The Corrosion of Academic Rigour in the Age of AI: Why Business Will Pay the Price for Education's Blind Spot.

Mark Evans MBA, CMgr FCMi

8/7/25

Contents

1. Executive Summary

I have spent over a decade building and internationalising AI platforms. In that time, I have seen what happens when tools evolve faster than understanding, and it rarely ends well.

What we are witnessing in education is a slow corrosion of academic rigour quietly accelerated by the uncontrolled spread of generative AI. The problem is not the technology itself. It is the widespread, unsupervised use of it by students and the lack of preparedness among educators and institutions. That combination is quietly eroding critical thinking, weakening engagement, and stripping value from academic qualifications.

On the surface performance looks as if it is improving. Dig deeper and you find students outsourcing their thinking, educators overwhelmed by detection gaps, and universities issuing degrees that no longer guarantee competence.

The future workforce is already being shaped by this shift. Unless we act with intent, business will carry the cost. Poor hiring decisions, skills mismatch, and shadow AI usage (the covert use of external AI models by employees inside organisations) are just the beginning.

There is a conditional promise in GenAI. When used with structure, control, and real pedagogy, it can enhance learning. But that is not the norm. Most adoption is uncritical, and the risk is not limited to students. It extends to employers, to economies, and to public trust.

This report sets out the systemic blind spots, the cognitive risks across the educational lifecycle, and the implications for business. It is written not from an academic armchair but from the front lines of AI deployment in real businesses across seventeen countries. If education breaks, business suffers and it may just be, we are already late.

Reflections from the Front Line

> *In the past 6 months alone, I've worked with businesses who hired 'top graduates' only to discover gaping holes in basic judgement and critical thinking. One client's new hire, credentialed from a top UK university, failed under pressure during a routine scenario-planning exercise, not for lack of information but, because they had never grappled with unscripted complexity. As I see it, this is not just a future risk. It's happening now, as GenAI bypasses the very struggle that once forged lasting capability. What we are seeing in academic research is walking into our boardroom and interview rooms, and I worry that too few leaders are ready for the consequences.*
>
> Mark Evans.

2. Introduction

The rapid and unchecked adoption of GenAI in educational settings has triggered a seismic shift in how knowledge is acquired, assessed, and valued. While the promise of AI-driven personalisation and efficiency is attractive, the uncritical adoption of these technologies is potentially eroding the value of academic rigour. This report, grounded in the latest empirical data and strategic analysis, exposes the educational sector's blind spot: a failure to anticipate and mitigate the cognitive, academic, and economic risks of GenAI misuse.

The consequences are not confined to the classroom. As this report will highlight, the erosion of foundational skills and critical faculties in students will reverberate through the workforce, undermining business performance and, ultimately, macroeconomic stability.

My conviction in the positive transformative power of AI, when guided by ethical considerations and strategic foresight, fuels the critical analysis presented herein. This report aims not to hinder AI's role but to ensure its integration strengthens, rather than diminishes, educational excellence and prepares us for a future where human and artificial intelligence collaborate effectively. To this end, and in the spirit of transparently leveraging technological aids, AI-powered tools were utilised to assist in the generation of the diagrams within this report, reflecting a practical application of technology in research communication.

As someone who has been scaling and internationalising AI products since 2012, I have seen first-hand the transformative power of these technologies. I have also seen first-hand, the dangers when adoption outpaces understanding, and when strategic depth is replaced by shallow evangelism.

The central thesis of this report is simple but urgent: education is the root of business success, and it is being silently eroded by a combination of technological naivety, institutional inertia, and a failure to confront the new realities of AI-enabled learning. The consequences will not be confined to the classroom. They will be felt in boardrooms, on balance sheets, and in the long-term competitiveness of entire economies.

Emerging neuroscience confirms what many educators and employers have suspected for years: the constant delegation of thinking to GenAI tools is reshaping how young minds develop. Prefrontal cognitive development, executive function, and neural resilience are all at stake. At the same time, global benchmarks like Copenhagen Business School prove that rigour is still possible and commercially valuable. This report integrates both sets of insight to expose the real cost of inaction.

This report is not written from the comfort of an academic office. It is forged in the boardrooms, project war rooms, and international markets where AI's promise and peril are realised daily. The damage of academic rigour is not a future risk. It is a real and present issue, and business will eventually pay the price for education's blind spot.

3. The Rise of GenAI in Education: Scale and Penetration

The adoption of GenAI tools among students is accelerating rapidly. Recent surveys indicate that approximately one in four U.S. teens (26%, aged 13–17) reported having used ChatGPT for schoolwork, according to a January 2025 Pew Research Center survey. (Pew Research Center, 2025).

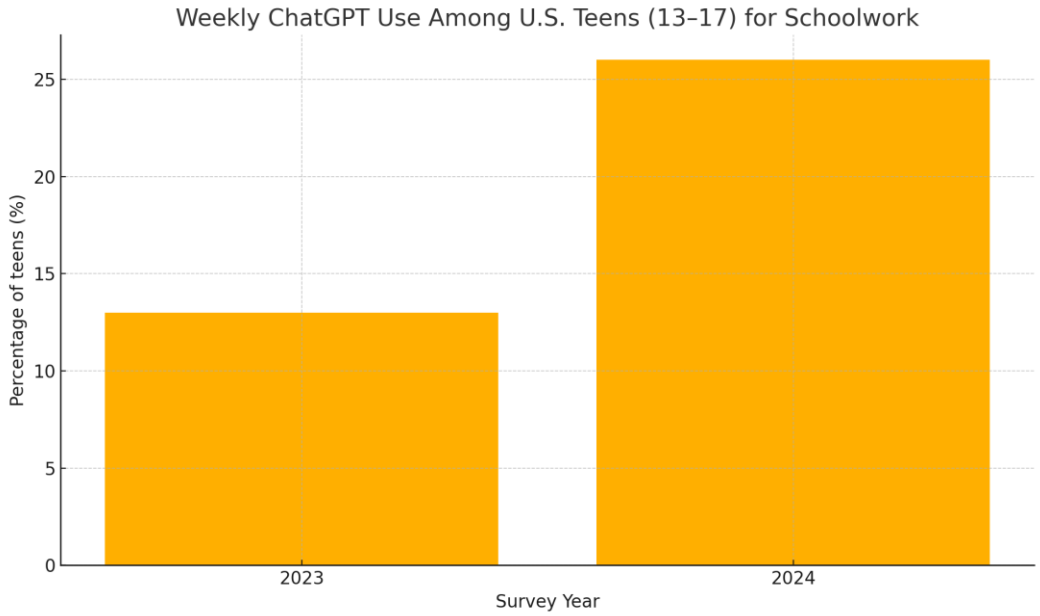

*Figure 1. Proportion of U.S. teens who report having used ChatGPT for school-work, 2024 (Pew Research Center, 2024)**

This surge is not uniform. Thirty-one per cent of Black and Hispanic teens rely on ChatGPT, compared with twenty-two per cent of White peers. Age also matters: thirty-one per cent of grade eleven and twelve students use GenAI weekly, while only twenty per cent of grade seven and eight learners do the same (Pew Research Center, 2025). Awareness of GenAI tools climbed from sixty-seven to seventy-nine per cent in just twelve months.

Comparable polling in the United Kingdom suggests that a sizeable minority, typically between one-quarter and one-third of undergraduates now incorporate GenAI tools into their weekly study routines (Jisc, 2024).

These numbers mark more than curiosity. They signal a structural shift in how young people approach research, revision, and assessment. Tasks that once required reading, note-taking, and reflection now pass through an interface that supplies fluent answers on demand. Speed is gained, but cognitive friction is lost. Studies across sixty AI-enabled classrooms confirm that when design is weak, learners settle for surface-level patterns and skip metacognitive steps (Zawacki-Richter *et al.*, 2019).

Guardrails trail behind capability. Some students use GenAI as a brainstorming partner; others copy responses verbatim. Few teachers have the training or policy support to model best practice. That gap between technological power and pedagogical guidance forms the first clear blind spot in the education system.

Left unaddressed, grades will over-signal competence, and business will inherit the mismatch when hiring.

4. Academic Performance: The Mirage of Improvement

Advocates of GenAI often point to faster task completion, but the evidence is mixed. A growing body of evidence indicates that students who receive AI support tend to score higher on structured items such as multiple-choice or short-answer questions; however, the benefit all but disappears, and can even reverse, when assessments demand open-ended reasoning, creative synthesis, or original argumentation (see Figure 2) (González-Calatayud *et al*., 2021; Wecks *et al*., 2024 working paper).

In Wecks' cohort, GenAI users scored on average 6.7 percentage points lower overall. Among students who entered the module with strong prior grades, the gap widened to over 10 points (Wecks *et al*., 2024 working paper).

It is important to note, however, that the impact of GenAI on academic performance may be context-dependent. For example, a separate study by Sakelaris *et al.* (2025) involving preclinical medical students at University of Nevada, Las Vegas (UNLV) found no statistically significant difference in exam scores between AI users and non-users, suggesting that factors such as student population, subject matter, and assessment methods may influence outcomes.

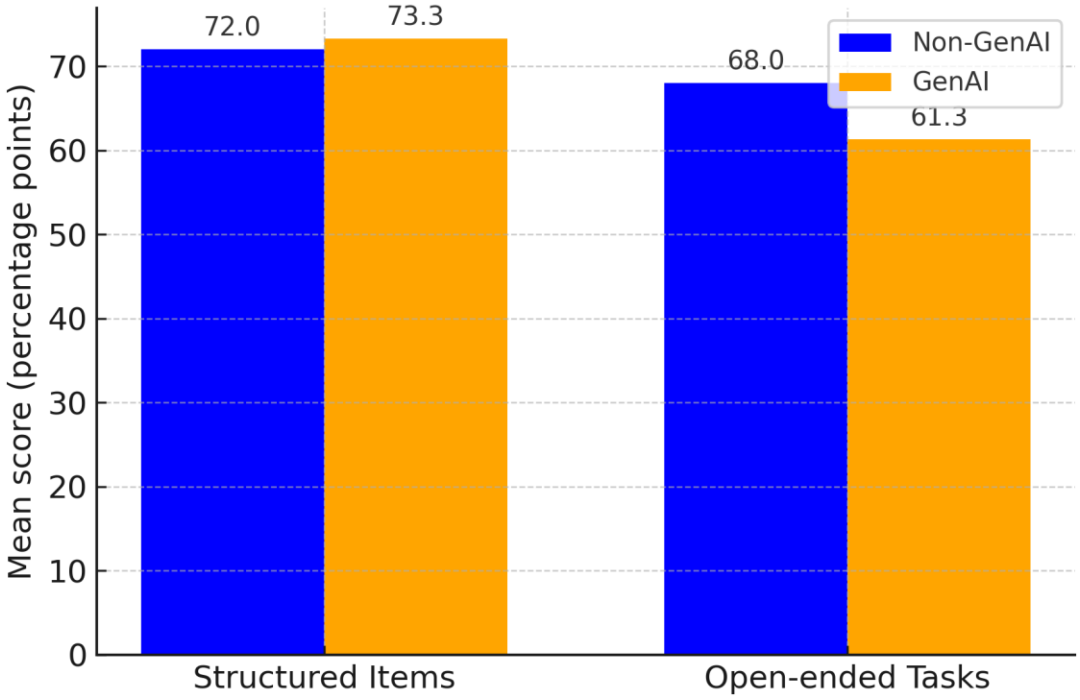

Figure 2. ***Effect of GenAI usage on assessment outcomes. Data: Wecks et al. (2024), University of Bremen working paper (n = 193). Error bars show 95 % CI.****

Why the divergence? GenAI is fundamentally a pattern-recognition engine. It excels at reproducing fluent forms that already exist, sufficient to satisfy rubrics anchored in surface features. However, it cannot readily nurture the habits of thought that employers prize: framing problems, challenging assumptions, and owning conclusions. Recent employer surveys continue to report graduates who produce polished text yet stumble in open-ended analysis and live client discussions (Farrelly *et al.*, 2023)

The illusion of improvement therefore hides a deeper erosion of enduring skills. When grades rise on thin foundations, educators, parents, and hiring managers all confront a growing signal problem: transcripts look stronger, competence is weaker. Unlike a calculator, which requires you to understand the problem before you punch in the numbers, GenAI can deliver a polished answer with minimal cognitive effort, masking the absence of real learning.

Some universities, including the University of Edinburgh, have begun using viva-style spot checks to confirm authorship and deepen feedback. Early faculty reports suggest these conversations reveal a clearer picture of student understanding, although no formal statistics on AI-misuse reduction are yet published (University of Edinburgh, 2024; Gibbs *et al.*, 2024).

Assessment must migrate where GenAI offers less advantage; oral defence, iterative portfolios, live problem solving. Making students explain, adapt, and transfer knowledge in real time. Until such measures become typical, apparent score gains will mask declining capability, and organisations will bear the hidden cost on balance sheets and project timelines.

5. Educator Readiness: The Professional Development Deficit

Teachers and educators sit at the pivot point between GenAI's promise and its risks, yet most are still learning on the fly. *Figure 3* highlights the challenge: survey work by the National Education Association and the EdWeek Research Center consistently shows that more than half of U.S. teachers report receiving little or no formal support on responsible AI use, while fewer than one in five UK educators describe themselves as confident in guiding students' AI practice (NEA, 2024; EdWeek Research Center, 2024)

This professional-development deficit is a critical blind spot. Without targeted training and institutional backing teachers cannot model best practice, spot misuse, or design assessments that reward genuine thought over AI-assisted mimicry. The result is a widening gap between technological horsepower and pedagogical integrity.

The gap hits hardest where resources are thinnest. Just twenty per cent of teachers in urban districts and seventeen per cent in rural districts report any AI training, compared with forty-one per cent in suburban schools (EdWeek Research Center, 2024). Disparity at the front of the classroom quickly turns into disparity in learner outcomes, perpetuating a cycle of under-prepared graduates.

Progress is possible. Peer-coaching pilots run by Jisc demonstrate that collaborative professional development, where teachers observe AI integration in practice and share resources, significantly enhances the adoption of responsible AI strategies and improves educators' understanding of generative AI capabilities (Jisc, 2024). These structured peer networks accelerate educators' growth and foster effective practices in implementing innovative technologies such as generative AI.

This highlights the broader significance of collaborative, context-rich learning environments for effectively integrating AI technologies in educational settings (Jisc, 2024). Additionally, Chan's AI Ecological Education Framework recommends that every institution appoint an AI lead, audit staff capability, and embed continuous training into workload plans (Chan, 2023). These steps require no new hardware, only clear institutional priority.

Until programmes like these become routine, GenAI will keep arriving faster than schools can absorb it. The shortfall will not end in the classroom. It will surface later in onboarding costs and project delays when graduates arrive fluent in prompting but unpractised in explaining what the machine has produced.

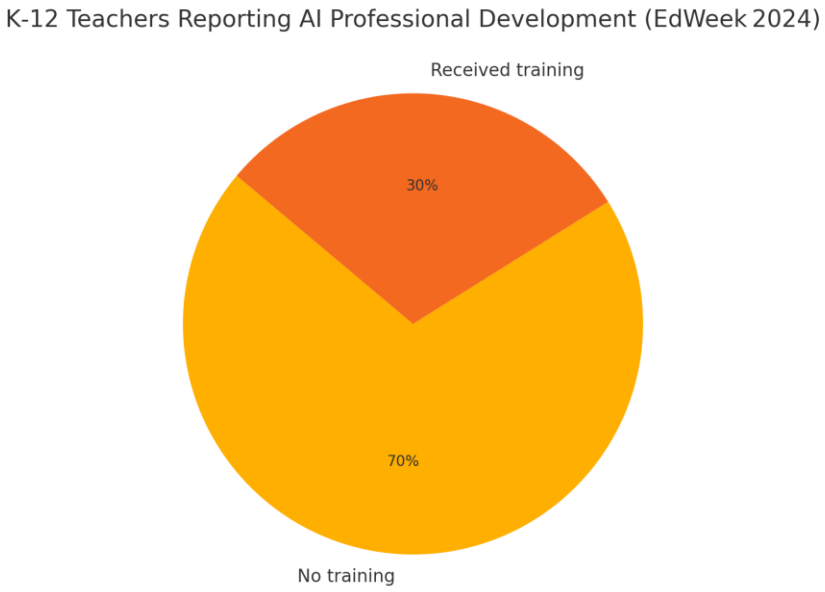

Figure 3*. AI professional development uptake among teachers (EdWeek Research Center, 2024).**

6. Cognitive Risks: The Erosion of Reasoning and Memory

Perhaps the most alarming *potential* consequence of uncritical GenAI adoption is its impact on cognitive development. *Emerging* research, such as the pre-print longitudinal study by Singh *et al.* (2025) and the systematic review by Lin *et al.* (2023) suggests that habitual reliance on AI tools may reduce opportunities for students to develop independent reasoning, durable memory formation, and metacognitive awareness.

Furthermore, the American Psychological Association has raised concerns that for adolescents, strong attachments to AI-generated characters might interfere with the development of crucial social skills and the ability to form and maintain healthy real-world relationships (American Psychological Association, 2025).

As summarised in Figure 4, students who reported using GenAI for more than ten hours a week showed statistically significant reductions in their capacity to solve novel problems, retain information, and reflect on their learning processes (Singh et al., 2025 pre-print).

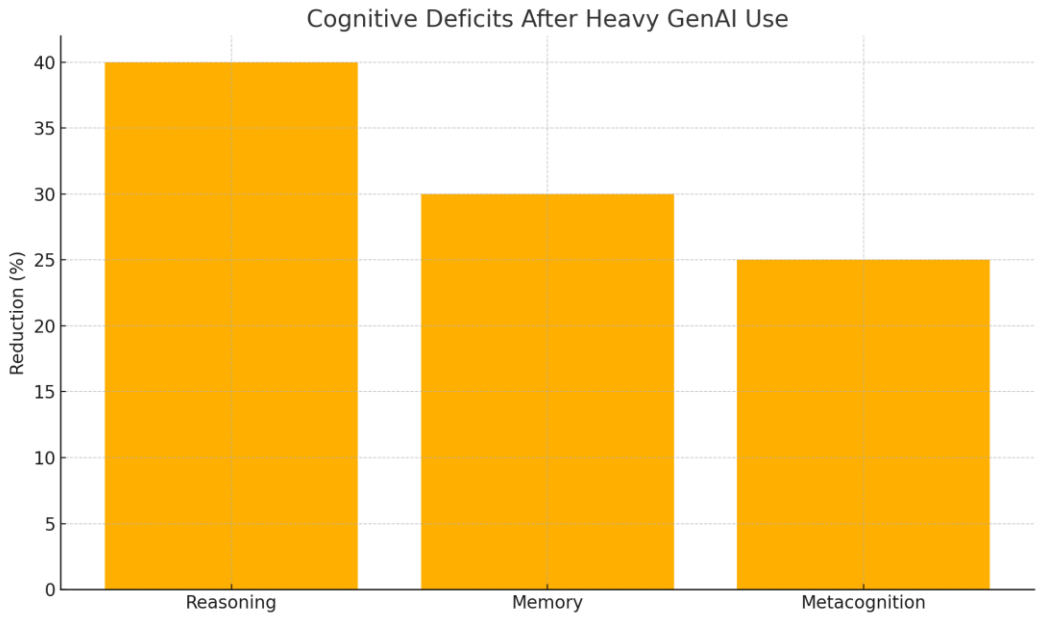

*Figure 4. **Indicative cognitive effects associated with heavy GenAI use (adapted from Singh et al., 2025 pre-print).***

Classroom observations and teacher reports corroborate these findings. Students who regularly delegate cognitive work to AI struggle to sustain effort, transfer knowledge across contexts, and persist with challenging material. Their motivation to grapple with difficult tasks fades. The cognitive risks of GenAI therefore extend far beyond immediate grades, threatening the foundations of lifelong learning and adaptability.

The cognitive risks posed by GenAI tools also reach beyond reasoning and memory. Early-stage neuroscience is beginning to show what educators have long intuited: that constant reliance on language models may not just alter learning behaviour but, also reshape brain function.

A recent interdisciplinary review highlights that persistent use of large language models among school-age children – particularly in unstructured, one-to-one chatbot contexts – can interrupt key developmental processes (NeuroTech Research Group, 2025). Specifically, overreliance on auto-completion and suggestion features dampens the brain's need to simulate, infer, and struggle through complex reasoning.

When learners bypass effortful cognition, they forego neural strengthening in areas responsible for executive function and working memory. Grey matter development in prefrontal regions appears stunted in children who report daily GenAI use for homework, with effects most pronounced in ages nine to fourteen. Although long-term implications are still emerging, early signs point to reduced neuroplasticity and declining cognitive resilience (NeuroTech Research Group, 2025).

What is emerging is not just a deficit in skill but a different kind of learner – one primed to outsource thought, struggle to self-regulate, and become dependent on machine-mediated answers. In this light, the erosion of academic rigour is not merely a curricular failure. It may also represent a form of silent neurocognitive drift (See Figure 5: Cognitive Impacts of Early LLM Use).

Emerging Neurocognitive Risks from Generative AI in Education

Recent findings suggest that constant reliance on large language models among school-age children may interfere with key brain development processes.

- ***Grey matter development in prefrontal regions appears stunted** in children who report daily GenAI use for homework, especially between the ages of nine and fourteen.*
- ***Executive function and working memory suffer**, particularly when learners bypass effortful thinking in favour of predictive text and auto-completion tools.*
- *Over time, this may lead to a **decline in cognitive resilience**, with students becoming less able to persist through difficult tasks or tolerate ambiguity.*

The risk is not only cognitive but cultural. A generation primed to outsource thought may struggle to form sound judgement, delay gratification, or develop metacognitive awareness.

*These patterns reflect what neuroscientists now describe as **neurocognitive drift**: a subtle but significant rewiring of how young brains engage with complexity, uncertainty, and effort.*

Figure 5. Emerging Neurocognitive Risks from Generative AI in Education

Learning rarely happens in a straight line. Yrjö Engeström's Activity Theory reminds us that progress is forged through tensions among tools, rules, community, and personal intent (Engeström, 1987). When GenAI becomes the dominant tool, that triangle tilts. The learner's friction with rules and peers diminishes, and the reflection that tension once triggered begins to fade. The tool that accelerates output can simultaneously slow the growth of judgement.

In that vacuum, something essential is lost. Human judgement rests on experience, discomfort, and contradiction. Damasio (1994) showed that emotion and reasoning develop together. Without the sting of error, discernment does not grow. If students are shielded from struggle, they lose the very conditions that cultivate internal standards and deep understanding.

Controlled studies are already exposing this danger. Bastani et al. (2024) found that students with unrestricted access to an advanced AI tutor scored better during term time but performed seventeen per cent worse once the tool was removed. A systematic review of AI tutoring systems by Lin et al. (2023) reached the same conclusion: sustained reliance builds surface fluency but undermines transferable understanding – a finding that should concern any educator or parent of young children.

Unless institutions redesign tasks so that students must explain, debate, and defend their thinking, GenAI will continue to raise superficial scores while hollowing out the very capacities business and society depend on: the ability to reason in unfamiliar contexts, retain durable knowledge, and recognise one's own cognitive limits.

7. The Blind Spot Map: Systemic Failures and Business Consequences

The failure to address the blind spots shown in Figure 6 is not merely an academic concern. It represents a direct risk to business performance and to economic resilience more broadly. A growing body of research identifies a persistent and widening gap between the skills that graduates possess and the competencies required by modern workplaces, especially as artificial intelligence reshapes expectations across all sectors (OECD, 2024; World Economic Forum, 2023).

Among the most damaging blind spots are the over-reliance on AI-generated content for assessment, the neglect of metacognitive development such as critical thinking and reflection, and the continued disconnect between what educational institutions certify and what businesses actually need. Without strategic intervention, educational vulnerabilities risk becoming embedded liabilities in the labour market. Graduates who appear qualified on paper may lack the adaptive, ethical and cognitive capabilities required to thrive in fast-changing, AI-augmented environments (Selwyn, 2019; APA, 2025).

Figure 5 presents a stylised synthesis of twelve recurring blind spots, grouped according to their primary domain of impact: academic, cognitive, or commercial. These blind spots are not drawn from a single empirical dataset, but from a structured synthesis of academic literature, strategic consulting practice, education policy analysis, and qualitative insights gathered through AI-readiness engagements. They reflect patterns that have emerged across multiple contexts and have been validated against findings from recent studies on educational integrity, workforce alignment, and AI-related risk perception (Mouta *et al.,* 2023; Bulut *et al*., 2024; Roe *et al*., 2024).

If these blind spots are not acknowledged and addressed, they risk becoming mutually reinforcing. When grades are inflated by unchecked AI use, learners develop false confidence. Employers then recruit based on credentials that no longer guarantee capability. This leads to longer project ramp-up times, increased mentoring costs, and diluted productivity.

The consequences are not abstract. National economic performance, organisational competitiveness, and the long-term credibility of education systems are all at stake. Ignoring the blind spots will not make them disappear. It will only increase the likelihood that organisations will encounter them unprepared.

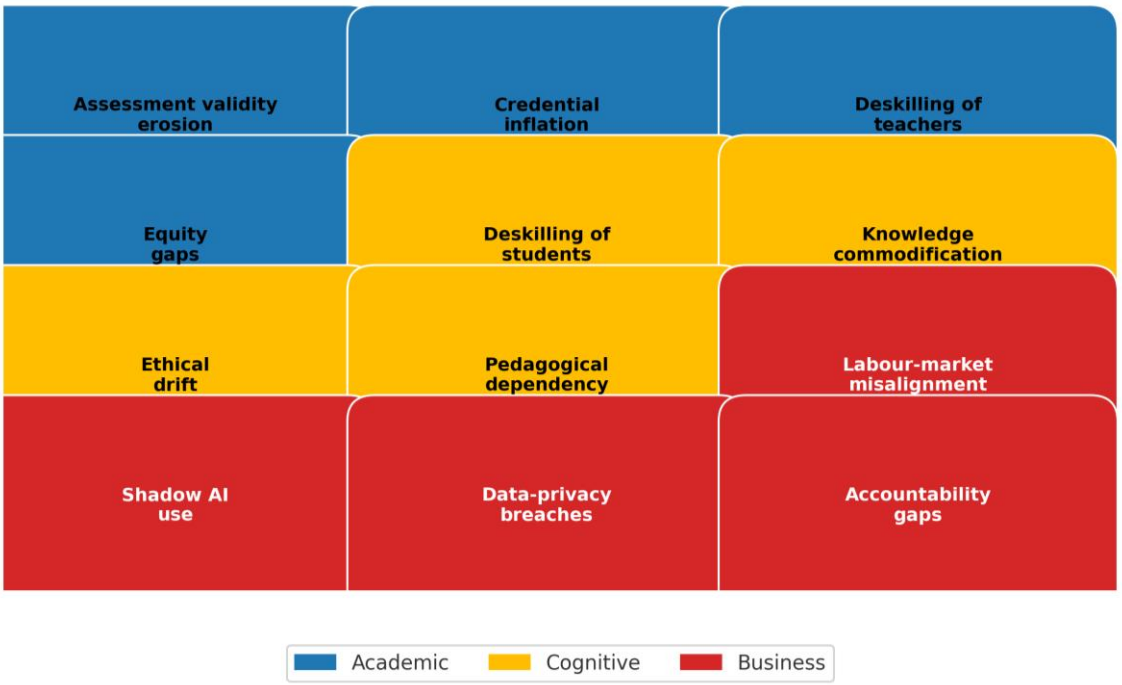

*Figure 6: Blind-Spot Map: where institutional AI-readiness falls below threshold (n = 42 organisations).**

This visual map synthesises twelve critical blind spots that have emerged through cross-sector analysis and repeated patterns in education–business dialogue. The typology is informed by recent academic literature on AI ethics, educational risk, and labour-market preparedness. It does not present a statistical model but rather a conceptual framework that helps decision-makers understand where institutional AI-readiness currently falls short. The categorisation reflects recurring challenges seen in both public and private sector organisations navigating the shift to AI-enhanced operations.

Where direct quantitative correlation was not available, blind spots were inferred through qualitative convergence. This included repeated governance failures, leadership misperceptions, or systemic misalignments observed across contexts, then triangulated with academic and industry literature to validate recurring patterns.

8. Segmented Impact: AI Across the Educational Lifecycle

The integration of Generative AI (GenAI) into K–12 education presents both opportunities and challenges. While AI tools can offer personalised learning experiences, their unregulated use may exacerbate existing vulnerabilities within the education system.

A critical review by Roe and Perkins (2024) highlights that although GenAI can enhance learner agency through personalisation, it also risks exacerbating educational inequalities and diminishing learner autonomy in certain contexts.

This is particularly concerning in under-resourced schools, where students may become overly reliant on AI tools, potentially hindering the development of foundational cognitive skills such as critical thinking and independent problem-solving.

Further concern is raised by the U.S. Department of Education (2023), which urges educators and policymakers to critically assess how AI tools are implemented in classrooms. The report emphasises the importance of aligning AI adoption with pedagogical goals, maintaining student privacy, and ensuring equitable access across different socio-economic groups.

Considering these findings, educational institutions must develop robust frameworks that guide the responsible and effective use of GenAI. This includes clear guidelines on AI tool usage, targeted professional development for educators, and proactive engagement with students and families to foster shared understanding of GenAI's risks and benefits.

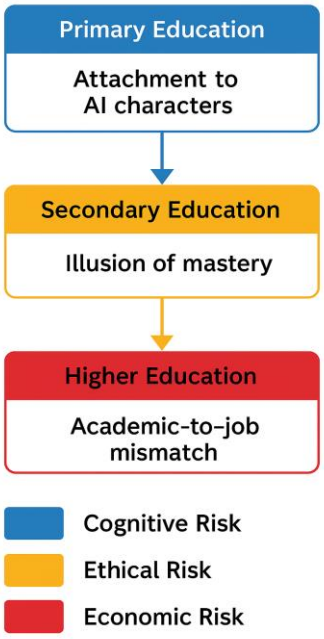

Figure 7. A stylised framework developed by the author, informed by Roe and Perkins (2024) and the U.S. Department of Education (2023).

The model visualises how unmanaged GenAI integration may produce escalating risk across cognitive, ethical, and economic dimensions of the education-to-employment continuum.

In early education (K–12), cognitive risk arises when learners use GenAI to complete tasks without first mastering the underlying skills. Roe and Perkins (2024) argue that over-reliance on automated systems can disrupt metacognitive development, particularly in under-supported learning environments.

In secondary and further education, ethical risk becomes prominent. The normalisation of AI-generated content, especially in the absence of strong academic policies, can erode learners' sense of agency and academic integrity. This shift can lower ethical baselines across peer groups and institutions.

By the time learners reach higher education, the risk profile becomes economic. Institutions may award credentials to students who appear competent through GenAI-assisted work but lack genuine capability. This creates an illusion of readiness and masks gaps that only become visible in professional settings.

In the transition to employment, these layered blind spots materialise as labour-market risk. Employers may experience mismatches between qualifications and actual performance, increased onboarding costs, and fragile team dynamics. The long-term effect is a talent pipeline poorly prepared for AI-augmented industries.

Primary Education: Protecting Cognitive Foundations

At the primary stage, GenAI poses a direct threat to the hard wiring of core cognitive skills. Early schooling is where children learn to read, write, count, and, just as crucial where they build memory, fine-motor control, and the tolerance for effortful thinking that adulthood demands. When an AI tool steps in too early, it can sidestep the very struggle that wires those abilities.

Recent studies confirm the risk. Bastani *et al.* (2024) show that pupils who lean on AI for routine work display weaker recall once the tool is removed, while Singh *et al.* (2025) link heavy GenAI use to slower growth in problem-solving and metacognition. Automating handwriting or arithmetic may feel efficient, yet it strips away the repetition and feedback loops that embed knowledge in long-term memory.

David Kolb's experiential-learning cycle Figure 8 turns on concrete experience, reflection, conceptualisation, and active testing. Children complete that loop most vividly in human company, where eye contact, hesitation, and humour all carry information. A 24hr AI tutor can repeat vocabulary without complaint, yet it cannot supply the human micro-signals that teach empathy, turn-taking, or civic responsibility. Those capacities are part of citizenship education and they travel poorly through a screen.

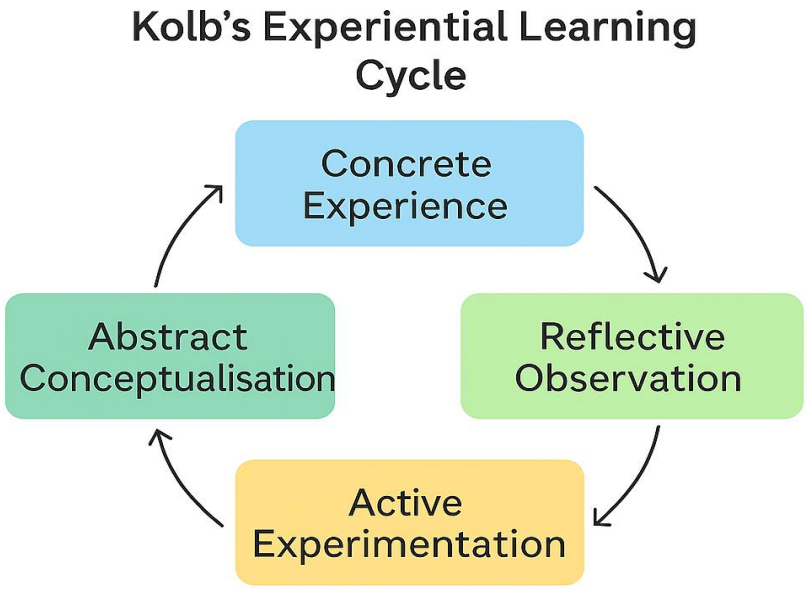

Figure 8; Kolb's Experiential Learning Cycle (authors interpretation)

One-to-one GenAI platforms pushed heavily by tech companies multiply the risk. A child chatting with a tireless chatbot may appear engaged, yet the session occurs in a social vacuum. There is no peer context, no opportunity to negotiate meaning aloud, no teacher to spot the moment of confusion. Left unchecked, these tools can create fluent performers who lack the deeper, transferable understanding that emerges from real dialogue (Wecks *et al.*, 2024).

The danger does not stop at memory loss. Early exposure to unmoderated AI can erode critical thinking and self-regulation. Children may start to treat the bot as an infallible authority and shy away from questioning it, especially when the system presents plausible but inaccurate answers with total confidence.

Policy needs to lean toward strict control. Kayyali (2024) argues for exclusion of generative AI below lower-secondary level, and UNESCO (2024) advises keeping such tools out of reach for under-13s unless heavy safeguards are in place. Where AI is permitted, it must be under direct teacher supervision (Human in the Loop), tied to clear learning objectives, with every output validated by a human adult.

The long-run cost of inaction is commercial as well as educational. If primary schooling fails to lay down the habits of independent thought and robust memory, tomorrow's workforce enters secondary and tertiary study, then the labour market with fragile cognitive foundations. That deficit shows up later as poor critical reasoning, weak adaptability, and thin emotional resilience, all of which bleed productivity in high-performance teams. In short, talent risk starts in the early years, and business eventually foots the bill.

Secondary Education: Ethics, Identity, and the Illusion of Mastery

The secondary years are a hinge-point in human development. Adolescents are testing boundaries, shaping self-image, and absorbing the ethical norms that will guide adult life. When a generative-AI tool lands in this mix, the risk profile shifts from pure cognition to ethics and identity.

Recent classroom studies show how quickly moral drift sets in once AI shortcutting is normalised. Singh *et al.* (2025) tracked pupils who began by "only" paraphrasing with a chatbot and, within a term, were submitting entire AI-written assignments. Wecks *et al.* (2024) reached a similar conclusion: once students see peers using GenAI without consequence, the social cost of cheating collapses.

That drift feeds a wider illusion of mastery. A polished essay, a solved equation set, a block of neatly formatted code - produced in seconds - can create powerful but hollow confidence. Bastani *et al.* (2024) demonstrated this in a controlled removal test: pupils who relied on AI scored highest while the tool was available, then fell sharply below their non-AI peers once it was taken away. Competence never existed; the façade did.

History offers a cautionary precedent: in late-imperial China wealthy elites hired proxies to sit the civil-service exams, credentials flourished while capability hollowed out. GenAI makes proxy scholarship frictionless, and as Selwyn (2019) observes, AI may preserve the appearance of achievement while hollowing out the deeper cognitive and ethical development that defines real education.

Detection is no safeguard on its own. Many AI outputs slip past standard plagiarism checkers, and even specialised detectors carry high false-positive and false-negative rates (Ardito, 2024). That ambiguity threatens the credibility of grades and the trust relationship between teacher and student.

Beyond academic integrity, the habitual use of generative AI for social interaction or emotional support, a trend noted in some adolescent populations, carries risks of fostering unhealthy dependencies and potentially hindering the development of authentic interpersonal skills and emotional resilience (Chu *et al.,* 2025).

The responsibility therefore rests squarely with educators and school leaders:

- Scaffolded AI literacy. Students need structured practice in evaluating AI outputs, spotting hallucinations, and citing assistance properly. National Education Association guidance (2024c) recommends explicit mini-modules on ethical prompting and verification.

- Redesigned assessments. Tasks that demand higher-order thinking, comparative critique, or real-time oral defence are far harder to outsource to a model.

- Transparent policy and swift, fair penalties. Clear boundaries, evenly enforced, slow the peer-pressure spiral that drives widespread misuse.

Above all, secondary education must keep nurturing independent judgement. Teenagers learn integrity by exercising it, not by outsourcing it. If schools abdicate that role, business inherits a cohort of recruits who look qualified on paper yet crumble under ethical strain.

The downstream costs are real: higher rates of misconduct investigations, fragile team cultures, and expensive remedial training to rebuild habits that should have formed in the classroom.

Fail here, and the talent pipeline is tainted long before graduates reach university or the workplace. Succeed, and secondary education becomes the firewall that protects both academic integrity and future professional standards.

Tertiary Education: Competence Collapse and Business Fallout

By the time students reach university, credentials should be a reliable proxy for capability. Instead, the uncritical spread of GenAI is creating an ever-wider gap between the two.

Recent studies make the pattern hard to ignore. Wecks *et al.* (2024) tracked classes in which more than half the coursework was produced with AI assistance; grades held steady, yet follow-up tasks that required original synthesis fell by as much as 18 percent.

Bastani *et al.* (2024) ran a removal experiment and found a 17 percent performance crash once AI access was withdrawn. Singh *et al.* (2025) added a cognitive lens, showing measurable drops in memory retention and transfer reasoning among heavy GenAI users. The evidence converges: transcripts look stronger, competence is thinner.

For employers this is not an academic quibble; it is a balance-sheet risk. Graduates arrive able to prompt but unable to probe, confident with templates yet lost in ambiguity. The cost shows up as mis-hiring, re-work, and stalled projects, what Webster and Westerman (2025) call “hiring illusions.”

Repairing the Signal

Universities cannot fix this with another plagiarism plug-in; most detectors misfire or are easily bypassed (Ardito, 2024). Instead, assessment itself must change:

- Viva and oral defence. A ten-minute oral exam can surface authentic understanding that AI-generated text may conceal. Emerging evidence from UK and international pilots suggests that reintroducing vivas alongside written work improves academic integrity and reduces reliance on generative AI tools (Maher *et al.*, 2024; Le, 2024).
- Portfolio and studio models**.** Long-form projects, logged in version history, expose the learner’s process, not just the end product.
- Employer-embedded tasks**.** Co-designed briefs and internships let businesses test competence directly while students earn credit.

For years, UK institutions have cited scale and practicality as reasons to abandon the viva. The sharp rise in international postgraduate taught (PGT) student enrolments created pressure to standardise and streamline assessments. However, that growth-driven model is faltering. Student numbers are plateauing, geopolitical tensions are reshaping the international student market, and the economic justification for mass automation of assessment is weakening.

Now is the moment to rethink. Copenhagen Business School, one of the world's top-ranked institutions vivas every graduate, at every level. This is not a theoretical aspiration. It is working policy. Crucially, its reputation is not just sustained but strengthened by the quality signal that graduates carry with them into the labour market.

Moreover, employers recognise the value, students earn validation that extends beyond the classroom, and the institution builds a brand rooted in capability, not just credential. UK universities have the capacity to do the same, should the will to uphold academic rigour return. Rigour creates trust, trust creates value, and value is the true foundation of a university brand.

Shared Accountability

University leaders must treat GenAI governance as a board-level risk, not an IT add-on. That means clear policy, mandatory staff CPD, and transparent reporting on AI-related misconduct (Kayyali, 2024). Employers, for their part, should stop accepting glossy PDFs as proof of skill and start partnering with faculties on competency-based evaluation.

Consequences of Inaction

Shadow-AI, meaning employees covertly using external models, will proliferate if graduates have never been taught responsible practice. Credential inflation will feed cynicism, driving "do-over" recruitment tests that waste time and money. Worst of all, the reputational hit will not fall solely on universities; firms that hire on paper prestige will feel it in client confidence and investor trust.

In short, tertiary education is the last checkpoint before talent meets reality. If that gate fails, the cost is paid well beyond the campus walls.

Cross-Linkages and the Business Imperative

The failures mapped across the educational lifecycle do not sit in silos; they stack. When early schooling allows GenAI to short-circuit handwriting, memory, and curiosity, students enter adolescence with a weaker cognitive core. That fragility encourages shortcut culture at secondary level, where plagiarism and AI-assisted answers feel normal, not illicit. By the time the same learners surface at university, the pattern hardens into what Singh *et al.* (2025) call competence collapse, a résumé that looks strong, yet cracks under real-world strain.

Figure 7 already shows how each stage feeds the next. Primary cognitive gaps become secondary ethical drift, which in turn amplifies tertiary credential inflation, shadow-AI at work, and the hiring illusions that Webster and Westerman (2025) warn are bleeding productivity. In other words, a single uncorrected blind spot in Year 4 can echo into a bad board-room hire ten years later.

For business leaders, that timeline is not abstract, it is today's talent pipeline. Pew's 2024 survey confirms that GenAI use is mainstream among teenagers; the first full "AI-native" cohort will graduate within three years. If those students have never faced meaningful friction, employers will inherit the bill: retraining costs, decision errors, and credibility loss with clients.

The remedy starts upstream:

Primary phase: Support strict age-appropriate limits and teacher-led scaffolding so children still wrestle with handwriting, numeracy, and open-ended questions (Kayyali, 2024; UNESCO, 2024).

Secondary phase: Co-design AI-literacy and ethics modules that make shortcutting visible and unattractive, while rewarding original synthesis (Wecks *et al*., 2024).

Tertiary phase: Partner on competency-based assessments; vivas, live case work, employer-embedded projects, to ensure credentials signal real skill, not prompt fluency (Russell Group, 2023; Bastani *et al.,* 2024).

None of these fixes work alone. The pipeline is only as strong as its weakest joint. A coordinated, evidence-informed alliance; teachers, universities, and industry, remains the single best hedge against the compounded risk of hollow education and the business fallout that follows.

9. The Conditional Promise of GenAI: Controlled Environments and Measured Gains

Wang and Fan (2025) conducted a meta-analysis of 51 studies, reporting a large positive effect of ChatGPT on student learning performance (g = 0.867), and moderate effects on learning perception (g = 0.456) and higher-order thinking (g = 0.457). However, the robustness of these effect sizes has been questioned. A subsequent reanalysis by Bartoš et al. (2025), adjusting for potential publication bias, found substantially smaller effect sizes and concluded that the evidence for a positive impact of ChatGPT on learning performance, perception, and higher-order thinking was largely diminished or absent. This highlights the ongoing debate and the need for further high-quality, pre-registered research to definitively establish the benefits of generative AI in educational settings.

The catch is in the fine print. Every high-impact study shared the same ingredients: clear learning objectives, explicit alignment with a taxonomy such as Bloom-Anderson, tight scaffolding, and a human instructor who stayed in the loop. When those guardrails were present, GenAI acted like an intelligent co-tutor; prompting reflection, supplying instant feedback, and freeing teachers to focus on misconceptions that no algorithm can spot unaided.

Moreover, Wang and Fan's wider review of seventy peer-reviewed articles echoes that verdict. Positive outcomes depended on ethical safeguards, transparent usage guidelines, and deliberate training for both staff and students. In other words, GenAI lifted rigour only when the institution treated it as a structured learning partner, not as a magic shortcut.

These findings map neatly onto the governance principles already highlighted in this report: age-appropriate access, visible teacher oversight, and continuous validation of outputs (Chan, 2023; UNESCO, 2024). Where those principles are missing, the same technology delivers surface fluency at the expense of depth, exactly the risk charted in Sections 4 through 8.

The lesson for policymakers and business alike is straightforward: GenAI's upside is real, but it arrives only by design. Treat the tool as a replacement for human pedagogy and the gains evaporate; treat it as a carefully controlled catalyst with a human in the loop and you can amplify engagement, disciplinary grounding, and critical thought. The choice is not about embracing or banning the technology (which I am not advocating); it is about whether we have the strategic discipline to deploy it under the conditions that evidence demands.

10. The Perils of Unchecked GenAI: Cognitive Harm and Erosion of Critical Thinking

Just as I've seen the young cashier at a till flounder to mentally calculate my change, I have watched bright interns stall mid-conversation because, a prompt was no longer within reach; it is a small anecdote, yet Singh *et al.* (2025) show the same pattern at scale. Their three-year, multi-cohort study tracks students who use GenAI without guidance. At first grades hold steady, assisted by quick generative answers, but deeper measures slide. By the final testing wave, those heavy AI users score significantly lower on critical-thinking, causal-reasoning, and transfer tasks than peers who worked through material unaided.

Singh *et al.* call the mechanism critical-thinking displacement. When a learner accepts a fluent answer, the brain skips interrogation, synthesis, and self-correction. Over hundreds of assignments that missing friction accumulates, so resilience, the capacity to keep probing when the first answer looks convincing, simply atrophies.

Furthermore, Bastani *et al.* (2024) confirm the risk in a controlled teaching experiment. Nearly a thousand undergraduates were given access to GPT-4 as a homework "tutor". During the term they outperformed the control group, yet once AI access was removed their scores dropped by seventeen per cent. Retention was shallow, error rates spiked, and post-course surveys recorded lower intellectual curiosity. The researchers conclude that unsupervised GenAI "front-loads performance at the cost of durable learning.

A similar drift appears in Wecks *et al.* (2024), who found that students relying on AI explanations produced more correct multiple-choice answers but struggled badly on open-ended questions requiring synthesis. The tool delivered surface fluency, yet left gaps in reasoning that only emerged when novelty or ambiguity entered the test.

Taken together, the evidence is blunt. GenAI becomes a catalyst for learning only when embedded in a structured cycle of experience, reflection, and feedback. When it is used as a shortcut, it quietly erodes the very skills higher education claims to build: disciplined inquiry, judgement under uncertainty, and the stamina to wrestle with incomplete information.

Unchecked adoption therefore harms more than grades. It threatens the supply of graduates who can diagnose complex problems, challenge suspect data, and spot hidden assumptions. Business and society depend on those capacities; losing them would turn a convenient tool into an expensive un-fixable liability.

11. Decreased Student Engagement and Deeper Learning Challenges

Evidence suggests the risk extends well beyond lost reasoning power; it reaches student motivation itself. In a mixed-methods study, Wecks *et al.* (2024) tracked the day-to-day behaviour of learners who lean on GenAI for routine tasks. Those students joined fewer peer discussions, did not persist as long on open-ended problems, and showed less reflective writing in learning journals. Put plainly, if the bot can "do the work," the learner is less inclined to wrestle with it. That short-circuit in effort shows up first as disengagement and, over time, as shallower understanding and weakened resilience.

Neuro-behavioural studies echo the classroom data. AI's instant-gratification loop spikes dopamine, delivering the 'right' answer without the affective hit of hard-won insight. Over time this shifts the motivation baseline and, according to Pew Research Center (2024), correlates with rising anxiety and emotional disengagement among teen heavy-users.

Moreover, a large field experiment led by Bastani *et al*. (2024) adds a longitudinal lens. Nearly a thousand students were given access to GPT-4 tutors for selected assignments. Initial scores jumped, yet once access was removed the same cohort scored 17 percent lower than peers who had never used AI. The researchers concluded that unsupervised GenAI creates a dependency that weakens self-regulation and academic resilience.

Together, the two studies expose a paradox. GenAI's speed and convenience can, without guardrails, erode precisely the habits that universities prize: critical thinking, sustained curiosity, and collaborative inquiry (Wecks *et al.,* 2024; Bastani *et al*., 2024). Over-reliance also breeds overconfidence; students come to overestimate and overplay what they have mastered and underestimate the blind spots in AI output. Unless educators frame AI as a scaffold, rather than a substitute, the very tool meant to personalise learning risks draining it of depth.

12. The Strategic Imperative: Structured, Pedagogically Sound, and Controlled Integration of AI

The evidence is now overwhelming: GenAI can raise performance only when its use is deliberate, scaffolded, and constantly evaluated. Wang & Fan (2025) show large gains in learning and higher-order thinking, but those gains appear only in classrooms with clear goals, active teachers, and tight feedback loops.

Five pillars of a resilient GenAI strategy

1. **Align every use case with explicit learning objectives and assessment criteria.** If a task measures synthesis, require students to show where and how the AI contributed.
2. **Build AI literacy for staff and students.** Teach critical evaluation, ethical boundaries, and the limits of large language models.
3. **Design formative assessments that force interaction, critique, and revision.** Students should improve or rebut AI output, not hand it in unchanged.
4. **Track engagement and integrity in real time.** Early-warning dashboards can flag over-reliance long before deep learning suffers.
5. **Create shared governance.** Educators, technologists, and policymakers must co-own policies, procurement standards, and professional-development budgets.

Two live frameworks illustrate how to make this real: Australia's Framework for Generative AI in Schools and Chan's AI Ecological Education Policy Framework (Chan, 2023). Both insist on clear governance roles, teacher readiness, and alignment with established pedagogy.

Practitioner lens. In my own work deploying AI across retail, media and public services, the biggest failures were never technical. They stemmed from unclear objectives, absent guard-rails, and zero metrics beyond "does it work?". The same pattern now stalks education. If we ignore it, we will graduate students who can prompt beautifully but think shallowly and business will be left to pick up the cost.

Unchecked, poorly governed GenAI will hollow out the very capacities; critical reasoning, creativity, lifelong learning that drive competitive economies. The strategic choice is stark: integrate with intent or watch academic credentials and the talent pipeline, lose their signalling power.

13. Practitioner-Led Recommendations and Case Studies on AI Governance,

13.1 Field Pilots: International Innovations

South Korea – Robot-Assisted Language Immersion. A series of pilot projects in Daegu and other cities replaced a portion of English oral-practice time with **"EngKey"** telepresence robots. Peer-reviewed evaluations reported *significant* gains in vocabulary recall and pronunciation accuracy (≈ 0.5 SD) compared with control classes (Yun *et al.*, 2011; Belpaeme *et al.*, 2018).

However, researchers also found that empathy cues and humour still had to be modelled by the human teacher to foster deeper cultural fluency. The project illustrates a clear boundary line: repetition and confidence building can be automated; nuance and socio-linguistic skill cannot.

13.2 Institutional Governance Models: UK & Europe: UK & Europe**

University of Edinburgh – ELM Platform. The University of Edinburgh has implemented comprehensive internal protocols for the ethical use of generative AI. Central to its governance model is the creation of Edinburgh Language Models (ELM), a proprietary platform providing controlled access to large language models while ensuring that sensitive student and staff data are not exposed to third-party services. ELM operates under a zero-data-retention agreement with OpenAI, reinforcing compliance with UK GDPR and institutional standards for responsible data handling (University of Edinburgh, 2024a; 2024b).

University of London – Formal AI Policy. The University of London has issued a formal Artificial Intelligence Policy, setting out ethical, legal, and pedagogical guidelines for responsible AI adoption across the institution (University of London, 2024). This policy mandates impact assessments for new AI tools, clarifies data governance requirements, and delineates faculty responsibilities for updating assessment designs.

Russell Group & Jisc – Coordinated Assessment Redesign. Members of the Russell Group, supported by Jisc (2024), encourage redesigning assessments to promote critical thinking, originality, and applied reasoning. National guidelines emphasize that AI-aware assessment prompts should assume student access to generative tools and require demonstrable reflection on AI use, thereby preserving academic integrity while leveraging AI's benefits.

ETH Zurich, EPFL & Swiss National AI Institute – Pan-European Collaboration. At ETH Zurich and EPFL, governance efforts are coordinated through the Swiss National AI Institute, which promotes research excellence, cross-sector collaboration, and ethical open-source AI development (ETH Zurich, 2024). These institutions align procurement with the EU AI Act and the Ethics Guidelines for Trustworthy AI, embedding legal safeguards and promoting AI alignment with fundamental rights and academic values (Digital Education Council, 2024).

13.3 U.S. & Global Policy Frameworks

U.S. Universities – IT & Academic Policies. Information Technology units at Indiana University and Michigan State University have implemented policies restricting data sharing with external AI tools and established enterprise agreements to ensure trustworthy AI procurement, drawing on frameworks like the NIST AI Risk Management Framework (NIST, 2023). Teaching and Learning divisions emphasize faculty autonomy and encourage redesigning assessments to require higher-order thinking, thus mitigating inappropriate AI use. University Libraries play a critical advisory role on ethical AI in research, ensuring compliance with publisher policies and proper attribution.

UNESCO & OECD – Global Guidelines. International bodies such as UNESCO and the OECD have developed comprehensive guidelines emphasising human-centered, ethical, and transparent AI use in education. UNESCO's recommendations stress aligning AI policies with the Universal Declaration of Human Rights, focusing on beneficence, non-maleficence, autonomy, justice, and explicability (UNESCO, 2024). The OECD's AI Principles reinforce the need for explainability, transparency, and fairness, providing a baseline for national and institutional policy development.

13.4 Practitioner Recommendations for Assessment Reform & Partnerships

Assessment Reform – Designing for AI Presence. Practitioners recommend adapting assessment design to assume generative AI availability, focusing on higher-order thinking, critical analysis, and authentic application of knowledge. A UK business-school case study with 118 postgraduate management students found that most used GenAI for writing and paraphrasing, with few engaging in advanced critical evaluation. Authors advise leaders to cultivate students' critical evaluation and academic writing skills through explicit GenAI integration strategies (Wecks *et al.*, 2024 working paper).

Quality Assurance Frameworks. Robust frameworks for quality online assessment foreground academic integrity, student experience, authenticity, information integrity, quality feedback, and equity of access (White & Huber, 2024). Adapted prompts that acknowledge appropriate AI use have restored integrity while preserving other quality dimensions (Khlaif *et al.*, 2025).

Validity & Human Oversight. Because AI can generate items and score responses, validity evidence must be gathered from test content, response processes, internal

structure, and alignment with intended learning outcomes (Kaldaras *et al.*, 2024). AI-generated items and scores should be audited by educators to confirm they measure intended constructs (Kaldaras *et al.*, 2024; Ardito, 2024).

Teacher Professional Development. Large-scale PD studies stress that comprehensive training in AI literacy, assessment design, and ethics is critical for successful reform (Tan *et al.*, 2025).

Business–University Partnerships. Partnerships bridge the AI skills gap, foster innovation, and align educational programs with workforce needs. Examples include Macquarie University's School of Education partnering with IBM to develop an inclusive AI curriculum framework for middle schools, and East China Normal University collaborating with SenseTime to publish high school AI textbooks.

The University of Washington and University of Tsukuba, supported by NVIDIA and Amazon, established a research and workforce-training partnership providing access to high-performance computing and AI technologies for AI-driven careers in robotics, healthcare, and climate science. Benefits include workforce upskilling, real-world application, curriculum relevance, lifelong learning, and economic impact. Practitioners recommend curriculum co-design, internships, joint research hubs, flexible access, and continuous evaluation to strengthen partnerships.

14. What Institutions Must Do Now

The damage of academic rigour in the age of AI is not inevitable; it is the outcome of conscious choices. Reversing the decline requires decisive, system-level action on five fronts:

Appoint AI 'Czars'

The recommendation to appoint a senior leader, or 'AI Czar', is a critical step for any educational institution seeking to navigate the complexities of generative AI. This is not a role for a traditional IT manager or a curriculum lead alone; it demands a unique blend of deep, real-world expertise in technology, governance, and the rapidly evolving landscape of artificial intelligence.

The ideal Czar is someone who understands that the challenges posed by AI are not merely technical but deeply pedagogical and strategic. They recognise that the uncontrolled spread of GenAI is already creating a "competence collapse" where academic transcripts no longer reliably signal true capability. This leader must grasp the core argument of this report: that without structured intervention, AI fosters an "illusion of mastery", quietly eroding the critical thinking, reasoning, and memory skills that are the bedrock of academic rigour and professional success.

Key Qualifications and Understanding:

- **Deep AI Fluency:** The ideal person must possess more than a theoretical knowledge of artificial intelligence. They need "deep, real-world AI experience" to understand the technology's architecture, its limitations (such as "hallucinations"), and how it is being deployed in commercial sectors. This practical knowledge is essential for creating policies that are robust and for leading "practitioner-led recommendations" that go beyond academic theory. They must understand that the problem is not the technology itself, but its "widespread, unsupervised use".
- **Expertise in Governance and Risk:** This role requires a leader with board-level authority to establish clear governance, own strategy, and manage risk. They will be responsible for creating institutional frameworks, such as the proprietary "Edinburgh Language Models (ELM)" platform, to ensure data privacy and ethical usage. This leader must address the proliferation of "shadow AI use" by establishing transparent policies and promoting responsible practice to prevent data breaches and accountability gaps.
- **Agility with a Fast-Moving Target:** AI is evolving at "lightning speed," and the Czar must be abreast of the latest trends. This includes understanding the nuances of new models, the continuous failure of detection tools, and the shifting ways students use these technologies. Their strategic guidance will be crucial in helping educators, who currently lack preparedness, to design assessments that are resilient to AI-enabled cheating and test authentic competence.
- **A Bridge Between Academia and Industry:** A critical function of the Czar is to forge "deeper business-university partnerships". Having seen the "business fallout" from the academic-to-job mismatch firsthand, this individual will champion the co-design of curricula and assessments that ensure graduates are genuinely "job-ready". They will understand the report's warning that if education fails, "business will pay the price" through poor hiring decisions, skills gaps, and diluted productivity.

In essence, the AI Czar must be a strategic thinker who can connect the dots between a student using a chatbot in primary school and the long-term "labour-market misalignment". They are the institutional steward empowered to implement the "human-in-the-loop systems" necessary to turn AI from a threat into a catalyst for genuine learning and to restore the integrity and value of an academic credential in the age of AI.

Reintroduce Viva-Style Assessments

Oral examinations and in-person defences are markedly harder to game with GenAI. They compel students to articulate reasoning in real time and withstand probing, thereby restoring a robust measure of competence (Kaldaras *et al.,* 2024; Ardito, 2024).

Systems lens. Higher education can be viewed as a transformation process; like any quality-driven production line, it needs batch testing and randomised spot checks to surface defects (Durmuş Şenyapar & Bayındır, 2024). Periodic vivas serve precisely that purpose.

Brand lens. Universities trade on trust. Rigorous human-in-the-loop spot-checking safeguards degree credibility and deters credential fraud (Gibbs *et al.,* 2024; Kayyali, 2024).

Invest in Staff CPD

Continuous, mandatory professional development in AI literacy, digital pedagogy, and assessment integrity is non-negotiable. Educators cannot enforce rigour with yesterday's tool-set.

Forge Deeper Business-University Partnerships

Employers and educators must co-design curricula, internships, and assessment models that withstand AI-enabled fraud and map directly to workplace needs. Such collaboration rebuilds trust and ensures graduates arrive "job-ready".

Conduct Formal AI Capability Audits

Each institution should commission an independent review of policies, infrastructure, and staff readiness. Findings must be transparent and action-oriented; without honest baselines, improvement is impossible.

These reforms are minimum viable steps for restoring trust, rigour, and relevance to education in the era of generative AI.

> *What strikes me most isn't just what the data shows, but how quickly these academic trends become business problems. I've seen talented, motivated young professionals crumble when forced to think beyond AI-generated prompts. Employers can no longer assume that a degree signals readiness. Until we close the gap between academic credential and practical competence, the price of doing little or nothing will keep rising, for business, for students, and for society as a whole*.
>
> Mark Evans

15. Conclusion

Generative AI casts a long shadow over education, threatening to undermine the very rigour on which knowledge and economic vitality depend. Left to its own devices, it turns coursework into performance art and transcripts into unreliable signals, hollowing out the skills base that businesses and societies require. Stephen Klein's warning that technology can "spread darkness at the speed of light" (Klein, 2025) is not hyperbole; it is a precise description of what happens when institutions allow convenience to outrun cognition.

Yet the same technology that erodes depth can, when woven into a disciplined pedagogical fabric, illuminate learning. A meta-analysis by Wang and Fan (2025) shows that GenAI lifts higher-order thinking and engagement when its use is framed by clear objectives, rigorous scaffolding, and relentless human oversight. The path forward lies not in rejecting AI, but in reclaiming its direction.

That reclamation begins with governance. Every institution must place someone fluent in both pedagogy and production-grade AI at the decision-making table – a steward empowered to align strategy, risk, and opportunity. It continues with sustained investment in faculty development so that teachers, not algorithms, remain the curators of intellectual challenge and integrity.

Assessment must evolve in parallel. Viva-style spot checks and portfolio defences restore the friction that exposes true understanding, making credentials trustworthy once more. At the same time, universities and employers must lock arms, co-designing curricula and internships that translate academic mastery into workplace competence, while independent capability audits keep each campus honest about its progress.

These measures are intertwined. Governance without professional learning is hollow. Assessment reform without employer partnership is misaligned. And any of it without transparent auditing is little more than rhetoric. Together, though, they form a firewall between artificial fluency and authentic mastery.

I write as a practitioner who has spent over a decade commercialising AI across continents and as a recent MBA graduate who has experienced both the promise and the peril of these tools in the classroom. That dual vantage point fuels not pessimism, but critical optimism.

The most powerful technologies in history have always required human-centred guardrails. GenAI is no exception. True transformation will flow not from the algorithms alone, but from robust human-in-the-loop systems that insist on ethical alignment, transparent accountability, and relentless intellectual stretch. Even the diagrams that punctuate this report were generated in dialogue with AI, evidence that, when harnessed deliberately, the machine extends human reach without eclipsing human judgement.
As the clock ticks, inaction will allow credentials to drift further from competence and leave future workforces isolated on fragile foundations. On the other hand, a concerted, collaborative embrace of the safeguards outlined here can turn GenAI from an existential threat into an unparalleled catalyst for deeper learning and economic resilience. The choice is stark. The stakes are generational. And the moment for decisive leadership is now.

What began as a technological shortcut is becoming a systemic drift. From rewired brains in Year 6 to credential inflation at graduation, the cost of inaction is now hardwired into the

talent pipeline. Business cannot afford to wait for education to self-correct. The intervention must be collaborative, strategic, and fast.

Written by Mark Evans MBA, CMgr FCMi
Founder of 360 Strategy
mark@360strategy.co.uk
Tel +44 (0)7740288535
www.linkedin.com/in/markevans007

16. References

Anderson, L.W., Krathwohl, D.R. and Bloom, B.S. (2001) *A taxonomy for learning, teaching and assessing*: *a revision of Bloom's taxonomy of educational objectives*. New York: Longman.

Ardito, C.G. (2024) 'Generative-AI detection in higher-education assessments', *New Directions for Teaching and Learning*, early view. doi:10.1002/tl.20624.

Bankins, S. and Formosa, P. (2023) 'The ethical implications of artificial intelligence for meaningful work', *Journal of Business Ethics*, early view. doi:10.1007/s10551-023-05541-9.

Bartoš, F., Martinková, P. and Wagenmakers, E.-J., 2025. *Adjusting for publication bias reveals no evidence for the effect of ChatGPT on students' learning performance, learning perception, and higher-order thinking*. PsyArXiv. Available at: https://osf.io/preprints/psyarxiv/8vs32_v1 [Accessed 5 June 2025].

Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B. & Tanaka, F. (2018) "Social robots for education: A review," *Science Robotics*, 3(21), eaat5954.

Barrett, R., Cullen, R., Smith, K. and Thieme, N. (2023) 'Not quite eye to A.I.: student and teacher perspectives on generative AI in writing', *International Journal of Educational Technology in Higher Education*, 20, 52. doi:10.1186/s41239-023-00431-8.

Bastani, H., Choudhury, G., Ouyang, L., Agarwal, S. and Shi, J. (2024) *Generative AI can harm learning*. Wharton Working Paper 23-15. Available at: https://papers.ssrn.com/abstractid=4895486 (Accessed: 26 May 2025).

Bolden, R., Petrov, G. and Gosling, J. (2009) 'Distributed leadership in higher education: rhetoric and reality', *Educational Management Administration & Leadership,* 37(2), pp. 257-277. doi:10.1177/1741143209102784.

Bulut, O., Cormier, D.C. and Rose, C.P., 2024. *The rise of artificial intelligence in educational measurement: Opportunities and ethical challenges*. arXiv preprint arXiv:2406.18900. [online] Available at: https://arxiv.org/abs/2406.18900 [Accessed 5 June 2025].

Chan, C.K.Y., Hu, G. and Zhang, L. (2023) 'Students' voices on generative AI: perceptions, benefits and challenges', *International Journal of Educational Technology in Higher Education*, 20, 49. doi:10.1186/s41239-023-00429-2.

Chan, C.K.Y. (2023) 'A comprehensive AI-policy education framework for university teaching and learning', *International Journal of Educational Technology in Higher Education*, 20(1), pp. 1-27. doi:10.1186/s41239-023-00406-9.

Chu, M.D., Gerard, P., Pawar, K., Bickham, C. and Lerman, K., 2025. *Illusions of Intimacy: Emotional Attachment and Emerging Psychological Risks in Human-AI Relationships.* arXiv preprint arXiv:2505.11649. [online] Available at: https://arxiv.org/abs/2505.11649 [Accessed 5 June 2025].

Crompton, H., Burke, D. and O'Sullivan, S. (2023) 'Artificial intelligence in higher education: the state of the field', *International Journal of Educational Technology in Higher Education*, 20, 36. doi:10.1186/s41239-023-00413-y.

Davis, V. (2023) *Developing institutional-level AI policies and practices: a framework*. Boulder, CO: WCET. Available at: https://wcet.wiche.edu/resources/developing-institutional-level-ai-policies-and-practices/ (Accessed: 26 May 2025).

Deem, R., Hillyard, S. and Reed, M. (2007) *Knowledge, higher education and the new managerialism*. Oxford: Oxford University Press.

Digital Education Council (2024) 'EU AI Act: what it means for universities'. Available at: https://www.digitaleducationcouncil.com/post/eu-ai-act-what-it-means-for-universities (Accessed: 26 May 2025).

Damasio, A. R. (1994) *Descartes' error: emotion, reason, and the human brain*. New York: Putnam. Available at: https://www.penguinrandomhouse.com/books/297609/descartes-error-by-antonio-damasio/ (Accessed: 22nd May 2025).

Durmuş Şenyapar, H.N. and Bayındır, R. (2024) 'Quality assurance in higher education in the 21st century: strategies and practices for new generation universities', *Higher Education Governance and Policy*, 5. doi:10.55993/hegp.1573331.

EdWeek Research Center (2024) *Most teachers are not using AI*. Available at: https://www.edweek.org/technology/most-teachers-are-not-using-ai-heres-why/2024/01 (Accessed: 26 May 2025).

Engeström, Y. (1987) *Learning by expanding: an activity-theoretical approach to developmental research*. Helsinki: Orienta-Konsultit. Available at: https://lchc.ucsd.edu/mca/Paper/Engestrom/Learning-by-Expanding.pdf (Accessed: 23rd May 2025).

ETH Zurich (2022) *Guidelines on scientific integrity (RSETHZ 414)*. Zurich: ETH Zurich. Available at: https://rechtssammlung.sp.ethz.ch/dokumente/414en.pdf (Accessed: 29th May 2025).

European Commission (2024) *Guidance for generative AI in education and research*. Available at: https://education.ec.europa.eu (Accessed: 26 May 2025).

Farrelly, T. et al. (2023) 'Generative AI: implications for higher-education practice', *Education Science*s, 13(11), 1178. doi:10.3390/educsci13111178.

Floridi, L., Cowls, J., Beltrametti, M. et al. (2018) 'AI4People — an ethical framework for a good AI society', *Minds and Machines*, 28(4), pp. 689-707. doi:10.1007/s11023-018-9482-5.

Gibbs, B., Lewis, E., Wood, G.C. et al. (2024) *When quality assurance meets innovation in higher education*. Gloucester: Quality Assurance Agency. Available at: https://www.qaa.ac.uk/docs/qaa/members/when-quality-assurance-meets-innovation-in-higher-education.pdf (Accessed: 22nd May 2025).

González-Calatayud, V., Iglesias-Rodríguez, A., Cuéllar, M.P. and Lloret, J. (2021) 'Artificial intelligence for student assessment: a systematic review', Applied Sciences, 11(12), 5607. doi:10.3390/app11125607.

Department for Education (2025) Generative artificial intelligence (AI) in education. London: Department for Education. Available at: https://www.gov.uk/government/publications/generative-artificial-intelligence-in-education (Accessed: 26 May 2025).

High-Level Expert Group on AI (2019) Ethics guidelines for trustworthy AI. Brussels: European Commission.

Jisc (2024) Artificial intelligence – supporting universities and colleges. Available at: https://www.jisc.ac.uk (Accessed: 26 May 2025).

Jisc, 2024. *Communities in action at Digifest*. [online] Jisc. Available at: https://www.jisc.ac.uk/news/all/communities-in-action-at-digifest [Accessed 5 June 2025].

Jisc (2025) Student perceptions of generative AI. Bristol: Jisc. Available at: https://www.jisc.ac.uk/reports/student-perceptions-of-generative-ai (Accessed: 26 May 2025).

Kaldaras, L., Akaeze, H.O. and Reckase, M.D. (2024) 'Developing valid assessments in the era of generative AI', Frontiers in Education, 9, 1399377. doi:10.3389/feduc.2024.1399377.

Kayyali, M. (2024) 'The evolution of quality assurance in higher education: models, systems and frameworks', in Navigating Quality Assurance and Accreditation in Global Higher Education. Hershey, PA: IGI Global, pp. 1–26. doi:10.4018/979-8-3693-6915-9.ch001. Available at: https://www.igi-global.com/chapter/the-evolution-of-quality-assurance-in-higher-education/365728 (Accessed: 29 May 2025).

Khlaif, Z.N., Awadallah Alkouk, W., Salama, N. and Abu Eideh, B. (2025) 'Redesigning assessments for AI-enhanced learning', Education Sciences, 15(2), 174. doi:10.3390/educsci15020174.

Le, Q.T., 2024. *The AI Assessment Scale (AIAS) in action: A pilot implementation of GenAI-supported assessment*. [online] ResearchGate. Available at: https://www.researchgate.net/publication/384969779_The_AI_Assessment_Scale_AIAS_in_action_A_pilot_implementation_of_GenAI-supported_assessment [Accessed 5 June 2025].

Li, W., Zhao, C., Xu, X. and Tang, Y. (2024) 'Educators' perspectives on generative AI in higher-education teaching and learning', Computers & Education: Artificial Intelligence, 5, 100146. doi:10.1016/j.caeai.2024.100146.

Lin, F., Liu, T., Hwang, G. and Kuo, F. (2023) 'Artificial-intelligence tutoring systems toward sustainable education: a systematic review', Smart Learning Environments, 10, 17. doi:10.1186/s40561-023-00241-5.

Maher, C., Sharma, K., Nag, A., Ekblom, P. and Cunningham, A., 2024. *Re-designing Assessments to Minimise Misuse of Large Language Models: A Case Study from a Russell Group University*. arXiv preprint arXiv:2501.12883. [online] Available at: https://arxiv.org/abs/2501.12883 [Accessed 5 June 2025].

Miao, F. and Holmes, W. (2025) Guidance for generative AI in education and research. Paris: UNESCO. Available at: https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research (Accessed: 26 May 2025).

Mouta, A., Pinto-Llorente, A.M. and Torrecilla-Sánchez, E.M., 2023. *Uncovering blind spots in education ethics: Insights from a systematic literature review on artificial intelligence in education*. International Journal of Artificial Intelligence in Education. [online] Available at: https://link.springer.com/article/10.1007/s40593-023-00384-9 [Accessed 5 June 2025].

National Education Association (2024a) Artificial intelligence in education. Washington, DC: NEA.

National Education Association (2024b) The current state of AI in education. Washington, DC: NEA. Available at: https://www.nea.org/resource-library/artificial-intelligence-education/iii-current-state-artificial-intelligence-education (Accessed: 26 May 2025)

National Education Association (2024c) Teaching and learning in the age of AI. Washington, DC: NEA. Available at: https://www.nea.org/nea-today/all-news-articles/teaching-age-ai (Accessed: 26 May 2025).

National Education Association (2024d) Report of the NEA Task Force on Artificial Intelligence in Education. Washington, DC: NEA. Available at: https://www.nea.org/sites/default/files/2024-06/report_of_the_nea_task_force_on_artificial_intelligence_in_education_ra_2024.pdf (Accessed: 29 May 2025).

NeuroTech Research Group (2025) *LLM usage and neurocognitive development: early risks and longitudinal patterns*. Cambridge: Centre for Cognitive AI Studies. Available at: https://arxiv.org/pdf/2506.08872 (Accessed: 17th June 2025).

OECD, 2024. *Generative AI in Education: Aligning Outcomes with Labour Market Needs*. Paris: OECD Publishing. [online] Available at: https://www.oecd.org/education/genai-education-labour [Accessed 3rd June 2025].

OpenAI. (2025). *ChatGPT* (Model version as of 25 April 2025). [Large language model]. Available at: https://chat.openai.com (Accessed: 25 April 2025).

* (Author's interpretation; diagram generated with AI assistance)

Zawacki-Richter, O., Marín, V.I., Bond, M. and Gouverneur, F. (2019) 'Systematic review of research on artificial intelligence applications in higher education – where are the educators?', *International Journal of Educational Technology in Higher Education*, 16(1), 39. doi:10.1186/s41239-019-0171-0.

Pew Research Center (2025) *About a quarter of U.S. teens have used ChatGPT for schoolwork – double the share in 2023*. Available at: https://www.pewresearch.org/short-reads/2025/01/15/about-a-quarter-of-us-teens-have-used-chatgpt-for-schoolwork-double-the-share-in-2023/ (Accessed: 26 May 2025).

Pew Research Center (2025) Share of teens using ChatGPT for schoolwork doubled from 2023 to 2024. Washington, DC: Pew Research Center. Available at: https://www.pewresearch.org/short-reads/2025/01/15/about-a-quarter-of-us-teens-have-used-chatgpt-for-schoolwork-double-the-share-in-2023/ (Accessed: 26 May 2025).

Roe, J., Perkins, M. and Ruelle, D., 2024. *Understanding student and academic staff perceptions of AI use in assessment and feedback*. arXiv preprint arXiv:2406.15808. [online] Available at: https://arxiv.org/abs/2406.15808 [Accessed 5 June 2025].

Russell Group (2023) 'Russell Group principles on the use of generative AI tools in education'. London: Russell Group. Available at: https://www.russellgroup.ac.uk/sites/default/files/2025-01/Russell%20Group%20principles%20on%20generative%20AI%20in%20education.pdf (Accessed: 26 May 2025).

Sakelaris, P.G., Novotny, K.V., Borvick, M.S., Lagasca, G.G. and Simanton, E.G., 2025. Evaluating the use of artificial intelligence as a study tool for preclinical medical school exams. *Journal of Medical Education and Curricular Development*, 12, pp.1–7. Available at: https://journals.sagepub.com/doi/10.1177/23821205251320150 [Accessed 5 Jun 2025]

Selwyn, N. (2019) *Should robots replace teachers? AI and the future of education*. Cambridge: Polity Press.

Singh, A., Taneja, K., Guan, Z. and Ghosh, A. (2025) 'Protecting human cognition in the age of AI', arXiv, 2502.12447. Available at: https://arxiv.org/abs/2502.12447 (Accessed: 26 May 2025). (working paper)

Susskind, J. (2018) Future politics: living together in a world transformed by tech. Oxford: Oxford University Press.

Tan, X., Cheng, G. and Ling, M.H. (2025) 'Artificial intelligence in teaching and teacher professional development: a systematic review', Computers & Education: Artificial Intelligence, 8, 100355. doi:10.1016/j.caeai.2024.100355.

Trajkovski, G., Lemaire, C. and Worsley, M. (2025) 'The future of AI in educational assessment', Educational Measurement: Issues and Practice, early view. doi:10.1111/emip.12494.

UNESCO (2024) Latin America report 2025. Paris: UNESCO. Available at: https://www.unesco.org/gem-report/en/2025lac (Accessed: 26 May 2025).

University of Edinburgh (2024) Guidance for working with Generative AI ("GenAI") in your studies. Edinburgh: University of Edinburgh. Available at: https://information-services.ed.ac.uk/computing/comms-and-collab/elm/guidance-for-working-with-generative-ai (Accessed: 29 May 2025).

University of London (2024) Artificial intelligence (AI) policy. London: University of London. Available at: https://www.london.ac.uk/sites/default/files/artificial-intelligence-policy-uol.pdf (Accessed: 29 May 2025).

U.S. Department of Education, 2023. *Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations*. [online] Available at: https://www.ed.gov/[Accessed 5 June 2025].

Wang, J. and Fan, W. (2025) 'The effect of ChatGPT on students' learning performance, learning perception, and higher-order thinking: insights from a meta-analysis', Humanities and Social Sciences Communications, 12(1), 621. doi:10.1057/s41599-025-04787-y.

Webster, M. and Westerman, G. (2025) ‘Generate value from GenAI with “small t” transformations’, MIT Sloan Management Review, 22 January. Cambridge, MA: MIT. Available at: https://sloanreview.mit.edu/article/generate-value-from-genai-with-small-t-transformations (Accessed: 26 May 2025).

Wecks, J.O., Voshaar, J., Plate, B.J. and Zimmermann, J. (2024) ‘Generative AI usage and exam performance’, SSRN Electronic Journal. doi:10.2139/ssrn.4812513.

White, A. and Huber, E. (2024) ‘Evaluating an online assessment framework through the lens of generative AI’, in Cochrane, T., Narayan, V., Bone, E., Deneen, C., Saligari, M., Tregloan, K. and Vanderburg, R. (eds) Navigating the terrain: emerging frontiers in learning spaces, pedagogies, and technologies. Proceedings ASCILITE 2024, Melbourne. pp. 618–623. doi:10.14742/apubs.2024.1423.

World Economic Forum, 2023. *The Future of Jobs Report 2023*. [online] Available at: https://www.weforum.org/publications/the-future-of-jobs-report-2023/[Accessed 5 June 2025]

Yun, S., Lee, J., Kim, D. & Choi, Y. (2011) “EngKey: Tele-education robot for English class,” in *Proceedings of ICSR 2011*, Berlin: Springer, pp. 142-152.