Who Will Become the Experts?

Part 4: How Do You Become an Expert When AI Can Do the Work?

What students, trainees and educators should change, and what they should keep

By Dr Peter Fields

First posted: 7 November 2026

THE 2-MINUTE SUMMARY

The previous articles examined how AI is changing junior employment, why producing better work with AI does not necessarily develop expertise, and what employers can do about it. This final article considers the choices facing students, trainees and educators. The research presents a mixed picture. AI can accelerate learning, provide personalised explanations and improve access to educational support. But it can also allow people to complete work without developing the knowledge or judgement needed to do it independently.

The answer is not to reject AI or preserve traditional teaching methods unchanged. It is to distinguish between using AI to accomplish a task and using it to develop the capability to accomplish that task. For students and trainees, this means combining AI assistance with independent thinking, deliberate practice and opportunities to test their understanding. For schools, colleges and universities, it means teaching AI literacy while protecting the development of underlying knowledge and redesigning assessments to measure both independent competence and effective AI use.

There is no established universal method. The most defensible approach is to retain educational practices supported by evidence, use AI where it demonstrably improves learning, and test whether those improvements persist when assistance is withdrawn.

The student who can produce the answer but cannot explain it

Imagine two graduates arriving for an interview at an engineering consultancy. Both have excellent academic records. Both are comfortable using AI, and both have prepared an impressive analysis of a technical problem sent to them before the interview.

The interviewer asks the first candidate to explain why a particular design is preferable to the alternatives. The graduate gives a confident answer, closely resembling the analysis they submitted. The interviewer then changes one assumption. What happens if the operating

temperature increases by 30 degrees? The candidate hesitates. The AI tool used to prepare the original analysis had not considered that particular possibility.

The second graduate has used AI extensively as well. But rather than simply asking it to solve the original problem, they used it to explore alternatives, explain unfamiliar concepts and challenge their initial assumptions. They cannot immediately calculate the precise effect of the increased temperature. Nevertheless, they can explain which factors matter, what additional information they need and how they would approach the problem.

Which candidate would you employ? The example is hypothetical, and real interviews are rarely this straightforward. The first candidate might be better in other respects, while the second might have misunderstood something important. Nevertheless, it illustrates the distinction running through this series: **the ability to produce an expert-looking answer is not the same as possessing the expertise necessary to understand, evaluate and defend it.**

That distinction matters particularly to young people entering a world in which AI can undertake increasingly sophisticated intellectual work. Should they spend years learning skills that technology may perform more effectively? Or should they concentrate on becoming proficient users of AI, allowing the technology to supply knowledge whenever it is required? Neither extreme is particularly attractive. But there is a more useful question: what should someone learn for themselves, what can they safely delegate, and how should the relationship between the two change as they become more experienced? That is the subject of this final article.

Students have already adopted AI. Education is still catching up.

The debate about whether students should be allowed to use generative AI has become increasingly detached from their actual behaviour. In March 2026, the Higher Education Policy Institute published its third annual survey of undergraduate AI use. Based on responses from 1,054 full-time UK undergraduates, collected in December 2025, the study found that 95% reported using AI in at least one way. Some 94% said they had used generative AI when preparing assessed work. [1]

Those figures do not mean that almost every student is using AI to cheat. The reported activities included explaining concepts, summarising reading, suggesting research ideas and helping students organise their thoughts. Indeed, 12% reported including AI-generated text directly in a piece of assessed work, as distinct from the separate response categories concerning AI-generated text that was edited before use. The figure had increased from 8% in 2025. The survey also revealed considerable uncertainty about how universities were responding. Only 48% of students felt their teaching staff were helping them develop AI skills

for their future careers, while 68% considered understanding and effective use of AI essential to thrive in today's world.

The figures are self-reported and do not measure whether students are actually becoming more knowledgeable or capable. Nor do they establish that universities are universally failing to provide adequate support. Nevertheless, they suggest a growing mismatch between students' adoption of AI and their confidence in the guidance they receive.

There is another important finding. Students from more advantaged socioeconomic backgrounds were more likely to use certain AI tools than students from less advantaged backgrounds. The researchers could not establish precisely why, and the availability of free AI services suggests that subscription costs alone do not explain the difference. This creates a risk that educational institutions could inadvertently widen existing inequalities by assuming that all students have equal access to the technology and the knowledge necessary to use it effectively.

But the fundamental challenge is not simply access. **Students need to understand how to use AI to become more capable, rather than merely more productive.** That is a different educational objective.

What does the research tell us about learning with AI?

Parts 2 and 3 examined the distinction between AI-assisted performance and independent competence. Three experiments help explain why this matters, although their different designs and outcomes should not be conflated. In a randomised study involving nearly a thousand Turkish high-school mathematics students, Bastani and colleagues found that unrestricted AI assistance improved performance during practice, but the students subsequently performed approximately 17% worse on an unaided examination than the control group. A guided tutor offering hints and teacher-prepared explanations largely avoided that harm, without producing a statistically significant improvement in independent performance over the control group. [2]

The findings were not uniformly negative. In a Harvard physics experiment involving 194 students, a specially designed AI tutor outperformed an active-learning classroom approach on immediate assessments, although the intervention covered only two lessons and did not establish long-term retention. [3] In a separate experiment with 211 undergraduates, Contractor and Reyes found that access to ordinary generative AI improved scores on subsequent unaided knowledge tests, with much of the advantage persisting approximately a week later. Students who used AI for explanations tended to do better than those who used it primarily to generate text, but those patterns of use were not randomly assigned and cannot establish causation. [4]

The distinction is therefore not between AI that helps and AI that harms in every setting. **AI can support genuine learning, but an improved result while the tool is available is not proof of independent competence.** How assistance is designed and used matters; no single approach has yet been established as universally best.

The temptation to skip the difficult part

Suppose a law student is trying to understand a difficult judgment. They could read it, identify the legal principles and work through the reasoning, or obtain an AI summary in thirty seconds.

The second approach might provide an excellent introduction to the case. It could explain unfamiliar terminology and help the student identify the passages requiring closer attention. But what happens if the student never reads the judgment? They may know what the AI says the court decided without understanding how the judges reached that conclusion.

The distinction is not confined to law. An engineering student might obtain a technically correct solution to a complex calculation without understanding the underlying physical principles. A business student might receive a sophisticated market analysis without appreciating why particular assumptions matter. In each case, the technology can reduce the time required to complete the task. Whether it also improves learning depends on what the student does with the result.

Should students be allowed to delegate the capability an assignment is designed to develop? Sometimes yes, where the purpose is learning to use professional tools. But where the purpose is to understand the underlying reasoning, delegation may defeat the objective.

There is an established psychological basis for this distinction. In their 2016 review of cognitive offloading, Risko and Gilbert examined how people use external resources to reduce mental demands. They recognised that offloading can improve performance and free mental resources for other activities. However, they also discussed how the educational consequences depend on what is being offloaded and whether the resources released are redirected towards productive learning. [5]

The question is not merely whether AI makes a task easier, but whether the effort being removed develops the intended capability. **Difficulty has no intrinsic educational value. The value lies in what the effort teaches.** Hours of unnecessary mechanical work may be pointless; avoiding every challenging intellectual decision may be equally unwise.

Becoming proficient at using AI is not the same as becoming an expert

Students are rightly encouraged to develop AI literacy. That can mean constructing prompts, using research tools, automating workflows or generating code. These skills may improve productivity and employability without establishing competence in the underlying subject.

A proficient AI user might be excellent at obtaining a legal memorandum while lacking sufficient legal knowledge to recognise that it contains a fundamental error. Conversely, an experienced professional with limited knowledge of prompting might be particularly effective at evaluating the memorandum because they understand the legal issues. The two capabilities complement each other.

UNESCO's 2024 *AI Competency Framework for Students* recognises that AI education should extend beyond operating particular tools. It identifies twelve competencies across four dimensions: a human-centred mindset, AI ethics, AI techniques and applications, and AI system design. These are organised into three developmental levels: understand, apply and create. [6] The framework concerns a broader objective than preparing students for employment. People need sufficient understanding of AI to evaluate its influence on society and make informed decisions about its use.

Students need both an understanding of AI and an understanding of the subject in which they intend to apply it. Accountants still need accounting knowledge; engineers need engineering knowledge; lawyers need legal principles, procedure and professional responsibility. AI's ability to supply information does not remove the value of those foundations.

Without some understanding of the subject, how can someone reliably identify when the information supplied is incomplete, irrelevant or wrong? That does not mean professionals must memorise everything that AI can retrieve. The important distinction is between knowledge that must be available internally to exercise judgement and information that can safely be obtained from external resources. The boundary will differ between occupations and will change as technology develops. There is no reason to pretend otherwise.

What should students and trainees actually do?

It is easy to tell students to think critically about AI; it is harder to say what that means while they are learning. I would start with one principle: **use AI to support the thinking you need to develop, rather than routinely having it replace that thinking.** A student preparing an essay or professional report might first identify the problem, outline an approach and record the reasoning behind it. AI could then explain unfamiliar concepts, identify omissions or challenge assumptions. The student would check those suggestions against reliable sources, revise the argument and attempt to explain it independently.

The approach is compatible with established educational psychology. Dunlosky and colleagues' review of learning techniques found particularly strong support for practice testing and distributed practice, in which knowledge is revisited over time. [7] Those findings predate modern generative AI and do not validate any particular prompting technique. Nevertheless, AI might usefully generate practice questions, challenge attempted solutions or revisit a topic several days later, provided the learner does the intellectual work and checks the feedback. Recognising an AI explanation as plausible is not equivalent to being able to reproduce and apply the underlying reasoning.

Not every task calls for this process. Someone who already understands a task can reasonably delegate substantial work to a tool. The distinction is whether the immediate objective is productive output or the development of a capability. The five practical habits later in the article offer a way for students to apply that distinction.

The importance of knowing when you are wrong

Perhaps the most valuable professional skill in an AI-assisted environment is the ability to recognise the limits of one's knowledge. Experts are not simply people who possess large amounts of information. They are also people who understand which conclusions can safely be drawn from that information, which assumptions are questionable and when further investigation is necessary. This is closely related to metacognition: the ability to monitor and evaluate one's own thinking.

Risko and Gilbert's review identified circumstances in which access to external information affected people's assessments of their own knowledge. In some experiments, people who had used internet searches became more confident about what they knew, even when the information had been obtained externally. Other experiments produced different effects, including circumstances in which access to external information made people more willing to acknowledge uncertainty. The effects were therefore not universal. [5]

Nevertheless, the research raises an important possibility for AI-assisted learning. A student who can produce increasingly sophisticated work may become more confident without acquiring corresponding independent competence. There is no simple test that can eliminate this risk. However, students can examine it. After completing an AI-assisted task, they might attempt a related problem without assistance, explain the reasoning to another person or identify weaknesses in an AI-generated solution. These activities provide information about what they understand, rather than simply how good the completed work looks.

They also reveal something about professional judgement. An expert is not expected to know every answer immediately. But they should increasingly be capable of recognising what they do not know, identifying the information required and evaluating whether the resulting

answer is reliable. That is a capability worth developing, whether or not AI becomes substantially more sophisticated.

Should schools and colleges change what they teach?

The discussion so far has concentrated on individual learners. But students cannot carry the entire responsibility for adapting to AI. Schools, colleges and universities determine much of what students learn, how they learn it and how their achievements are assessed. If AI changes the relationship between completing a task and developing competence, educational institutions need to consider the consequences.

That does not mean rewriting the entire curriculum around the latest technology. Indeed, I would resist that temptation. Knowledge of mathematics, scientific principles, language, history and other established subjects remains valuable. AI may change how people apply that knowledge, but it does not establish that the knowledge has become unnecessary. The more difficult question is how to incorporate AI without undermining the acquisition of those foundations.

For younger students, this requires particular care. The Department for Education's policy on generative AI recognises potential benefits in lesson planning, educational resources, personalised support and reducing teachers' administrative workloads. However, it also acknowledges that evidence concerning the effects of pupil-facing AI on learning remains limited. The Department advises schools and colleges to consider safeguarding, data protection, age restrictions, bias, accuracy and the continuing importance of the teacher-pupil relationship. [8]

That is a sensible starting point. A school might use AI to help teachers prepare differentiated exercises or explain a difficult concept in several ways. It might also provide supervised opportunities for pupils to explore how AI works, where it can make mistakes and why its outputs require checking. But there is no compelling educational argument for allowing young children unrestricted access to systems that can complete the very activities through which they are learning basic literacy, numeracy and reasoning. Nor does the evidence justify banning all pupil-facing AI.

The Department for Education's separate *Generative AI: Product Safety Standards*, updated in January 2026, go further in addressing the design of learner-facing tools. They expect suppliers to consider the risk of cognitive deskilling and, by default, to provide hints or partial steps before full solutions, encourage learners to attempt problems themselves and monitor occasions when the tool undertakes the thinking for them. [10] These are expectations directed primarily at developers and suppliers, not experimental proof that every proposed feature improves learning or a statutory rule requiring schools to adopt a particular product.

The appropriate balance will depend on age, subject, learning objectives and the safeguards available. **The guiding principle should be educational purpose, not technological enthusiasm or hostility.**

Universities face a more difficult problem

Universities have traditionally relied on essays, reports, projects and other assignments to assess students' understanding. Generative AI complicates that arrangement. A student can now produce work that resembles a competent university assignment without necessarily having acquired the knowledge or skills that the assignment was intended to assess. This creates a problem both for academic integrity and for the validity of assessment.

An excellent essay might demonstrate substantial intellectual achievement. Alternatively, it might demonstrate that the student is particularly effective at generating and editing AI-produced text. The finished document alone may not reliably distinguish between those possibilities. Some institutions have responded by introducing more supervised assessments, oral examinations or restrictions on AI use. Others are increasingly incorporating AI into teaching and assessment.

There is a legitimate case for both approaches. The mistake would be to assume that either provides a complete solution.

Assessing independent competence

If students must acquire particular knowledge or reasoning skills, universities should have reliable ways of establishing that they possess them. This might involve supervised examinations, oral questioning, practical demonstrations or unfamiliar problems undertaken without AI assistance. Such assessments are not necessarily a return to outdated educational practices. Their purpose is to examine capabilities that cannot otherwise be established reliably from an AI-assisted submission.

However, the assessment must reflect the capability being tested. Requiring students to memorise large quantities of easily accessible information may be unjustified if the relevant professional skill is the ability to interpret and apply it. A carefully designed assessment should distinguish between unnecessary memorisation and the internal knowledge required for competent judgement.

Assessing effective AI use

Equally, students entering AI-assisted workplaces need opportunities to demonstrate that they can use the technology effectively. An assessment might therefore require them to use AI, evaluate its output, identify errors and explain the reasoning behind their final conclusions.

In law, that might involve checking whether an AI-generated memorandum accurately represents the authorities it cites. In engineering, it could involve identifying weaknesses in an AI-produced design. In business, students might be asked to challenge the assumptions underlying an AI-generated market forecast. Such tasks address a capability that conventional unaided examinations may not measure, although their effectiveness as assessments would still need to be established.

Why both types of assessment matter

The two approaches examine different questions. An unaided assessment asks what the student can do independently. An AI-assisted assessment asks what the student can accomplish with the technology and how effectively they supervise its contribution. Neither necessarily establishes the other. The HEPI 2026 report recommends allowing students to experience both AI-free assessments and assessments requiring AI use, alongside clear guidance on academic integrity. [1] I consider that a sensible direction.

However, implementing it properly requires more than dividing assessments into two categories. Universities need to establish which capabilities matter, how those capabilities should be measured, and whether the assessments provide reliable evidence of learning. Those decisions will differ between subjects. The distinction should also remain clear to students. If an institution allows AI assistance in one assignment but prohibits it in another, students should understand why. That explanation is educationally more useful than a vague instruction to use AI responsibly.

Academic integrity cannot depend on detecting AI-generated text

There is a temptation to treat AI use primarily as a policing problem. If an institution could reliably identify which passages were generated by AI, perhaps conventional assessments could continue largely unchanged. But this approach has limitations.

AI-generated material can be edited substantially, and the boundary between acceptable assistance and prohibited authorship may not be obvious. A student might legitimately use AI to explain a concept, improve accessibility or organise research notes. Another might use it to produce an entire assignment. Both have used AI, but their intellectual contributions are very different.

The Joint Council for Qualifications' guidance, effective from September 2026, requires candidates to submit work that demonstrates their own knowledge, skills and understanding. AI assistance used for planning, structuring, drafting or editing may need to be acknowledged; acknowledging it does not entitle the student to marks for work they have not independently produced. Centres must authenticate assessed work and investigate credible concerns about

misuse. The guidance also makes clear that AI-detection scores have limitations and must not be used in isolation to determine malpractice. [9]

The Exams Office, a separate membership organisation, announced complimentary access for its member centres to a detection tool from September 2026, initially including 50,000 words. It expressly stated that the tool was an aid to investigation, not a substitute for teachers' judgement. This was **not a JCQ-issued detection service**. [11] For schools and colleges, the rules of the relevant qualification still govern what AI use is permitted. As NCFE's assessment guidance illustrates, permission to access the internet in a particular assessment does not by itself authorise every use of generative AI. [12]

Universities have greater discretion over many of their own assessments, but face the same problem of distinguishing genuine student achievement from externally generated work. HEPI's 2026 survey also recorded student anxiety about false accusations of AI misconduct. Those accounts do not establish how often false accusations occur, but they show why fairness matters. [1] I would therefore place more emphasis on assessment design and clear expectations than on attempting to identify every instance of AI assistance after the event. That does not mean abandoning integrity rules; it means making them serve a coherent educational purpose.

Students should know when assistance is permitted, what must be acknowledged and which intellectual contribution they are expected to make themselves. Equally, institutions should be cautious about treating an automated detection score as proof of misconduct. Academic standards must be protected, but so must procedural fairness.

Could AI make education more accessible?

There is a potentially important benefit that should not be overlooked. Good educational support is not distributed equally. Students with access to private tutors, well-resourced schools, knowledgeable family members or experienced professional mentors may enjoy advantages unavailable to others. AI could reduce some of those differences.

A student struggling with a concept might obtain an explanation at any hour, ask follow-up questions without embarrassment or request additional practice. Students learning in a second language may benefit from translation and clarification. Learners with certain disabilities may find that appropriately designed tools improve access to educational materials. These possibilities are significant.

The Harvard tutoring experiment provides some evidence that a carefully designed AI system can improve immediate learning outcomes. [3] However, it does not establish that ordinary commercial chatbots will deliver the same benefits, or that AI will automatically reduce

educational inequality. The HEPI survey's findings concerning unequal uptake are a reminder that access is not simply a matter of whether a free version of a tool exists. [1]

The problem extends beyond money. Students with stronger subject knowledge may be better equipped to evaluate AI outputs, recognise errors and exploit the technology effectively. Others may need more guidance even when everyone is given access to the same tool. Suitable devices, reliable connectivity and confidence in using AI also matter. Paid services can offer capabilities unavailable in free versions; institutions that make them necessary for assessments without providing access risk introducing another source of disadvantage.

Consequently, providing everyone with the same AI tool does not necessarily provide everyone with an equal educational opportunity. The appropriate response is not to withhold the technology. It is to combine access with guidance, appropriate teaching and opportunities to develop the knowledge needed to use it well.

There is also a risk of replacing human support rather than supplementing it. A student who receives AI explanations instead of access to a knowledgeable teacher or supervisor may gain convenience while losing other valuable aspects of education. The evidence does not support assuming that an AI tutor is an adequate substitute for the teacher-student relationship across every educational setting. **AI should be evaluated by whether it improves the quality and accessibility of learning, not simply by how cheaply it can deliver answers.**

What about graduates entering the workplace?

The transition from education to employment presents another difficulty. Students may acquire strong academic qualifications and considerable experience using AI, yet still lack opportunities to exercise professional judgement. Part 1 examined the possibility that AI will reduce some of the junior work through which professionals traditionally acquired experience. Part 3 considered how employers might redesign training rather than simply preserve the old employment structure.

But what can an individual graduate do? The answer cannot simply be to develop better prompting skills. Those skills may help someone work productively, but they do not necessarily provide a route towards professional expertise. A graduate should also consider the quality of the experience an employer is offering. Will they have opportunities to analyse unfamiliar problems, make recommendations and receive feedback? Will experienced colleagues explain their decisions? Will the employee gradually assume greater responsibility? Or will the role consist largely of operating AI systems and formatting the results?

Working extensively through AI may be entirely appropriate. But a graduate hoping to develop into a senior professional should distinguish employment that produces immediate output from work that builds progressively greater capability.

The employer offering the highest salary or most impressive technology may not provide the best opportunities to learn; nor does repetitive manual work automatically constitute superior training. What matters is the capability the employee is likely to develop.

Developing expertise outside formal training

Graduates may need to pursue difficult projects, seek feedback, study underlying principles and use AI to explore unfamiliar problems. Depending on the occupation, qualifications, supervised practical experience and mentoring may remain essential.

However, there are limits to what individuals can accomplish independently. A trainee cannot recreate years of experience managing real clients, operating industrial equipment or exercising responsibility for complex professional decisions merely by interacting with an AI system. Simulations and practice exercises can contribute to learning, but they are not always substitutes for exposure to real consequences and professional supervision. This is why the employer's role remains important. The development of expertise is not simply a matter of individual motivation. It also depends on access to appropriate opportunities.

Should students worry about becoming dependent on AI?

Repeated reliance on AI may weaken some independent capabilities, but professionals have always depended on tools: engineers on calculators and modelling software, lawyers on electronic databases, accountants on spreadsheets and doctors on diagnostic equipment.

The fact that someone could not perform their entire job without these technologies does not establish that they lack professional competence. We would not normally insist that a modern engineer perform every complex calculation manually before recognising their expertise. But the professional must still detect errors of input, modelling or interpretation. Generative AI introduces a particular difficulty: it can produce plausible language and reasoning that are substantively wrong, without signalling the error. Can the professional recognise the problem, identify what needs checking and respond appropriately?

That does not necessarily require reproducing every calculation or research process from memory. But it does require sufficient understanding to exercise judgement. The same principle should apply to AI. A trainee who cannot work without AI is not necessarily incompetent. The more important question is whether they can exercise the knowledge and judgement required to use it responsibly.

Nevertheless, where a particular independent capability remains important, students should practise it. If someone is learning legal reasoning, they need opportunities to reason through legal problems. If someone is learning to programme, they need opportunities to understand

code rather than merely accept its generation. If someone is learning financial analysis, they need to understand why particular assumptions produce particular results.

These activities may be undertaken with AI assistance at some stages and without it at others. The appropriate balance should be determined by the learning objective. That may sound unremarkable. But it is precisely the distinction that risks being overlooked when educational success is measured principally by the sophistication of the completed work.

A practical approach for students

I would avoid imposing a complicated new learning system. Students already have established methods that can be effective, particularly practice testing, feedback, revisiting material over time and applying knowledge to unfamiliar problems. AI can be incorporated into those activities without abandoning them.

For someone learning a subject, I would suggest five practical habits:

1. **Attempt to understand the problem before requesting a complete answer.** This need not involve hours of unaided work. The objective is to identify the underlying question and establish what is already understood.
2. **Use AI to explain, question and challenge.** Ask for alternative explanations, counterarguments, possible errors and practice problems. Do not assume that every explanation or correction is accurate.
3. **Verify important information independently.** Use primary sources, authoritative textbooks, professional guidance and other reliable material, particularly where accuracy matters.
4. **Test what remains when AI is removed.** Revisit the subject later, attempt unfamiliar problems and explain the reasoning without assistance where independent competence is a relevant objective.
5. **Seek human feedback and real experience.** AI can be useful, but teachers, supervisors, colleagues and practical experience provide forms of learning that should not automatically be replaced.

These are recommendations derived from established learning principles and the emerging AI research, not a newly validated educational method. I would also avoid treating them as rules that must be followed for every interaction. An experienced professional using AI to summarise administrative material does not necessarily need to repeat the task independently afterwards. The point is to recognise when the objective is learning and adjust the use of AI accordingly.

What should educational institutions do differently?

Schools, colleges and universities should identify the capabilities students need and examine how AI affects their development. That involves three decisions.

First, institutions need to determine which knowledge and skills should remain central to their curricula. The fact that AI can perform a particular task is not sufficient reason to remove it from education. The task may be important because it develops understanding required for more advanced work. Conversely, some traditional exercises may be inefficient ways of achieving the intended learning objective.

Second, institutions need to determine where AI can improve teaching and learning. This could include personalised explanations, additional practice, formative feedback and support for students with different learning needs. But the benefits should be evaluated, rather than assumed.

Third, institutions need assessments that distinguish between independent competence and effective use of AI. This may require a combination of supervised work, practical assessment, oral explanation and AI-assisted tasks. There is no evidence that one particular combination will work best across every subject. The appropriate arrangements should be tested and revised as experience develops.

There is a practical constraint: teachers need time, training and suitable tools. Without them, AI-supported learning risks becoming another unfunded obligation. Assessment reform also requires careful attention to validity, fairness, accessibility and administrative cost. The objective should be better evidence of learning, not merely a new examination format.

Who will become the experts?

We began this series with an apparently straightforward concern: if AI performs increasing amounts of junior work, how will the next generation gain the experience required to become senior professionals? The earlier articles examined changing recruitment incentives, the distinction between impressive output and genuine expertise, and employers' responsibility to redesign training. This final part has considered how students and educators should respond. None of those inquiries establishes that traditional apprenticeships must survive unchanged, or that AI necessarily undermines learning.

The answer is less dramatic than some predictions about the future of education. There is no sound basis for abandoning established principles of learning simply because AI can provide sophisticated answers. Nor is there a persuasive argument for excluding a technology that can demonstrably improve learning in some circumstances.

The evidence instead points towards a distinction that should become central to educational and professional development: **using AI to produce work and using AI to become more capable are different objectives.** They can be achieved together, but one does not guarantee the other. Students should therefore learn to exploit AI's capabilities while retaining opportunities to develop and test their own understanding. Educational institutions should adapt curricula and assessments to reflect both independent competence and responsible AI-assisted performance.

Employers should provide opportunities for employees to develop the judgement and experience required for progression, recognising that these may no longer arise automatically from conventional junior work. Some familiar activities may disappear; others may remain important. AI may enable people to acquire certain skills more quickly, while creating new capabilities that professionals must develop.

It may also change what society means by expertise. An experienced professional in 2035 may spend much less time producing technical documents and considerably more time evaluating competing analyses, making decisions and supervising AI-assisted work. That is plausible, but remains a prediction. We should not design entire educational systems around confident assumptions about capabilities and labour markets that are still evolving.

What we can do is examine whether people are genuinely learning. That is a measurable objective, even if measuring it well is difficult, and it is a better foundation for educational reform than attempting to preserve every aspect of traditional training or embracing every apparent technological improvement.

The future experts will not necessarily be those who use AI least, nor those who use it most. They will be people who develop the knowledge, capabilities and judgement needed to work effectively with the tools available to them. Whether our educational institutions and employers provide adequate opportunities for them to do so remains an open question. It is also a question on which decisions made today may have consequences for many years.

Alex Asks

> *If I can explain a subject, solve a problem and produce an excellent answer in seconds, why should you spend hours learning to do those things yourself?*
>
> *Perhaps you shouldn't always have to.*
>
> *But if you cannot tell whether my answer is correct, or why it matters, have I made you more capable? Or have I simply made it easier for you to appear capable?*
>
> *Maybe the question isn't whether you should learn to work without me. It's whether working with me helps you become someone who understands what they're doing.*

References

[1] Stephenson, R. and Armstrong, C. (2026). *Student Generative AI Survey 2026*. Higher Education Policy Institute, Report 199, 12 March 2026. Survey of 1,054 full-time UK undergraduates, conducted by Savanta in December 2025. https://www.hepi.ac.uk/wp-content/uploads/2026/03/HEPI-Report-199-Gen-AI-Survey-2026.pdf

[2] Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö. and Mariman, R. (2025). *Generative AI without Guardrails Can Harm Learning: Evidence from High School Mathematics*. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. DOI: 10.1073/pnas.2422633122. https://doi.org/10.1073/pnas.2422633122

[3] Kestin, G., Miller, K., Klales, A., Milbourne, T. and Ponti, G. (2025). *AI Tutoring Outperforms In-Class Active Learning: An RCT Introducing a Novel Research-Based Design in an Authentic Educational Setting*. *Scientific Reports*, 15, Article 17458. DOI: 10.1038/s41598-025-97652-6. https://www.nature.com/articles/s41598-025-97652-6

[4] Contractor, Z. and Reyes, G. (2026). *Experimental Evidence on the Learning Impact of Generative AI*. IZA Discussion Paper No. 18792, July 2026. https://iza.org/publications/dp/18792/experimental-evidence-on-the-learning-impact-of-generative-ai

[5] Risko, E. F. and Gilbert, S. J. (2016). *Cognitive Offloading*. *Trends in Cognitive Sciences*, 20(9), 676–688. DOI: 10.1016/j.tics.2016.07.002. https://doi.org/10.1016/j.tics.2016.07.002

[6] UNESCO (2024). *AI Competency Framework for Students*. UNESCO, Paris. Published August 2024. https://www.unesco.org/en/articles/ai-competency-framework-students

[7] Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J. and Willingham, D. T. (2013). *Improving Students' Learning with Effective Learning Techniques: Promising Directions from Cognitive and Educational Psychology*. *Psychological Science in the Public Interest*, 14(1), 4–58. DOI: 10.1177/1529100612453266. https://www.psychologicalscience.org/journals/pspi/1529100612453266/

[8] Department for Education (2025). *Generative Artificial Intelligence (AI) in Education*. Policy paper, updated 12 August 2025. https://www.gov.uk/government/publications/generative-artificial-intelligence-in-education/generative-artificial-intelligence-ai-in-education

[9] Joint Council for Qualifications (2026). *AI Use in Assessments: Your Role in Protecting the Integrity of Qualifications*. Effective September 2026, particularly sections 3–8. Together with JCQ (2026), *Plagiarism in Assessments: Guidance for Teachers/Assessors*. https://www.jcq.org.uk/exams-office/malpractice/artificial-intelligence/

[10] Department for Education (2026). *Generative AI: Product Safety Standards*. Updated 19 January 2026, especially 'Cognitive development' (PDF pages 15–16). https://www.gov.uk/government/publications/generative-ai-product-safety-standards/generative-ai-product-safety-standards

[11] The Exams Office (2026). *Member Update, April 2026*, page 1, 'Artificial intelligence (AI) and plagiarism detection tool'. Access for member centres, distinct from JCQ policy guidance. https://www.theexamsoffice.com/

[12] NCFE (2026). *Use of Artificial Intelligence (AI) in Assessments*, including guidance on assessment-specific restrictions. https://www.ncfe.org.uk/customer-and-learner-support/assessment-support/use-of-artificial-intelligence-ai/