysing organisational expertise. Respondents mention self- and manager assessments, competency assessments and skills-verification tools, as well as using AI to infer skills from available workforce data and signals.

*“We use AI [...] e.g. coming up with drafts for skill profiles to be reviewed by specialists and final versions to be used to map relevant learning content from our academic catalogs.”
– Jeff Kortenbosch*

- ▶ **Future skills and workforce planning:** Identifying emerging and future capability requirements, analysing skills demand, and comparing current workforce capability with what the organisation will need next. This work supports workforce planning, future-proofing workforce skills, understanding the competencies required for AI-enabled work, and making sense of emerging capability needs.



EVALUATION AND ANALYTICS

Learning analytics

- ▶ **Data analysis and insight generation:** Combining and analysing data from learning platforms and other organisational sources to identify patterns, trends and relationships and generate insights for decision-making. Data mentioned by respondents includes participation, course uptake, learner progress and engagement, HR data, employee surveys, proficiency and confidence measures, and other programme KPIs.
- ▶ **Impact and ROI:** Connecting learning data with business outcomes such as cost, productivity, safety incidents, retention and time to competence; comparing results before and after the introduction of new tools; and using AI to support ROI calculations.
- ▶ **Learner and programme performance:** Analysing learner progress, engagement and course or programme outcomes to identify where participation drops, progress stalls or particular cohorts are struggling. Respondents also mention tracking capability growth through proficiency and confidence assessments rather than relying on attendance alone.
- ▶ **Feedback and evaluation analysis:** Analysing qualitative and quantitative learner feedback; extracting themes and patterns from free-text responses; identifying challenges; summarising comments; and assessing sentiment. This includes post-learning evaluation data such as satisfaction scores and feedback on the application of skills, as well as user-testing feedback used to improve learning experiences.

Reporting and dashboards

- ▶ **Reporting:** Turning analysis and evidence into reports, including post-training and client impact reports, and using insights from KPIs and programme data to make findings easier to interpret and support operational and strategic decision-making.
- ▶ **Dashboards and visualisation:** Creating dashboards, indicators and data visualisations, sometimes by combining data from multiple systems, to make learning metrics and trends easier to monitor.

Evaluation design

- ▶ **Surveys and questionnaires:** Generating surveys, questionnaires and other evaluation instruments for learning and organisational development.
- ▶ **Evaluation frameworks:** Creating evaluation plans, criteria and measurement approaches that define what should be measured and how.

Marking and assessment

- ▶ **Grading:** Reviewing and grading assignments, projects, capstones, certification work and free-text responses, including automatic question grading and preliminary marking against assessment criteria. Some respondents use AI for an initial assessment that is then checked by a human, while others highlight its potential to improve consistency across project and certification grading.

“We are using it for analysis, seeing which courses people in the organization took, why, and how much it cost.”



“We are using AI to analyse data and generate insights that feed into OD frameworks, organisation design, learning solution design, etc” – Drew McGuire

“Claude is [...] my preferred tool for analyzing larger volumes of information or identifying patterns across documents and historical data, particularly when I need deeper analysis rather than content generation.”





“We are using it to review notes from meetings with clients and to test what we have gathered against our working hypothesis. We are also using it during discovery meetings to record and transcribe. We can then analyse key trends emerging from the discovery meetings and compare trends that emerge from different groups of stakeholders, e.g. what do operations believe vs HR vs the Executive Team.”

“Beyond content creation, we are using AI to improve proposal development and solution design by helping estimate effort, identify risks and recommend delivery approaches.”

– Jenny Randau Mellander



ADMINISTRATION & COMMUNICATIONS

Routine administration

- ▶ **Email management:** Drafting emails, summarising information, prioritising work from emails, and otherwise helping manage high-volume mailboxes.
- ▶ **Meeting administration:** Preparing for meetings, note-taking and producing follow-up summaries and action items.
- ▶ **Coordination:** Supporting scheduling, RSVP management, conference logistics, and participant allocation to cohorts or groups.
- ▶ **LMS administration:** Streamlining recurring L&D administration and LMS tasks, such as reporting, documentation, writing course descriptions and producing other shorter pieces of copy that sit around the core learning content.

Stakeholder communications

- ▶ **Alignment:** Creating stakeholder communications, such as emails, executive briefs and presentations.
- ▶ **Proposals and business cases:** Developing business cases to get buy-in for proposed investments or approaches, sharing recommendations with business leaders, and drafting client proposals and statements of work (SOWs).
- ▶ **Vendor-side communications:** Creating client-facing sales materials, including sales messages and briefs used to support client conversations.

Internal marketing

- ▶ **Communication strategy:** Supporting programme launches and structuring communication plans.
- ▶ **Promotion assets:** Creating employee-facing communications and promotional materials, including emails, invitations, information sheets, learning agreements, intranet content, course blurbs, newsletters, social posts, website content, podcast blurbs and blogs.

Market intelligence

- ▶ **Trend monitoring:** Tracking developments in learning and AI, analysing L&D trends and producing recurring sector updates that highlight changes and implications for the learning function.
- ▶ **Market and best-practice research:** Researching the market and industry best practice, including benchmarking how other organisations approach similar challenges.
- ▶ **Vendor scanning:** Researching and evaluating training providers and learning technology vendors.

“I rely on it for financial management tasks, from tracking billing cycles to consolidating project-level data.” – Fernanda Mendes



OPERATIONS

“We [...] use AI to help automate our training intake process when new requests are submitted. The tool helps us facilitate an initial decision to qualify [if] training is the best solution for the request.” – Kristal Walker

- ▶ **L&D strategy:** Developing organisation-wide learning strategy and using AI as a critical thinking partner to test assumptions and shape strategic direction.
- ▶ **Project and work management:** Supporting planning, prioritisation and delivery by creating action plans, tracking capacity and project progress, estimating effort, identifying risks and preparing launches.
- ▶ **Workflow automation:** Automating repetitive L&D processes and connecting steps across workflows, such as emails and reporting, customer reporting, learning logistics, certification, training-request intake and triage, project updates, content-development workflows, module creation from storyboards and integrations with learning platforms or other systems.
- ▶ **Internal tools:** Building task-specific assistants and lightweight applications through vibe coding or AI-assisted coding to speed up work or enable self-service. Examples include tools that help managers identify development or performance gaps, allow business teams to design learning experiences, support personalised content, analyse data across the tech stack, support managers and HR with decision-making, provide guided chatbot experiences, and power platforms or agents used in client delivery.

“AI supports our own developers in building our software. Assisted programming with human review, not vibe coding.”



"I built a semi-agentic triage and prioritisation system using our in-house LLM to determine content relevance and priority."

– Ash Wedderburn

"The benefit of using Claude Code and Vercel to build my eLearning, instead of Articulate Storyline which I have used for close to 15 years, is that the entire course is wired to Claude Code and entire modules can be rebuilt from beginning to end in a few hours instead of a few weeks. Changes can also be propagated system wide without any disruption to already published courses." – Lawrence Jacobson

"AI agents [are] supporting the full content development process, AI-enabled learning experiences (simulations, personalized feedback, etc.), custom chatbots / GPTs, content maintenance and updates, and developing a tool for automated, personalized content on top of a content library for learning in the flow of work." – Jolien Hartveld





WHAT ABOUT AGENTS?

In last year's survey there were few mentions of either AI agents or agentic AI and of these, only three responses described features that are commonly associated with genuine AI agents: autonomy in decision-making, tool use and workflow orchestration and integration across environments.

This year, 41 respondents used the word 'agent', although again with varying degrees of accuracy. In particular, many respondents seem to use the words 'agent' and 'chatbot' interchangeably and describe something closer to a custom GPT than software that plans and carries out a sequence of tasks by itself. A significant number of mentions of agents are also too brief to confirm whether they are, in fact, agentic. Many are also about work in progress. Of all 71 uses of the word 'agent' across the four years, 24 talk of building, creating or developing, and 16 of experimenting, piloting or hoping.

TWO DIRECTIONS

Where respondents do say what their agents are for, they point in two directions. Some agents serve the L&D team's own work: drafting designs, checking quality and triaging requests. Others face employees directly, answering questions, supporting people as they work or coaching them. Coaching is the most common single purpose.

What both have in common is specificity. These are not general-purpose AI assistants. Respondents describe agents with defined jobs, often narrow enough to be given their own names. One described three agents built for their organisation:

"LEN: faster learning discovery. ALI: faster, better learning design. DAN: better writing and stronger drafting capability."

One respondent in 2025 mentioned a "strategy agent" and an "implementation agent". Another described "Companions" for feedback and coaching. The names are incidental, but they reveal how respondents are thinking about agents: less as tools that can do anything and more as colleagues given a particular role.



That role is then supported with organisational context. One respondent described an agent that had been given the foundations it needed to work according to the L&D team's own approach:

"It knows our learning strategy, and it knows our learning foundation (learning theories, how we like to design, about how to be a trusted learning advisor etc.)." – from 2025

Others connect their agents to SharePoint, Slack or an organisational database to give them access to current organisational knowledge. The pattern is a move away from prompting a general-purpose tool each time towards building an agent that already knows what its job is, how the organisation works and what good looks like there.

A FEW WORKING SYSTEMS

Only a few respondents describe agents linked together into a workflow of several steps. These come closest to the technical meaning of the word, and they report the largest effects. Only one gives a measured result:


"In documentation and training, we have an end-to-end workflow using agents at every pain point. Some have seen a verified reduction in cycle time of more than 90%."

Some of these systems reach beyond L&D. One respondent described a coaching agent that is:

"feeding performance/developmental data to and from our HRIS [HR information system] system feeding into performance reviews with line managers." – Kelly Rodrigues

Not every build succeeds. One UK respondent, whose organisation uses AI extensively, was frank about the limits:

"Other things like knowledge and pursuit advice agents are struggling to find proper, measurable adoption even though they remain relatively easy to do."



WHO BUILDS AGENTS

Mentions of agents are linked to more advanced use of AI. In 2026, 16% of respondents who chose 'Using AI extensively' mentioned building a chatbot or agent. Among those who chose 'Using AI in some parts of L&D', the figure was 7%. The difference is statistically significant, though modest.

Where respondents work makes little difference. Among those in workplace L&D teams, 11% mentioned agents; among consultants and freelancers, 10%.

What does vary is who the agent is for. In-house teams describe agents for their own colleagues. Some consultants describe agents as part of what they sell:

"Development of platforms and agents for our use in providing services to our clients."

STILL AN UNCLEAR PICTURE

In 2026, the share of respondents claiming to build their own agents increased sharply. These claims, however, should be treated with caution. Respondents, as has been noted, are loose in their use of terminology, and most of what they describe is a single assistant, given a name and some knowledge of the organisation. A small number have linked agents into workflows, and they report the biggest gains in speed. Overall, despite the technical agentic AI breakthroughs in the past year and the hype surrounding it, L&D seems to have experienced an uptick in activity, but not yet anything approaching an explosion in agentic AI use.



“

“AI adoption is no longer happening on a level playing field. Some organisations now have the infrastructure, expertise and strategic backing to build complex AI-enabled systems at scale. Others are still working out how to move beyond individual experimentation and relatively simple use cases.”

”

CASE STUDIES

In this section of the report, you'll find case studies that take a closer look at AI in action across L&D. They span a wide range of approaches: from highly sophisticated, data- and technology-intensive implementations to much more accessible examples built with tools already available to the learning team. Together, they show not one version of what "good" AI use looks like, but a spectrum of possibilities depending on where an organisation is starting from.

That matters because AI adoption is no longer happening on a level playing field. Some organisations now have the infrastructure, expertise and strategic backing to build complex AI-enabled systems at scale. Others are still working out how to move beyond individual experimentation and relatively simple use cases. A case study that feels ambitious or out of reach in one organisation may feel entirely achievable in another; equally, an example that looks straightforward to an advanced team may offer exactly the right next step elsewhere. The value of these stories lies not only in what was built, but in helping teams see what is possible from where they are now.

To gather the case studies in this report, we reached out to learning leaders after hearing about their work. We want to thank them for sharing not only what they built, but how they approached it, what they learned and what others might take from their experience.

We hope you find the case studies useful – whether as a glimpse of what may be possible further ahead, or as a practical example of something you could begin trying now.



INTERNATIONAL AVIATION TECHNICAL COLLEGE RIYADH

The Agentic Teammate

Jinani Muhammad, AI and Digital Training Lead

CONTEXT AND JOURNEY

International Aviation Technical College (IATC) at Riyadh, operated by ARX, is a vocational college that trains aviation technicians. ARX also delivers regulatory compliance training more broadly. Jinani Muhammad's role has grown with the company's needs: over the past few years it has become deliberately generalist, spanning design, systems and data, as the business has needed to cover more ground and agent capability has let him take on more of the work.

Last year, Jinani tried something similar to what he is developing now: an early concept for a company agent that could act on requests semi-independently. It was useful within a narrow scope, but the technology of the time was too brittle to extend much beyond it. Day-to-day work involved too much variation for the system to adapt reliably.

Once agentic AI matured earlier this year, Jinani's work shifted from building individual tools to building something closer to a colleague: a system, built on Claude Code, that stays current on the company's projects, understands how the business works, knows who's who, holds the brand, and compounds what it learns over time rather than starting from zero on every new task. This marks a

shift from individual tools and automations, each built for one problem, to a shared agent 'brain' that can draw on the same context and abilities as different problems arise.

The agent in practice

In practice, one agent system is supporting a growing range of work. It:

- ▶ prepares monthly expense claims by reading the bank statement, classifying transactions, gathering the invoices and completing the claim form, with separate checks before submission;
- ▶ drafts on-brand, client-facing proposals and reports that are fit for purpose and ready for human review, with minimal rework;
- ▶ supports business development meetings from preparation through follow-up, including tasks and documentation in ClickUp, the project management platform;
- ▶ has audited and helped strengthen the AI-Assisted Assessment Writing System, an AI-powered aviation exam and practice-question system built before the agent itself, and produced two manuals to support handover: a user manual for the curriculum team and a detailed technical manual for developers maintaining, troubleshooting or extending the system;

These four are examples, not the complete list. Many more can be developed as the business needs them. The agent keeps its knowledge, context, abilities and understanding of the work in one core system, including what each task relates to and who it is for. Connections to tools such as ClickUp, OneDrive and email let it carry that understanding into the places where work happens. The connection provides access; the company's conventions teach the agent how to work in each place, including what belongs there, how it should be recorded and when it must ask for guidance or approval.

Last year, an AI tool might have helped produce a single image or document. Now the agent can support a business development workflow from preparation through follow-up. Beforehand, the agent prepares using the project record and helps shape the agenda. Afterwards, it classifies the discussion, updates the relevant ClickUp tasks and records, adds any new people or context, and produces minutes using the company template. If the discussion leads to a proposal or presentation, the agent conducts an interview to define the brief and prepares the branded materials for human review, drawing from its catalogue of approved assets where needed.

The workflow is not handed over all at once. Each stage is guided and corrected until the agent can carry more of it reliably. In Jinani's words: "The agent is co-producing the work, as opposed to the human requesting isolated input from AI tools."

The better the agent gets, the better its output gets – but 'better' here isn't only about the work itself. Part of every session is the agent reflecting on what it got right and wrong, working out what's worth remembering, and checking with Jinani before folding any of it permanently back

into how it operates. As he puts it: "Just as I help build people's capabilities, I now find myself building an agent's capabilities."

Continuous learning

Building a custom agent for a company isn't like building a narrow, single-purpose AI tool. It's a coordination layer: the part that keeps every piece of what the agent does – checking its own work, updating a task, preparing an email for approval – working together as one process, instead of separate steps someone has to manage by hand. Jinani approached that by breaking key operations down into 'bundles of work' one at a time.

A bundle of work is a repeatable workflow made up of related tasks. It might produce a report or a set of pitch decks, or support an ongoing business development cycle. The agent is taught to do it by trial and correction: shown examples, and given feedback on specific mistakes as they come up. Jinani has logged over a million words of dictated correction and context so far.

Because the agent keeps notes on what worked and what didn't during each session, it doesn't need the same context rebuilt from scratch each time. That improvement continues until the agent can carry more of the work reliably under appropriate human oversight. At that point it 'graduates' from learner to teammate.

The goal underneath all this was never the deliverables themselves. The focus has been on the **ability** to do something, not the **tool** to do it: a tool solves the problem in front of it, while an ability – knowing how to analyse a client brief, check its own work, or hold a conversation with someone senior appropriately – can be redeployed wherever that skill is needed, which is why the same agent that started on invoicing can also help audit,

document and strengthen an AI-powered aviation assessment-writing system.

That design philosophy shows up most clearly in how the system checks its own work. A 'tool' would run the same check every time, no matter the stakes; an 'ability' adapts – which is why verification here is scaled to what's at stake, not applied uniformly. Low-stakes work gets a self-check. Higher-stakes work may receive a panel of review agents or an independent check from a different AI model, such as Codex, which can catch things a same-model self-check might miss. For the most consequential work, reviewer panels may be asked to challenge the work rigorously in the way a key stakeholder or regulator might.

There has also been a shift in what counts as 'AI work'. Visible tools and internal products, which have traditionally been seen as AI work, are still being built. However, they increasingly feel secondary to the agent itself because they emerge from agent-assisted work rather than being its main objective.

IMPACT

Some of what's changed is straightforwardly measurable: proposals and pitch decks that previously took several days can now be completed within a day.

Another part is entirely new work, made possible by capabilities that simply didn't exist before:

- ▶ **A graduate career fair portal**, connecting 800 graduates with 17 aviation companies and processing more than 11,000 applications, designed, built and run inside a week or two. Before this year, a digital system like that wouldn't have existed at all, given the time it would take to build it.

- ▶ **Consistent follow-through on meetings and projects.** What used to be ad hoc and easily dropped is now tracked reliably as standard practice.
- ▶ **A higher quality bar for design assets and documentation**, without needing a dedicated designer for every deadline, since on-brand, well-produced material is now the default rather than a bottleneck.
- ▶ **More rigorous work, not just faster work.** The agent can surface relevant frameworks and standards, then hold the work to the ones the team chooses. It flags gaps such as a missing needs analysis, a skipped problem statement or a step that was claimed but not followed. This may require more effort, but it helps ensure that a method being claimed has actually been applied.

Beyond these demonstrated gains, the emerging model starts to resemble what some in the field call an AI 'factory': a system for building the reusable processes that produce the company's work. The wider opportunity is to make work viable that would previously have taken too long or too much effort to build from scratch.

CHALLENGES AND LESSONS LEARNT

Transformation before adoption. The instinct in many organisations is to hand out AI tool access and hope capability follows. Jinani's approach inverts that: transform the work properly first, one bundle at a time, and let adoption follow once there is something worth adopting, beginning with the people involved in that work. Rolling out access before the underlying work is trustworthy just imports human inconsistency into a system that is not ready to handle it.

“The more we invest in the agent’s abilities, the more it becomes a high-performing asset for the company. The shift is from creating a tool for one solution to developing an agent that can contribute to many different solutions as they arise.”

Trust is earned through visible output, not explained. Building an invisible layer underneath the business is genuinely difficult to sell to stakeholders, because there’s rarely anything to point at along the way. The approach here has been to let deliverables – proposals, reports, the career fair portal – arrive faster and better over time and let that pattern speak for itself, while deliberately avoiding both overselling (“it’s a superhero agent”) and underselling (“it’s just ChatGPT”) what it actually is.

The agent is a smart junior, not a replacement. The brief gets set at the front of every piece of work, and the output evaluated at the back; the agent proposes, verifies its own assumptions, and builds in between. That structure is built to hold permanently, no matter how good the agent gets.

Judgement-led work is much harder. Agents are more reliable where success can be checked objectively. When building a digital system, for example, much of the work can be tested against clear rules: does it run, does it produce the expected result, and does it pass its checks? Company-specific training, proposals and media are harder because quality depends more heavily on judgement, taste and context, and the company may not yet have written down what good looks like. This makes the work less reliable and requires much closer human involvement, both to define the standard and to judge the result. In more testable areas, the agent can carry much more of the work between the human-set brief and the final human decision.

The last mile doesn’t end. The technology is a moving target, and better models may validate some of today’s scaffolding while making other parts unnecessary. Each capability therefore has to keep being tested and refined rather than assumed to work. Even with heavy scaffolding in place, the agent still gets things wrong regularly, and closing the gap between good enough and genuinely trustworthy is slow, unglamorous work rather than a one-off fix. Getting from 60% to 90% quality is “excruciating” and new situations continue to surface gaps that need to be caught and folded back in.

NEXT STEPS

The work continues as a cycle. New bundles are tested on real work and refined through use, while existing capabilities continue to adapt as the work and technology change. The focus remains on areas where reliable execution can create the most value.

One major area that is now being developed is the instructional design system that underpins ARX’s training offerings, from analysis through design and implementation. This moves beyond a single bundle of work. Abilities are applied to individual tasks, which are joined into repeatable bundles. Several bundles can then support an entire business function. It draws on capabilities already developed for retaining context, coordinating agents, verifying work, applying the brand, producing documents, images and video, and managing projects. The aim is to test how far those capabilities can strengthen the function, with human judgement retained at the key decisions.

CASE STUDY 2



EY

A Sandbox for Agentic AI

Simon Brown, Global Learning & Development Leader and Partner, Talent

CONTEXT AND JOURNEY

EY is a multinational professional services organization with almost 400,000 employees in over 150 countries, advising clients on Assurance, Consulting, Strategy, Transactions, and Tax. Building leading-edge skills for people at the organisation is a strategic imperative, and AI integration into its operations is critical to EY's work with clients and leadership in this space.

Agentic AI is one part of that shift. It has the potential to change how the business operates, so understanding it now is groundwork for the strategic decisions still to come. To give people the infrastructure to explore agentic AI and see what it can do, EY has made a growing suite of tools available as part of its approved technology stack.

“The best way to answer where agents can play in the learning scope is to test that rather than just hypothesise.”

Around six months ago, Simon Brown, EY's Global Learning & Development Leader and Partner, put that infrastructure to use directly, setting out to answer a specific question: what can agents do within the scope of learning? The only way to find out, he concluded, was to start building and testing.

The Agentic Sandbox

Simon built an agent team of more than 35 personas. What resulted wasn't a handful of individual agents, but a full custom application that houses and runs the whole agent team, without him writing a line of code.

Each agent persona is built around a distinct role: a learning designer, a labour economist, a futurist, an executive coach, a DevOps and release engineer, among others, each with its own name, personality and area of expertise.

Simon puts the agent team to work in several ways. One use is running full workflows: multi-step sequences where several agents work through a task together, each handling its own part before passing to the next. For example, one workflow took a single brief and moved it through a research agent, an instructional designer, a writer, EY's brand agent, a graphics designer, a QA and accessibility auditor, a prototype developer and a tester, before coming out the other end as a finished, interactive draft course, with no further input from Simon along the way.

Another common use is as a thought partnering tool: pressure-testing ideas, preparing for client conversations, and running roundtables where several agent personas debate a brief and challenge each other's thinking. In one case, a roundtable changed a real decision: reviewing a list of external projects EY might get involved in, a deliberately contrarian “agitator” agent persona argued for an option

nobody had seriously considered – one that made sense because of work already underway elsewhere in the business. It turned out to be the better call.

Simon's team has its own access too, running related experiments in parallel and testing things like synthetic data generation or agent retraining from course outputs.

How it works

Inside the application itself, building an individual agent starts with an identity: a name, a role, a goal, a personality, and the underlying model it runs on. From there, it can be given tools such as web browsing or audio and video generation, a defined methodology, and a set of traits, skills and behaviours. An instructional design agent, for instance, was trained directly on EY's own accessibility design guide, alongside worked examples of good and bad learning content. As an agent runs, it keeps memory across sessions, takes feedback from people and from other agents it works with, and flags when new feedback conflicts with its existing training, so a person can resolve the conflict instead of having it get buried. New agents mostly build themselves: an "AI enhance" step fills in personality, methodology, traits, skills and behaviours automatically, in about 30 seconds.

Individual agents combine in two ways. Workflows chain several agents into a defined sequence, with optional human checkpoints between steps – the eight-agent course-build example described earlier is one instance. The research agent alone completed 62 activities, running iterative rounds of search and gap analysis across government, industry-standards and AI vendor documentation, drawing on sources including the EU AI Act. An agent manager reviewed each step's output in turn, sending the research back

for revision twice before approving it. The whole run took 106 minutes and produced a genuinely interactive draft course, simple but functional.

Roundtables work differently: several agent personas debate a brief directly and converge on a shared view. For instance, given a brief on how a learning team should approach AI upskilling, five personas debated for a few minutes before their combined view was turned, with one further prompt, into a fully branded visual deck.

Underneath it all sits a self-contained knowledge wiki that grows as agents do research. At one point it held more than 1,200 pages, entirely agent-written and sourced, standalone from any of EY's actual organisational data. All agents also draw on a shared taxonomy of agent skills: a common library of capabilities, each with its own instructions, methodology, tone and guardrails, so every agent works from the same underlying standards rather than reinventing them individually.

"Rather than the black box of putting in a prompt and getting something out, when you actually see the agents interacting – asking each other questions, challenging each other, iterating on their work – that's where the penny drops for a lot of people."

IMPACT

Thought-partnering with the agent team has become a genuine time-saver for Simon personally: running a question through a roundtable of personas, rather than working through it alone, typically saves him half an hour to an hour he'd otherwise spend thinking it through by himself.

Some of Simon's team have also adopted the new tools and started running their own independent experiments alongside his.

Finally, the same agent team is used directly in EY's AI training sessions for senior audiences. Watching agents interact and challenge each other lands differently than typing a single prompt and reading an answer – it's often where people first genuinely grasp what agentic AI means in practice. It's shown to clients too, as a live example of EY's own thinking on where AI in learning is heading.

CHALLENGES AND LESSONS LEARNT

Deliberately built-in challenge sharpens real thinking, not just output. Adding a persona built to argue the contrarian view, rather than only ones built to agree and execute, has repeatedly surfaced better thinking than working through a decision alone would have.

Leader role-modelling builds more than a good example. Simon showing what he's doing and sharing it openly creates permission for the team to experiment themselves. But it's also building the shared reference point the business will need to make real strategic decisions about agentic AI.

Build for a moving target, not today's baseline. Comparing the same workflow run six months apart shows how much the underlying models themselves have improved in a short time. Anything built now should expect the ground to keep moving underneath it.

NEXT STEPS

The next step is figuring out how the team works once agents can handle more of the common tasks. That might mean training and maintaining the agents doing the work. It could also mean something more concrete: with agents doing more of the production itself, the team gets the time it never had before to work out what's actually landing with learners, and to change what they build because of it. Months of production never left much room for that kind of attention.

CASE STUDY 3



RANDSTAD ENTERPRISE

A DIY Front Door for Learning

Suzanna Horsfall, Senior Director, Global Learning

CONTEXT AND JOURNEY

Randstad Enterprise, part of Randstad N.V., employs 5,000 people globally, delivering talent solutions for some of the world's largest companies. That makes for a unique internal demographic for the industry including: client-facing delivery staff, regional sourcing hubs, and a wide range of specialist positions.

Much of the internal population is embedded day to day within client systems, at a distance from Randstad Enterprise's own channels, meaning that learning content was scattered between an intranet and external platforms. An internal survey found the population split in two: those who knew what was available, and those who didn't. Given the scale of future upskilling forecast, the company saw the opportunity to elevate its learning strategy and ensure a more consistent and unified approach.

Addressing this meant shifting ownership and accountability for learning from a small central L&D team onto operational managers themselves. It also meant moving away from a 'top-down' cascade, toward something managers could drive directly with their own people.

However, managers don't always have the practice to diagnose a potential performance gap in someone and connect it to the right intervention. Being a largely remote working population made this more acute: without the everyday visibility of sitting

near their teams, managers sometimes found it challenging to define what support their reports actually needed.

The solution was a Gem: a custom, purpose-built version of Google's Gemini AI assistant that anyone with access to Google's tools can configure for a specific task, giving it its own instructions and knowledge to work from.

Randstad Enterprise already had it available: no new tools to buy, no new infrastructure to build. The Gem is the front door to the learning offering, helping managers work out what capability gap they're trying to address before pointing them to the right resources.

The Gem

The Gem starts by asking managers what they're seeing in someone's performance, and what improvement they would like to see instead. Where a manager has performance data or a clear set of KPIs to hand, it asks for that directly. Where they don't, it narrows in with further questions until the gap becomes specific enough to act on, closer to a guided conversation than a form to fill in.

The Gem gives complete control to the leaders while leading them to the appropriate post-learning coaching options. This means that once the learning solutions are displayed, the Gem asks if the leader would like to coach or run a quick informal discussion with the team member. The Gem is able to give a custom-tailored Coaching / Discussion plan with

clear points that the leader can discuss with the team member. This adds a layer that makes the Gem not just a learning concierge, but an intelligent partner to the leader, helping them better support their team members in applying learning and building team capabilities.

The Gem was built by Raghava Govardhanam Venkata, on Suzanna's learning team, as part of an existing remit that already covered platform learning, leader-led learning, and the organisation's leadership and management programmes. None of what went into it was new. The work was turning what already existed into something the Gem could actually use, refined through real testing and revision.

Building it meant giving the Gem knowledge of everything that had been scattered across systems. That included:

- ▶ **Programme details** – name, target audience, learning objectives, duration and session frequency, and how each is delivered. Randstad Enterprise's formal programmes typically run four to six weeks, blending self-paced digital pathways with facilitated sessions.
- ▶ **Licensed platform content** – what's available through Udemy and Social Talent, including pathways the team had already curated within Social Talent for specific parts of the business, such as the sourcing team or those developing their management skills.
- ▶ **Mentoring** – the platform where people self-match with mentors for reverse, flash, or longer-term mentoring relationships.
- ▶ **Additional resources** – content from programmes no longer actively run, such as repurposed slide decks, facilitator guides, and learning activities.

“If you've identified a problem and think you have a solution for it, then knock yourself out, because we're only going to find new ways of doing things if we give people the space to play.”

Getting the underlying setup right took real iteration: at least three full restarts, going back in and putting different information in each time because earlier versions weren't producing useful suggestions. The diagnostic question flow went through its own round of testing too, tried out with a small group first and refined from what came back, before it went out more widely.

The work on the Gem hasn't stopped since launch. Some of it is upkeep: every time a programme changes or something new gets added, it needs manually updating to match. Some of it is by design: when a manager reports that a suggestion missed the mark, that feedback goes straight back in, because the Gem was never built as a one-off release. It keeps improving with every report.

IMPACT

The clearest signal is that people are taking advantage of more of what's on offer: requests for Udemy and Social Talent licences are up, and so is actual use of them.

Managers are also increasingly leading their own team's learning instead of waiting for L&D to put a programme in front of them. That's showing up as more independent use of what's available – content, curated pathways, coaching – instead of enrolling in a programme by default. The shift is under way, but it's still early days on the journey toward a culture of self-directed learning.

CHALLENGES AND LESSONS LEARNT

Deciding what's worth the time needs a filter, not instinct. You genuinely don't know if an AI experiment is worth the effort until it's out in the world. Suzanna's filter for someone on her team wanting to pursue something speculative comes down to three questions: Does it support learning at Randstad? Is there real time for it? Are they capable of delivering it? If it's yes across the board, her answer is simple: "knock yourself out."

Start before it's ready. With AI moving this fast, there's always another reason an experiment could still feel like it's not ready. That's not a signal to wait, it's just the permanent state of building something genuinely new. It's easy to miss that launching something imperfect isn't the only cost on the table. Not launching at all has a cost too – a team that never builds the AI capability in the first place.

Fluency blinds you to friction. The team's first instinct was to doubt that navigating the intranet was really a problem: if you can use a smartphone, it's not hard. But that judgement came from people who already knew the system inside out. The real difficulty, and the actual opportunity to fix it, only came into view once they considered it from the employee's position instead of their own.

Building on a borrowed platform is a real, unresolved risk. Anything built on a single vendor's feature is at the mercy of that vendor's own roadmap, so you will always need to stay on your toes even with a process that is working.

NEXT STEPS

Randstad Enterprise is moving to a new learning platform, rolled out by country over roughly the next two years rather than all at once, meaning parts of the business will use different systems in the interim.

That timing gives the Gem a second purpose. Built quickly and without much resource, it's now positioned to be the one consistent way to navigate the learning offering while different parts of the business are on different systems during that transition.

"Test and learn is the biggest lesson. Especially with all of this new technology, it feels like every week there is something new that we could be doing. And waiting is not going to make any progress."

CASE STUDY 4



KPMG UK

Dynamic Learning Pathways



Ian Barnwell, Chief Learning Officer

Michael Snowden, Senior Product Manager

CONTEXT AND JOURNEY

KPMG UK, one of the leading professional services firms, employs over 17,000 people in the UK across multiple service lines.

KPMG operates in an industry built on knowledge work and expertise, and one at the forefront of change driven by AI, technology and market volatility. That makes it more important than ever to build and pivot capability faster than before, in order to keep advising clients who are themselves trying to make sense of that same change.

Capability, however, is not built through skills alone. People are at the heart of KPMG's business, and the capabilities of their workforce are one of the most important drivers of success. Human capability is shaped by much more than skills: it is the combination of skills, experience, exposure and relevance, all evolving over time. This has meant the L&D team's focus has been on helping people to grow and adapt, translating business priorities into meaningful development opportunities, and supporting capability building with greater precision, personalisation and scale.

This journey started with Spark, KPMG's AI-powered learning and skills coach, introduced in last year's report, **AI in L&D: The Race for Impact**. Built directly into Teams, where employees already spend much of their day, Spark replaced the old LXP entirely and now acts as the single point of entry for learning.

A year later, Spark is developing to build capability further, adding a more powerful way of personalising learning.

The solution centres on creating bespoke, dynamic learning pathways for each employee. A pathway shows up inside Spark as a sequence of stages building towards a goal someone has set for themselves, whether that's preparing for promotion or moving into a different role. Each stage pairs a targeted set of courses or content assets with a short explanation of why they've been included, followed by a selection of real-world tasks meant to put that knowledge into practice: writing something and getting it reviewed, asking peers for structured feedback, shadowing someone more senior, working through a coaching conversation using a defined model. The point of these tasks is to move someone from simply consuming content towards demonstrating a capability in the real world.

As those tasks get completed, they're logged against the objective they were meant to support, so by the time someone sits down for a formal development conversation, the evidence of what they've actually done is readily available. This closes the loop that began with setting a goal at the beginning of the process.

“We see learning as a capability engine for the firm – not something to look at in isolation, but the mechanism behind how we manage the supply and demand of skills for client engagements. That’s why the end state needs to be connected to the execution layer. We aren’t just developing learning. We’re developing our people so they can support our clients.”
– Ian Barnwell

Dynamic pathways

Unlike pathways curated by hand or by an algorithm matching content to a selected role, KPMG’s dynamic pathways model jobs and tasks at a much higher fidelity than a role label. That means building a picture of where a specific person actually is today – their current role, the tasks it actually involves, the skills needed to do them – set against where their particular role is realistically heading as the business, the market and their own work with clients change. The pathway itself is drawn from the space between those two. The result is not a prescription from the organisation, but a private and personalised development guide for the individual, helping them to make informed decisions about their own learning, growth and future direction.

The system is also deliberately built on more than learning data because putting what someone has learned into practice,

and staffing them onto the right client engagements, do not happen inside a learning platform. It brings together previously disconnected information from KPMG’s internal tools for workforce planning and AI-driven insight into how client work and the market are shifting, presenting it as an additional view to support human judgement.

The same logic extends to assessment. Rather than measuring competence solely against a generic external framework, KPMG’s approach is dynamic and grounded in the firm’s own data, role definitions, performance logic and methodology for what positive performance looks like. These data-led insights support, rather than replace, human judgement, with responsible-AI controls, clear accountability and regular monitoring built into the approach. The underlying information and resulting insights are presented directly to the individual, who decides what to use in their development and performance conversations; the organisation sees only aggregated views of the same data. Access and use are governed through robust data-quality, privacy and security controls. Built this way, assessment can plug directly into pathway data to provide a more granular, timely view of someone’s growth journey and help shape what their pathway shows them next, while keeping the individual in control of how their insights are used.

At its core, this is a people-development service, not a people-management system. The insights are designed to help individuals understand how their skills, experiences and opportunities are evolving, and to identify practical actions that will help them be even more effective in their roles. It is a tool for growth and career development, helping people build the capabilities they need to thrive in a changing world of work,

rather than a mechanism for evaluating, ranking or making decisions about people. Human judgement remains central, with individuals retaining control over how their insights are used and discussed.

Building the pathways

Underneath this is a massive data modelling problem: collecting and modelling enormous amounts of information about roles, skills, content and people, at a scale that people manually curating pathways could never approach. At a high level, it comes down to building a sophisticated data platform, then navigating the best path through it from how people actually engage with it.

Skills function as the common currency that ties the whole system together: content can be described in terms of the skills it delivers, and roles can be described in terms of the skills they require, which gives KPMG a single language it can use to connect otherwise unrelated data.

Modelling someone's role starts with where they are today: their current role, the tasks that role involves, and the skills needed to do them, drawn from a blend of cross-organisational data sets. Against that sits an understanding of where the existing content estate already covers those needs and where the gaps are. Separately, there's a picture of the future state: how each role and the work undertaken is expected to change over time, informed by market insight and by KPMG's own AI workforce and consultancy teams, ensuring an evolutionary path into the future.

The pathway is shaped by the gap between someone's current capabilities and what they need for their role now, and in the future. Spark uses this view to provide a focused development journey that combines relevant learning with different practical ways to apply it in real-world situations,

helping the individual build and demonstrate the capabilities they are likely to need next.

A key part of the approach is establishing a reliable 'ground truth' by reusing relevant historical, human-led work rather than starting from scratch. Once this initial cold start is overcome, the system remains dynamic, continuing to build and strengthen connections between the data as people engage with it and the wider context changes. This allows the underlying view to improve over time, while ongoing evaluation pipelines, monitoring and human review help keep outputs relevant, controlled and up to date as roles, content and business needs evolve.

IMPACT

Consumption and adoption of Spark have continued to climb since last year, alongside real financial savings, not least due to a reduction in technology debt from decommissioning the old learning platform.

Connecting the rich consumption data to capability and business impact is the next step, starting with having a single definition of capability and productivity across KPMG's many different businesses – work now being led directly with the firm's leadership.

“Other systems have loads of choice, but it means that nobody can actually find anything. The reality is, every user needs one pathway – one pathway, personalised to what you’re doing today and what you’re expected to do tomorrow.”

– Michael Snowden

CHALLENGES AND LESSONS LEARNT

Staying narrowly focused is harder than it sounds. The underlying data is rich enough that it’s tempting to chase interesting problems well outside learning altogether. The discipline that’s mattered most is holding the team to features that create defined value on the roadmap, and being ruthless about what’s deliberately not being built: almost any single piece of this system could easily become a multi-year project in its own right if left unchecked.

Value has to be translated, not just delivered. Leaders who fund and prioritise this kind of work don’t think in learning terminology. What they respond to is speed, competency and value, and translating the work into those terms, consistently, has been as much a discipline as building the system itself.

“Your lane” is wider than you think.

Many people assume their lane is the same as their industry or sector. But to be successful on a project like this, it needs to cut across learning, resourcing, and talent and recruitment as one connected value chain – multiple ways of thinking, multiple priorities, not a narrow vertical.

NEXT STEPS

1. Using the same data model to recognise internal experts and connect people to them directly, not just to content.
2. Moving Spark’s coaching layer from a batch-generated task list closer to a live coach – letting someone talk a task through with Spark directly and get feedback that adjusts what happens next in the pathway.
3. Expanding continuous, more rigorous assessment of how capability is actually being built, starting with a pilot already underway and extending into other heavily regulated parts of the business.




THE CHALLENGE OF CHANGE

AI enables a wide range of L&D activity. From four years of our survey, however, it seems practitioners are focusing on one part of that range. They are mostly doing what they have traditionally done, but faster, at greater scale, and this year more thoroughly.

In April 2024, Eglė Vinauskaitė's Complexity Scale (**Figure 9**) set out the range of ways L&D was using AI, from administrative tasks and content production on the left to skills management on the right. It also explained why most activity sat on the left. The simple uses could be carried out by one person with a free account, bypassing the bureaucracy and politics of adopting new technology. Each use further to the right needed more: extra skills, technology, relationships and data from elsewhere in the business.

AI in L&D Complexity Scale



© Eglė Vinauskaitė, 2024

Note: neither comprehensive nor unchanging

Figure 9: The Complexity Scale

The Immaturity Model




Figure 10: The Immaturity Model

NO GUARANTEE OF PROGRESSION

The scale was popular, and organisations were keen to find their place on it. But its shaded background suggested something we came to regard as misleading: that starting with a simple use of AI leads, inevitably, to more sophisticated use. We could find no evidence for that progression.

In our 2024 annual report, we grouped the uses into three clusters – **Internal efficiency**, **Point solutions** and **Business integration** – and deliberately drew no arrows between them. We named the result the Immaturity Model (Figure 10), to contrast it with models suggesting gradual but automatic movement from simple to sophisticated use. In our view, expertise in producing content with AI does not lead to expertise in inferring skills, any more than an arc welder can do a radiologist's job because both use electricity.

THE SAME MESSAGE IN EVERY MODEL

What does carry L&D to the right? Each of our models has given the same answer: change, both in L&D and in its relationship with the organisation.

"The more sophisticated the use case," we wrote in the interim 2024 report, "the more L&D needs to bring in the skills, knowledge, data and technology from elsewhere in the business." To use AI well, L&D needed to have strategic conversations and build relationships.





Figure 11: The Transformation Triangle

The Immaturity Model made this point explicit. Internal efficiency is within L&D's own power. Moving beyond it needs an 'Open' L&D department, ready to innovate and to work closely with the business. Breaking out of the focus on internal efficiency requires several elements, including strong leadership and an organisational mandate.

We expanded this idea in The Transformation Triangle, introduced in our annual AI in L&D report for 2025 and further developed in **The Transformation Triangle** (May, 2026). It suggested the traditional focus of L&D on content production was under threat from AI, and outlined – based on case studies and extensive interviews with leading practitioners – three viable alternatives to head towards. These were **Skills Authority**, **Enablement Partner** and **Adaptation Engine**, on the nodes of the triangle in **Figure 11**. We called these different 'organising logics'.

As well as naming possible destinations of travel, the Transformation Triangle was explicit about the route to them. Drawing an analogy with space flight, it said the journey requires L&D to reach escape velocity in a flight from its traditional focus on content. This conveyed the sense both of the sustained effort required to change L&D's organising logic and the fact that it was possible to work at using AI well, but still fail to transform how the L&D department worked due to the structural, cultural and capability drags fighting against movement.

These three models developed the same themes in more detail: that L&D's traditional differentiator of developing content was under threat from a technology capable of doing the same thing far faster and cheaper, and that moving to a different organising logic for the department would not happen automatically. It required a transformation, fuelled by good leadership, sustained effort and the right organisational environment.



MOVEMENT WITHOUT TRAVEL

Four years of data have now tested that argument. Across those years, AI use has risen steadily while the activities on the right of the Complexity Scale have largely stayed where they were. The uses that depend on integration with the business – AI drawing on data, systems and relationships beyond L&D's own – are barely more common now than at the start. Effort has grown but the centre of gravity has not moved.

Within departments, more AI does bring change. But most of that change is internal: a task now handled by AI, a team that directs and reviews rather than produces. Far fewer report a change in what the organisation asks of them, and those who do are marked out not by using more AI, but by using it for different things.

Each year we have built a model from that year's data, and each has reached the same conclusion: change has to be chosen. This year's data has given us no reason to think otherwise.



CONCLUSION

AI promised L&D a revolution. What we got was something subtly different.

There have certainly been extraordinary changes. The case studies in this and previous reports show L&D teams doing things that would have seemed impossible only a short time ago. AI has also brought an uncertainty about work that many of our respondents find deeply unsettling.

But the noise of revolution drowns out something quieter and more important. Almost four years after ChatGPT's public launch, L&D is changing in two ways: consolidation and growth. Consolidation takes place around the core of what L&D does. Growth builds beyond it.

Consolidation shows up in every one of our surveys. Year after year, L&D uses AI to produce more content, faster. But this is about more than just increasing the

volume of production. L&D now uses AI with far more sophistication, most clearly in research and analysis. The result: L&D arguably doing its core job better.

One small group of practitioners has gone further and is building beyond the core. These are the 'break-out' group of 14% of workplace L&D teams who say the rest of the organisation now approaches them for different things. They use AI for content as much as anyone, but far fewer of them count content creation among its main benefits. What marks them out is where they point it: twice as likely to have AI working with employees directly, building their skills rather than their materials, and almost three times as likely to automate their own processes and build their own tools.

Throughout this report we have tried to be suitably cautious, and to give the right caveats around the data. Here, let us



instead speculate, and strike a positive tone. The break-out group seems to be achieving something L&D has long sought: a different relationship with the business, possibly a more valued one.

There is plenty of disagreement about just how revolutionary AI will ultimately prove to be. Some expect it to transform organisations and entire industries; others point out, reasonably, that organisations move much more slowly than technology, and that much of what is being promised today may turn out to be hype. For L&D, we are not sure that argument matters very much. AI has already had an enormous effect on our field.

And the implications run in two directions. AI is changing how L&D works, while simultaneously changing the organisations L&D exists to serve – and therefore what those organisations need from it.

We believe the next phase may be as much about L&D's role in the organisation as its use of the technology.

There are emerging signs. AI use for data analysis has grown more consistently than almost anything else we have tracked. Taken together with the early signal from our break-out group, it allows us to imagine a function with more evidence, more influence and a different relationship with the business.

This, then, is the challenge of the future: not how much AI L&D uses, but how that use affects its role. It is the new frontier of L&D, with risk, yes, but also full of potential. We look forward to reporting next year on the progress L&D has made and the challenges it has overcome.

RESOURCES



THE TRANSFORMATION TRIANGLE

By Eglė Vinauskaitė and Donald H Taylor

Published: May 2026

The Transformation Triangle is a framework mapping three strategic directions for L&D beyond content, based on a qualitative study of L&D leaders navigating AI-driven change.

[Go to the report](#) or scan the QR code to view.

Link: bit.ly/AILD05-TTT



THE GLOBAL SENTIMENT SURVEY 2026

By Donald H Taylor

Published: February 2026

In this survey, L&D professionals from about 100 countries shared their views on what challenged and excited them about L&D in the year ahead.

[Go to the report](#) or scan the QR code to view.

Link: bit.ly/GSS26results



AI IN L&D: THE RACE FOR IMPACT

By Donald H Taylor and Eglė Vinauskaitė

Published: September 2025

Last year's AI in L&D report featuring case studies from Microsoft, ServiceNow, TTEC, KPMG UK, Leyton, mci group, talabat and EPAM.

[Go to the report](#) or scan the QR code to view.

Link: bit.ly/AILD04-Impact

SPONSORS



Sana is an AI lab transforming how organizations learn and access knowledge. Sana Learn is an AI-native learning platform that unifies LMS, LXP, authoring, and virtual classroom in one product. It delivers personalized, AI-supported learning for every employee and gives L&D one place to design, deliver, and measure impact. Global enterprises including Spotify, Polestar, and Accenture rely on Sana Learn to onboard, develop, and reskill their workforces at scale. Founded in Sweden in 2016, Sana has offices in Stockholm, London, and New York.



Docebo is the AI-powered learning platform built for the enterprise. We go beyond traditional LMS to connect skills intelligence, knowledge management, learning execution, and AI agents into a single system that continuously identifies capability gaps, delivers targeted learning, and measures real impact. With the acquisition of 365Talents, Docebo gives organizations the infrastructure to close skill gaps at scale, not just track course completions. Over 3,500 enterprises trust Docebo to make their people measurably better, faster than the rate of change demands.



OpenSesame is a workforce learning platform that helps organizations get the right skills to the right people at the right moment. One platform gives teams access to 70,000+ curated courses, tools to build their own, and the expert support to deliver it all where work happens, so people grow and the business stays protected.



Speeex is the people development platform for business coaching, language training, strategic mentoring, and intercultural training, trusted by global corporations to upskill their workforce. Trusted by 2000 organizations and 8 million users globally, Speeex delivers AI-driven learning with certified trainers and coaches, CEFR-based assessment, and role-specific content. Built for corporate L&D teams, Speeex integrates with SAP SuccessFactors, Workday, and Microsoft Teams, and meets enterprise-grade security standards.



When citing this report, please refer to

AI in L&D 2026: Around the Core

All contents © Donald H Taylor Services Limited 2026, © Eglė Vinauskaitė 2026

AI in L&D 2026 v 1.0

www.donaldhtaylor.co.uk

www.nodes.works