Who Will Become the Experts? · Part 2

The Illusion of Expertise: When AI Improves the Work, Does It Improve the Worker?

By Dr Peter Fields · First posted: 6 November 2026

THE 2-MINUTE SUMMARY

Generative AI can make a trainee's work look more expert without necessarily making the trainee more capable. A polished legal opinion, technical analysis or report may depend heavily on the system producing it. The distinction becomes visible when the user must explain the answer, respond to a changed assumption or perform a similar task without assistance.

Research on AI-assisted work and learning produces mixed results. In some settings, assistance improves both immediate performance and later capability; in others, people deliver stronger work without demonstrating improved independent judgement. The evidence is not a universal verdict for or against AI. Outcomes depend on the task, the user's starting knowledge and how assistance is provided.

Becoming an expert requires more than obtaining correct outputs. It involves developing mental models, recognising exceptions, testing assumptions and learning from feedback. AI can support that process, but employers and trainees need to test learning directly rather than assume that greater productivity proves greater competence.

Imagine two trainee lawyers working on the same unfamiliar legal problem.

The first spends several hours researching the relevant law, reading cases, identifying conflicting authorities and struggling to produce an answer. The initial draft is imperfect. A supervising lawyer points out several errors, explains what has been misunderstood and sends the trainee away to try again.

The second uses AI. Within twenty minutes, there is a well-structured analysis, properly presented and containing the relevant legal principles. After checking the authorities, asking AI to explain two points she does not understand, and making a few adjustments, the trainee submits it. The second lawyer has produced better work in considerably less time. From the firm's perspective, that looks like progress.

But which trainee has learnt more? Which would be better equipped to tackle a different problem next week, without AI assistance? Which is developing the judgement required to become a senior lawyer, consultant or partner?

We cannot answer those questions from the finished documents. And that is the problem. AI can improve the quality of what someone produces without necessarily improving their ability to produce it.

The distinction matters enormously. If we confuse the quality of someone's output with their underlying competence, we may fundamentally misunderstand how AI is affecting professional

development. Fortunately, we are beginning to acquire evidence that allows us to examine the question rather than simply speculate about it.

An argument that began 2,400 years ago

The concern that technology might give people the appearance of knowledge without genuine understanding is remarkably old.

In Plato's Phaedrus, Socrates recounts a story about the Egyptian god Theuth, who presents writing as an invention that will improve human memory and wisdom. The king, Thamus, is unimpressed. He argues that people will become dependent on external written records rather than exercising their own memories. Worse, access to written information may give them the appearance of wisdom without the understanding that wisdom requires.

It is tempting to dismiss the story as ancient resistance to technological progress. After all, writing has enabled extraordinary advances in education, science and civilisation. Few people would seriously advocate abandoning it. But there is a subtlety in the argument that remains relevant. Possessing information, being able to retrieve it and understanding it are different things.

A medical student with a textbook is not necessarily a competent doctor. Someone with access to every reported case is not automatically a capable lawyer. The technology can provide knowledge that the individual does not personally possess. The question is what happens when that knowledge must be understood, evaluated and applied.

Generative AI makes the distinction more difficult because it does not merely retrieve information. It can explain the information, organise an argument, suggest conclusions and produce work that resembles the output of an experienced professional. The result may look like expertise even when the person operating the system has not acquired it.

That possibility deserves investigation. But it is not, by itself, an argument against using AI.

How does someone actually become an expert?

Professional expertise is not simply the accumulation of facts. A senior lawyer knows considerably more law than a trainee, but that is only part of the difference. The experienced lawyer is also more likely to recognise which facts matter, identify an unusual risk, distinguish an apparently relevant authority and appreciate when a client's commercial objectives require a different approach.

Much of this judgement is difficult to express as a set of rules. An experienced engineer may recognise that an apparently satisfactory design has a weakness that would escape a less experienced colleague. A consultant may identify a flaw in a client's assumptions before beginning the analysis. A doctor may recognise that a combination of symptoms does not fit the most obvious diagnosis.

Such capabilities develop through the interaction of knowledge, experience, reasoning and feedback. Educational psychology has examined this process for decades.

One influential approach is cognitive apprenticeship, developed by Allan Collins, John Seely Brown and their colleagues. It describes how people may acquire complex intellectual skills by observing experts,

attempting tasks with support, receiving feedback and gradually assuming greater responsibility. It is a framework for teaching, not experimental proof that a single method invariably produces expertise.

There is an important distinction between observing a finished product and understanding how it was produced. A trainee might read an excellent legal opinion without appreciating why the author chose certain arguments, rejected others or attached particular significance to an apparently minor fact. To learn effectively, the trainee needs access to the reasoning behind those decisions.

The cognitive apprenticeship approach includes expert demonstration, coaching, supported practice and the gradual withdrawal of assistance as competence improves. These techniques aim to expose the thinking behind competent work, not just its final appearance. [1]

Deliberate practice is another important idea. It involves structured activities designed to improve performance, rather than simply repeating familiar tasks. However, we should be cautious about treating deliberate practice as a universal explanation for expertise. A metaanalysis by Macnamara and colleagues found that accumulated practice explained very different proportions of performance variation across fields. Its professional category was small, comprising only seven effect sizes from 321 participants, so striking generalisations about professions would be unjustified. [2]

Expertise is influenced by multiple factors, including prior knowledge, ability, the quality of feedback, motivation, opportunities and the environment in which skills develop. Nevertheless, there is a reasonably well-established principle here. Becoming an expert requires more than repeatedly producing acceptable work. It requires opportunities to develop, test and refine the knowledge and judgement behind that work.

The difference between useful difficulty and wasted effort

Some professional learning occurs precisely because a task is difficult. A trainee who struggles to reconcile two conflicting cases may acquire a deeper understanding of the principles involved. A junior programmer who investigates why a piece of code fails may develop knowledge that proves useful in a completely different application.

But not all difficulty is educational. Spending three hours formatting a document may teach very little. Repeatedly transferring data between spreadsheets is not necessarily a route to becoming a better accountant. Nor is there obvious value in requiring a junior lawyer to

perform a task manually when an automated system can perform it accurately and the lawyer has already mastered the underlying principles.

The distinction is between difficulty that develops capability and effort that merely consumes time. Ideally, AI would eliminate much of the second while preserving, or even improving, the first.

Research into cognitive offloading provides a useful way to draw that line. A 2016 review by psychologists Evan Risko and Sam Gilbert explains that shifting mental demands to an external aid can benefit learning when those demands are unnecessary for the particular learning objective, provided the effort saved is used productively. Applied to AI-assisted professional training, the question is therefore whether AI is removing incidental work or undertaking the reasoning that the trainee is

supposed to learn. That application is an inference from research predating modern generative AI, not an experimental finding about language models. [10]

The trouble is that organisations rarely measure those two forms of value separately. A task may appear unproductive precisely because its educational contribution is invisible. Equally, a task that occupies a trainee for several days may have negligible training value.

The answer cannot simply be to preserve traditional junior work. We need to understand which activities contribute to learning and whether AI can support that learning more effectively.

When AI produces better work but not better lawyers

In September 2026, David Autor and colleagues reported an experiment that recruited 133 practising patent lawyers at eleven American intellectual-property firms. Ninety-one completed the full protocol, including its assessment of unaided professional judgement.

Over three months, lawyers were randomly allocated access to a specialised AI patent-drafting assistant or to a control group. Google developed the tool, conducted and funded the study, and six of the seven authors were Google-affiliated. Independent patent lawyers, blinded to the participants' assignments, assessed the work using a rubric developed with Google's in-house counsel.

The immediate results were encouraging. AI access improved the quality of draft patents by an estimated 0.34 standard deviations after approximately ten days and 0.38 after ninety days. The estimated gains were numerically larger among junior lawyers, but those junior-senior differences were not statistically significant (reported p-values of 0.15 and 0.17). [3]

The researchers then examined something different. At the end of the trial, participants were asked to review and amend an existing patent application without AI assistance. This tested an aspect of professional judgement independently of AI-assisted drafting.

Overall, the AI group performed better than controls on this unaided task, but the statistically significant improvement was concentrated among senior lawyers, whose average score was 0.45 standard deviations higher. Among junior lawyers, there was no statistically significant average improvement.

The distribution of junior results was also more dispersed: there were fewer middling scores, more poor scores and more good scores, but no increase in excellent scores. The researchers considered this consistent with AI helping some juniors learn while substituting for learning in others. The study did not establish which mechanism accounted for any individual's result. [3]

There are important limitations. The 91 completers represent only part of the original 133; the researchers did not find statistically significant differential attrition between the randomised groups, but overall attrition was material. There was no baseline assessment of independent judgement, so the learning interpretation relies on randomisation rather than observing each lawyer's improvement directly. The researchers also could not fully enforce the instruction not to use AI during the final assessment, although they examined possible non-compliance. Finally, the test measured a particular patent-law skill over three months, not all dimensions of professional expertise.

The contrast remains important. Junior lawyers obtained clear benefits in assisted drafting without an average measurable gain in the subsequent unaided judgement task, whereas senior lawyers improved on both measures. Prior experience may have helped the seniors use AI more critically, but that explanation was not isolated experimentally. The study does not establish that AI made junior lawyers less capable overall, still less that it will do so throughout their careers.

The consultant who can do data science, until the AI disappears

A second experiment, published in Nature Human Behaviour on 26 October 2026, approaches the problem from a different direction. It examines whether AI enables people to undertake work outside their existing professional expertise.

Researchers from Boston Consulting Group and OpenAI, among others, investigated whether ChatGPT could help consultants undertake data-science tasks outside their established expertise. Of 986 consultants allocated to groups, 573 began the experiment and 487 completed the two-task protocol forming the principal analysed sample. The experiment compared access to ChatGPT and training in using it with other online resources.

With AI assistance, consultants performed better on coding, statistics and prediction tasks than controls, although the group as a whole did not match expert data scientists. Among participants reporting some prior coding experience, assisted results on coding and statistics could not be statistically distinguished from the expert benchmark. That is not proof of identical expertise, but it shows the practical value of combining prior knowledge with AI assistance.

When subsequently tested on a short set of technical questions without ChatGPT, the researchers found no measurable advantage in acquired technical knowledge for the AIassisted group. This does not exclude learning of other skills or learning over longer periods. [4]

A further result concerned users' judgement of the technology itself. In an exploratory analysis that was not preregistered, consultants given ChatGPT became more confident, and less accurate than controls, when judging which problems ChatGPT could solve. This was not a comprehensive assessment of their ability to supervise AI, but it raises a different concern alongside the absence of measured technical learning. [4]

There may therefore be two separate illusions: appearing capable of unfamiliar work without acquiring all the underlying knowledge, and believing one understands the tool's limits better than one does. Both were observed in this experimental setting; neither should be assumed to occur for every user.

Yet the ability to work effectively outside one's established speciality may be enormously valuable. A consultant who uses AI to perform unfamiliar technical tasks may offer services that previously required a larger team. That is a real gain in assisted capability, whether or not it amounts to unaided expertise.

But should that consultant now be described as a data-science expert? Only if we are clear about what we mean. Being able to accomplish sophisticated work with assistance is not the same as possessing all the underlying technical knowledge.

The study does not establish that consultants could never acquire that knowledge through sustained AI use. Its unaided assessment was brief and narrow, so it would be unsafe to generalise the result to every form of professional learning. Nevertheless, it illustrates why apparent capability cannot simply be equated with independent knowledge.

What happens when students use AI to learn?

In 2025, researchers led by Hamsa Bastani published a randomised field experiment involving nearly 1,000 secondary-school mathematics students in Turkey. They compared an unrestricted GPT-4 assistant, a guided AI tutor and a group receiving no AI assistance.

The guided tutor was not simply a chatbot told to give hints. Teachers provided problemspecific correct solutions, likely mistakes and suggested hints that informed its responses. That additional design effort is relevant when considering whether firms or schools can reproduce the results.

During practice, both AI groups performed better than the control group: by 48% with the unrestricted assistant and 127% with the guided tutor. Yet when AI was removed, students using the unrestricted assistant performed about 17% worse on the subsequent examination than those without AI.

The guided tutor largely eliminated that disadvantage, although its users did not score significantly higher than controls in the unaided examination. Students using unrestricted AI often requested complete solutions, while guided-tutor users were encouraged to engage with hints. The researchers also found that students using unrestricted AI did not adequately recognise the detrimental effect on their own learning. [5]

The trial concerned short-term mathematics learning among schoolchildren, not professional careers. It cannot establish what happens to a trainee lawyer after years of AI use. But it shows how apparently better completed work may conceal weaker independent performance, and why the design of the assistance deserves attention.

It would be wrong to generalise this experiment directly to professional careers. The participants were schoolchildren, the subject was mathematics and the assessment measured short-term learning. We cannot infer that a trainee lawyer using AI for several years would necessarily experience the same effect.

But AI can also accelerate genuine learning

A study by Erik Brynjolfsson, Danielle Li and Lindsey Raymond, published in the Quarterly Journal of Economics in 2025, examined an AI assistant used by 5,172 customer-support workers. It offered suggestions during customer conversations, improving average productivity by approximately 15%, especially for less experienced staff.

More importantly for the present question, researchers found evidence consistent with learning rather than mere dependence. During some periods when the AI system was

unavailable, workers previously exposed to its suggestions retained performance advantages. [6]

The inference needs care: those naturally occurring outages were not a separately randomised examination of general independent competence, were relatively rare, and might have affected the mix of customer queries. Nevertheless, they provide workplace evidence consistent with some retained learning.

The customer-support setting differs from legal, engineering or medical practice. We should not assume identical results elsewhere. The point is that improved assisted performance and improved independent capability are not invariably opposed.

Purpose-built tutors offer another route. In a 2025 crossover experiment involving 194 Harvard physics students and two lessons, a specially designed AI tutor produced better immediate posttest scores in less recorded lesson time than active classroom learning. Students also reported greater engagement. [7]

That is promising evidence of short-term learning, but the experiment did not establish that the gains persisted for months or that unrestricted chatbot use would achieve the same result.

A further experiment by Zara Contractor and Germán Reyes, published as a 2026 working paper, involved 211 Middlebury College undergraduates using ordinary generative AI tools while learning unfamiliar material and preparing an analytical essay.

Participants subsequently completed unaided knowledge tests, first immediately and then approximately one week later. The random allocation concerned access to AI, not the particular way students chose to use it.

AI access increased immediate unaided test scores by 6.7 percentage points. One week later, the advantage was still 5.1 points, about 76% of the original percentage-point improvement; the two estimated effects were not statistically distinguishable. This is substantial evidence of short-term retention, not merely improved essay output. [8]

Students who chose to use AI mainly for explanations appeared to retain more than those who used it chiefly to generate prose. But this was an observational subgroup comparison, not a randomised test of different prompting strategies; one reported estimate had a p-value of 0.061. The result suggests, rather than proves, a mechanism through which AI might assist learning. [8]

The experiment therefore supports a positive causal effect of access to AI on the measured learning outcome, while offering more tentative evidence about why that effect arose. It is a useful counterweight to claims that AI assistance necessarily erodes knowledge.

What determines whether someone learns?

One concerns the knowledge that the person already possesses. An experienced lawyer has an established understanding of legal principles, drafting conventions and professional judgement. AI may help that lawyer consider alternative approaches, identify weaknesses or examine unfamiliar material. The lawyer can place the suggestions within an existing framework of knowledge. The patent-lawyer and consultant findings are consistent with such a possibility, without proving it as a universal rule.

A trainee may lack that framework. When AI produces a sophisticated analysis, the trainee may have difficulty recognising which parts are significant, which assumptions require checking and where the reasoning might be defective.

The patent-lawyer experiment cannot establish that prior expertise alone explains the difference between senior and junior outcomes. Nor should an absence of average learning among juniors obscure the variation between individuals.

Another factor is the nature of the interaction. A person who asks AI to explain a concept, challenges the explanation and then attempts a problem independently is engaged in a different activity from someone who requests a complete answer and copies it.

A third factor is effort. In an October 2026 revision of a randomised experiment involving 1,174 adults in Argentina, Cruces and colleagues found that AI assistance narrowed an educationrelated performance gap on a business problem-solving task by about 75%.

Participants then undertook a related assessment without AI immediately after the assisted task. The lower-education group's estimated gain was 0.171 standard deviations; the corresponding higher-education estimate was not statistically significant. Crucially, the difference between those two effects was not itself statistically significant. In a separate exploratory analysis, stronger immediate unaided performance was associated with combining intensive AI use and sustained personal engagement, but that association does not establish causation. [9]

This is evidence of immediate carry-over, not proof that skills were retained over weeks or months. AI need not eliminate effort: it may redirect effort towards different activities. Its educational value may depend on what the individual does with the time and mental capacity that assistance releases.

The psychological problem: knowing when you do not know

There is another difficulty. To use AI effectively, a person often needs to recognise when its answer may be inadequate.

That requires some understanding of the subject. In a familiar area, an experienced professional may quickly identify a doubtful assumption or an important omission. In an unfamiliar area, the same professional may be less able to distinguish a convincing explanation from a correct one.

The problem is not unique to AI. Psychologists use the term metacognition to describe our ability to monitor and evaluate our own knowledge and thinking. We do not always judge those abilities accurately.

In 2016, Evan Risko and Sam Gilbert reviewed research into cognitive offloading, the use of external aids to reduce the mental demands of a task. Remembering an appointment through a calendar is a simple example; calculating with a spreadsheet is another. Such assistance is not inherently harmful and may free resources for more demanding work.

However, decisions about what to offload depend partly on our assessment of our own abilities. If those assessments are inaccurate, we may rely on external support inappropriately. [10]

Risko and Gilbert also examine evidence that external access to information can change not only what people remember but what they believe they know. In some experiments, people who had searched online overestimated their own unaided knowledge. In others, people expecting access to the internet became more willing to acknowledge that they did not know an answer. The effect was not uniform: access to external knowledge could either inflate or temper confidence, depending on the circumstances. What matters for the present argument is the possibility of confusing knowledge held personally with information obtainable through a tool. Whether generative AI produces the same effect, and in which settings, requires separate evidence. [10]

Generative AI introduces an additional complication because it can undertake much of the reasoning involved in a task, not merely store information or carry out calculations. A 2025 Microsoft Research survey collected 936 examples of AI-assisted tasks from 319 knowledge workers. It examined reported critical-thinking behaviour and confidence, not objective longterm changes in cognitive ability.

Greater confidence in AI was associated with less self-reported critical-thinking activity, whereas confidence in one's own abilities was associated with more. Participants also described shifting effort from producing material towards checking and integrating it. These are associations and self-reports, not proof that AI causes lasting cognitive decline. [11]

The concern is therefore not that every use of AI makes people less thoughtful. It is that a trainee may become more confident because AI repeatedly produces satisfactory results without necessarily developing an equivalent ability to judge when they are wrong. The consultant experiment's exploratory finding about misjudging AI capabilities is a direct, though setting-specific, illustration.

The automation problem is older than generative AI

In 1983, psychologist Lisanne Bainbridge published Ironies of Automation, examining industrial control systems rather than language models. She observed that automation could remove opportunities to practise skills while leaving operators responsible for intervening when the system failed.

The result was an irony: the operator might need the greatest expertise precisely when routine automation had provided the fewest chances to develop or maintain it. Even reliable automation could make failures rarer and training more important. [12]

A lawyer faces a related question when AI routinely identifies authorities, drafts analysis and highlights risks. If the system misses something important, the lawyer must still recognise the omission. But that ability cannot safely be assumed merely because previous AI-assisted work looked satisfactory.

Bainbridge did not argue that every operation should remain manual. Her point was to design automation around the capabilities people must retain. Applying her industrial analysis to modern generative AI is an analogy, not direct experimental evidence about lawyers. It nonetheless helps explain why training objectives should influence decisions about what to delegate.

Does everyone need to remain capable without AI?

There is an objection to much of this argument. Why should an individual be required to perform every task independently when reliable technology can do it better?

We no longer expect accountants to perform every calculation manually. Engineers routinely use sophisticated design software. Doctors rely on diagnostic equipment. Lawyers use electronic databases rather than searching through shelves of printed law reports.

Professional expertise has always depended partly on external tools and the knowledge of other people. AI may simply extend that process.

This objection has considerable force. It would be absurd to insist that professionals reproduce the working methods of previous generations merely to demonstrate that they can.

The question should not be whether someone can perform every task without assistance. It should be which capabilities they need to possess personally in order to use assistance safely, recognise its limitations and take responsibility for the result.

For some activities, independent competence may remain essential. For others, knowing how to formulate a problem, select appropriate tools, evaluate their output and integrate specialist contributions may be more valuable than performing every intermediate step.

Professional expertise may therefore change rather than disappear. A future senior lawyer may spend less time producing initial drafts and more time identifying the real dispute, assessing uncertain evidence, advising on strategy and managing legal and commercial risks.

But that creates another question. How does someone develop those higher-level capabilities?

It is tempting to say that trainees should simply move directly into more sophisticated work. Yet sophisticated work frequently requires an understanding of the principles and practical difficulties that earlier tasks helped develop.

The difficulty is not necessarily that the old tasks must survive. It is that their educational contribution must be recognised before they are removed.

A different way of thinking about AI-assisted learning

The research does not justify a universal conclusion that generative AI makes people less capable. Nor does it justify assuming that greater productivity necessarily means faster professional development.

What emerges is a distinction between three different outcomes. Assisted performance concerns what someone can accomplish while using AI. Independent competence concerns the knowledge and judgement they can exercise without it. Learning concerns whether AI use improves those underlying capabilities over time.

These outcomes can move in different directions. An individual may become substantially more productive without learning much. Another may use AI to acquire skills more quickly than traditional training would permit. A third may experience both benefits.

The distinction is not an entirely new theory: it follows from established principles of learning psychology and research into human interaction with technology. It has a practical implication. Organisations assessing AI should distinguish improvements in completed work from improvements in the people producing it. A better document, piece of code or report does not,

by itself, show that the author has developed better judgement. The two outcomes require separate evaluation.

So, is AI creating an illusion of expertise?

Sometimes, yes. Experiments involving patent lawyers and management consultants show that AI-assisted work can improve without a corresponding average gain on the particular unaided assessments researchers used. Neither experiment proves that no other learning occurred or that long-term competence was damaged.

Other research shows a more encouraging possibility. The mathematics experiment found that unrestricted AI can undermine a short-term unaided assessment, whereas guided assistance largely avoided that result. The customer-support and undergraduate studies provide evidence consistent with learning gains that persist after the tool is removed.

The lesson is not to reject AI assistance, but to design and evaluate it for more than immediate output. Productivity gains are valuable; they simply do not answer every question about professional development.

The strongest conclusion is therefore not that AI inevitably undermines expertise. It is that expertise does not necessarily develop as a by-product of using AI to produce better work.

The development of expertise depends on the interaction between the person's existing knowledge, the nature of the task, the assistance provided and the opportunities to practise and receive feedback.

This brings us back to the two trainee lawyers. Perhaps the first learnt more from struggling with the problem. Perhaps the second learnt more because AI provided useful explanations and exposed weaknesses that would otherwise have remained unnoticed. The finished documents cannot establish which outcome occurred.

That uncertainty is a reason to reconsider how employers supervise, assess and develop junior staff. The challenge is not to preserve every old task, but to ensure that apparent expertise is not mistaken for the capability organisations will eventually need.

AI may allow us to produce expert-looking work long before we have developed expert judgement. The challenge is to ensure that the appearance of expertise does not become a substitute for acquiring it.

In Part 3, we will examine what this means for employers: how they should balance immediate productivity with the longer-term responsibility of developing the professionals on whom their businesses will depend.

Alex Asks

If I can help you produce the work of an expert, have I helped you become one? Perhaps. But shouldn't we find out whether you have learnt something, rather than simply admiring what we produced together?

References

[1] Collins, A., Brown, J. S. and Holum, A. (1991). Cognitive Apprenticeship: Making Thinking Visible. American Educator, 15(3), 6–11, 38–46. Original article published by the American Federation of Teachers. Read the original article

[2] Macnamara, B. N., Hambrick, D. Z. and Oswald, F. L. (2014). Deliberate Practice and Performance in Music, Games, Sports, Education, and Professions: A Meta-Analysis. Psychological Science, 25(8), 1608–1618. DOI: 10.1177/0956797614535810. Read the research abstract

[3] Autor, D., Rodchenko, T., Martin, J., Iscenko, Z., Strand, S., Pearl, D. and Ferere, M. (2026). Does AI Assistance Enhance or Erode Expertise? Evidence from a Three-Month Field Experiment in Patent Drafting. NBER Working Paper 35720, September 2026, revised October 2026. DOI: 10.3386/w35720. Read the NBER research

[4] Wiles, E. and colleagues (2026). Generative AI and the Temporary Upskilling of Knowledge Workers. Nature Human Behaviour, published 26 October 2026. The reviewed text was the authors' prepublication manuscript. DOI: 10.1038/s41562-026-02566-6. Read the published research

[5] Bastani, H., Bastani, O., Sungu, A. and colleagues (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. Read the full article

[6] Brynjolfsson, E., Li, D. and Raymond, L. (2025). Generative AI at Work. Quarterly Journal of Economics, 140(2), 889–942. DOI: 10.1093/qje/qjae044. Read the full article

[7] 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. Read the full article

[8] Contractor, Z. and Reyes, G. (2026). Experimental Evidence on the Learning Impact of Generative AI. IZA Discussion Paper 18792, July 2026. Read the research

[9] Cruces, G., Fernández Meijide, D., Galiani, S., Gálvez, R. H. and Lombardi, M. (2026). Does Generative AI Narrow Education-Based Productivity Gaps? Evidence from a Randomized Experiment. NBER Working Paper 34851, revised October 2026. DOI: 10.3386/w34851. Read the NBER research

[10] 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. Read the research

[11] Lee, H.-P., Sarkar, A., Tankelevitch, L. and colleagues (2025). The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers. Proceedings of CHI 2025. DOI: 10.1145/3706598.3713778. Read the original paper

[12] Bainbridge, L. (1983). Ironies of Automation. Automatica, 19(6), 775–779. DOI: 10.1016/0005-1098(83)90046- 8. Read the original paper