

## Higher Education and Generative AI in the Era of the Fourth Industrial Revolution

Kehdinga George Fomunyam,

ICON, Durban University of Technology

Orcid ID: 0000-0003-2510-3498

### Abstract

The Fourth Industrial Revolution (4IR), marked by the integration of technologies that merge physical, digital, and biological domains, positions Generative Artificial Intelligence (GenAI) as a key driver of societal and economic change. Higher education institutions stand at a pivotal point, balancing the prospects and hurdles presented by these advanced systems. This article examines the interplay between higher education and GenAI within the 4IR framework. It posits that GenAI represents a profound shift, demanding a re-evaluation of teaching methods, research practices, and the university's fundamental purpose. The discussion covers three key areas: (1) the need for pedagogical adaptation, shifting from conventional evaluations vulnerable to AI outputs toward approaches emphasizing critical analysis, AI proficiency, and human-AI partnerships; (2) GenAI's function as a research enhancer, enabling broader access to intricate analyses and innovative explorations, while prompting debates on integrity, bias, and ownership; and (3) the institutional imperative to adapt curricula to 4IR workforce needs, equipping students for environments where human-machine synergy prevails. The article asserts that for higher education to sustain its role in fostering innovation and growth amid the 4IR, a shift from reactive measures to forward-thinking strategies is essential. This involves establishing strong ethical frameworks, committing resources to skill development for faculty and students, and redesigning education to prioritise irreplaceable human attributes that AI cannot duplicate.

**Keywords:** Generative AI, Higher Education, Fourth Industrial Revolution (4IR), Large Language Models, Chat GPT, DeepSeek, Grok, Future of Work, AI Ethics

## **Introduction: The 4IR, the GenAI Crisis, and the Search for a Framework**

The 21st century is increasingly shaped by the dominant rhetoric of the Fourth Industrial Revolution (4IR), which aims to describe a significant technological transformation. Unlike earlier revolutions driven by steam, electricity, or digital computing, the 4IR is characterised by a “fusion of technologies that is blurring the lines between the physical, digital, and biological spheres” (Schwab, 2017, p. 7). This revolution is said to progress exponentially rather than linearly, fueled by combined advances in artificial intelligence (AI), robotics, the Internet of Things (IoT), and quantum computing. In this context, the public launch of Generative Artificial Intelligence (GenAI), especially OpenAI’s (2023) GPT series, has significantly accelerated its development. These models, capable of generating innovative and complex text, images, and code, have swiftly shifted from specialised tools to widely accessible platforms (Bozkurt, 2023). This quick adoption has made GenAI a key driver of broad societal change, with the potential to reshape industries, labour markets, and fundamental human pursuits, from arts to healthcare diagnostics.

Higher education (HE) institutions see this moment as an existential crisis. Universities, traditionally centres of research, innovation, and human capital development, are at the heart of this transformation. The rise of powerful GenAI tools presents a stark duality. They offer significant opportunities to personalise learning at scale, improve administrative efficiency, and democratise complex research tasks once reserved for specialised labs (Mhlanga, 2023). However, they also introduce urgent challenges that threaten the core of academic work: undermining assessment models (Cotton et al., 2023), enabling sophisticated academic misconduct, and raising serious concerns about ethics, bias, and data privacy (UNESCO, 2023). So far, institutions have responded mainly reactively and in a panicked manner. However, this crisis narrative can obscure a more nuanced reality. As Selwyn (2022) advocates in a postdigital view, technology in education is not just an external force acting on institutions; it is a complex, intertwined part of social, cultural, and political practices. The challenge of GenAI is not just technical; it is pedagogical and human. The anxiety around AI-generated essays, for instance, indicates less about the technology's power and more about the limitations of current assessment methods, which often focus on the final product rather than the learning process—including critical thinking, research, and drafting.

This paper contends that Generative AI is more than just an incremental addition to university technology; it represents a disruptive sociotechnical assemblage requiring a complete rethink

of pedagogy, research methods, and the university's core purpose. Its integration calls for moving beyond a purely technologically deterministic view, which sees us as passive to inevitable change, toward a critical sociotechnical approach. This perspective considers GenAI as an active entity intertwined with human motives, institutional power, and teaching decisions (Latour, 2005). From this angle, the main challenge for higher education is not about AI-proofing curricula or preventing cheating. Instead, it involves balancing the technology's potential to augment human capabilities with its risk of reinforcing inequalities (Noble, 2018). It requires careful choices about what to automate and what to safeguard, along with a deliberate redesign of universities as spaces fostering co-intelligence, the partnership between human and artificial intelligences (Mollick, 2024).

This paper investigates three primary domains through this critical sociotechnical lens. First, it will explore the imperative for pedagogical evolution, arguing that the crisis of assessment is, in fact, an opportunity to move beyond flawed models and embrace frameworks that foster critical AI literacy and process-oriented learning. Second, it will examine the dual role of GenAI as a research accelerator and the complex epistemological and ethical questions it raises regarding integrity, bias, and the very definition of scholarly knowledge. Finally, it will address the strategic challenge for universities to align with the 4IR workforce, arguing that institutions must resist neoliberal pressures to become mere skills factories and instead double down on their mission to cultivate the uniquely human skills of critical, creative, and ethical reasoning.

### **The Pedagogical Juncture: Beyond the 'Cat-and-Mouse' Game of Assessment**

The most immediate, visible, and arguably panicked response to Generative AI within higher education has been centred on pedagogy, specifically on the perceived collapse of traditional assessment. For generations, the university essay, the take-home exam, the lab report, and the coding assignment have been the cornerstone instruments for evaluating a student's comprehension, critical analysis, and synthesis of knowledge (Boud & Molloy, 2013). The arrival of GenAI, which excels at producing fluent, coherent, and often highly plausible text on demand, directly threatens the integrity and diagnostic validity of these methods. This challenge has sent a shockwave through faculties worldwide, raising widespread and legitimate concerns about academic misconduct and the very meaning of student-generated work (Cotton et al., 2023; Mbakwe et al., 2023).

This assessment crisis is not a new issue caused by AI; instead, it is a pedagogical booster that has revealed longstanding, often overlooked weaknesses in an assessment system based on proxy indicators of learning, rather than the learning process itself. Moving forward requires shifting from a defensive stance—focused on detecting and banning AI use—to a proactive pedagogical realignment. This involves two main steps: first, redesigning assessments to be resilient and inclusive of AI, and second, developing critical AI literacy as a core skill for all students. Initially, institutions responded to tools like ChatGPT mainly with bans (García-Peñalvo, 2023). This was soon followed by an AI-detection arms race, with universities adopting detection software to identify students using these tools. However, this strategy is both technically unreliable and pedagogically harmful.

Technologically, AI detection is a failing enterprise. The tools have demonstrated significant issues with accuracy and reliability. Studies have shown a troubling rate of false positives, where human-written text is incorrectly flagged as AI-generated (Gao et al., 2023). Most disturbingly, this unreliability is not neutral; it is deeply biased. Groundbreaking work by Liang et al. (2023) demonstrated that detection tools are systematically biased against non-native English speakers. The linguistic patterns, syntactic structures, and lower perplexity writing common to those mastering a second language are often misidentified as AI-generated, leading to false accusations that disproportionately harm international and multilingual students. This creates an unacceptable ethical and equity crisis, where our solution to academic misconduct actively discriminates against a vulnerable student population. Furthermore, as GenAI models become more sophisticated and are increasingly integrated into standard writing tools (e.g., Microsoft 365, Google Workspace), their outputs will become technically indistinguishable from assisted human writing, rendering the entire detection project obsolete.

From a pedagogical perspective, adopting a prohibitive stance is counterproductive. It portrays the student-educator dynamic as adversarial and overlooks a vital opportunity for meaningful engagement. As scholars like Mollick (2024) and Lodge & Bonsanquet (2024) have noted, such an abstinent approach ignores the fact that these tools will be widespread in graduates' careers. Banning them at the university level does a significant disservice, comparable to teaching accounting with an abacus instead of a spreadsheet. It hinders the ability to teach students how to use these powerful tools effectively, ethically, and critically.

The concern about the death of the essay (Marche, 2022) is a distraction. It tends to diagnose the student for using a tool while overlooking the teaching approach that encouraged its use. If

a task can be fully and satisfactorily completed with a GenAI tool, the issue is not the tool itself; it is the nature of the task. This indicates that the task probably focused on producing standard text rather than assessing critical thinking. Therefore, this crisis serves as a much-needed reminder to move away from superficial, product-focused assessments and to adopt models that involve students in the complex, messy, and distinctly human process of learning.

A more robust and sustainable response requires a fundamental pedagogical shift away from simple content recall and toward assessment models that are inherently resilient to, or even productively integrate, AI. This moves beyond the vague notion of authentic assessment and requires adopting specific, evidence-based frameworks that prioritise process, dialogue, and human-AI collaboration. First, there is a compelling case for a renewed focus on process-oriented pedagogy. Rather than evaluating only the final, polished product, this approach assesses the entire workflow of knowledge creation. Students may be required to submit a process folio or research log to document their learning journey. This might include:

- Initial brainstorming and mind-mapping (which may or may not use AI).
- The prompts they engineered to query GenAI models are complex.
- Their critical evaluation of the AI-generated outputs (e.g., identifying hallucinations, biased statements, or factual errors) is essential.
- A record of how they revised, rejected, or integrated that output into their own, original argument.

In this model, the use of AI is not considered cheating; it is a documented part of the process. The assessment target shifts from the text to the student's critical engagement with it. As Mollick (2024) suggests, the educator's role shifts from being a judge of quality to a coach of process, assisting students in learning to work with AI as a co-intelligence partner. Second, the rise of GenAI necessitates a return to dialogic and embodied assessment methods. The most AI-proof assessment is a live, synchronous conversation. This suggests a revival of the viva voce (oral defence), not just for doctoral dissertations but as a regular component of undergraduate and graduate modules (Holmes, 2023). A short, 15-minute viva allows an instructor to probe a student's understanding of their own written submission, asking them to:

- ❖ Paraphrase a complex argument from their paper.
- ❖ Defend their methodological choices.
- ❖ Explain how they connected two different sources.
- ❖ Reflect on what they found most challenging.

These tasks are beyond what an AI can do for the student. While resource-intensive, this method is arguably the most effective way to verify authorship and assess deep, internalised understanding (Lodge et al., 2023). In-person methods like traditional invigilated exams, in-class presentations, and problem-based learning within supervised studio settings are also gaining importance. Additionally, pedagogy should shift from designing AI-proof tasks — those that are beyond AI capabilities — to AI-inclusive tasks that require the critical use of AI. This fosters human-AI collaboration. For example, students could use GenAI to generate three different solutions for a complex case study. Their challenge would be to create a fourth, original solution that critiques and synthesises the AI-generated attempts, identifies each solution's flaws, and justifies their own superior approach. Another example involves giving students an AI-generated literature review to fact-check and critique, including spotting hallucinated citations, recognising biases, and identifying missing key sources. Such tasks directly measure higher-order critical thinking skills crucial in an AI-saturated environment.

This new pedagogical framework, whether process-oriented or AI-inclusive, relies heavily on a foundational skill: Critical AI Literacy. Merely teaching students how to operate GenAI (functional literacy) is not enough and can be risky. It is like giving someone a car without instructing them on traffic rules or the physics that control its movement. Therefore, Critical AI Literacy should be a core meta-skill for all students, regardless of their field (Baidoo-Anu & Ansah, 2023; Southworth et al., 2023).

This critical literacy surpasses basic prompt engineering by fostering a deeper, more conceptual understanding. First, students need to understand the technology behind these models. They should know, in simple terms, that these are not sentient beings, brains, or search engines, but are probabilistic systems—advanced pattern-matching tools that generate the most probable next word based on training data (Bozkurt, 2023). This knowledge helps counteract anthropomorphism and explains why models sometimes hallucinate (fabricate plausible-sounding information) or produce biased outputs. Second, students must recognise bias. They should be explicitly taught that GenAI models inherit biases from vast, uncurated internet datasets, which contain historical and systemic human prejudices related to race, gender, culture, and ideology. Critical questions like Whose perspective are missing? Moreover, what assumptions are embedded? Should be encouraged. Third, ethical use and data privacy are essential: Students need clear guidelines on ethical principles- distinguishing legitimate use, like brainstorming or grammar checks, from misconduct such as presenting AI work as their own. Additionally, they should understand that inputting sensitive personal or proprietary

research data into public AI tools may risk a data breach, as such information might be stored and used for future training (UNESCO, 2023).

Ultimately, the core challenge of GenAI in education is not technical; it is human. It pushes higher education to move away from measuring proxies of learning and focus on evaluating teaching directly. This transition urges us to replace an adversarial detection approach with a collaborative method centred on critical engagement. This prepares students for a future where their worth is in guiding, curating, and critically analysing content created by AI, rather than just producing content themselves.

### **The Research Nexus: Acceleration, Integrity, and Epistemological Shifts**

Beyond the immediate classroom disruptions, Generative AI is fundamentally transforming how knowledge is generated. Traditionally, academic research has relied on human inquiry. However, AI models now rapidly analyse large datasets, identify new patterns, and simulate complex systems at speeds and scales far beyond human capabilities (Agrawal et al., 2024). This creates a duality: on one hand, GenAI functions as a powerful tool to speed up research, making intricate computational tasks more accessible and fostering discovery pathways. On the other hand, it prompts critical, unresolved questions about academic integrity, authorship, equity, and even challenges the core epistemological principles of scholarly work. GenAI is more than merely a tool; it functions as an active sociotechnical agent within the research process. Its incorporation demands a careful review of not only our research methods but also our definitions of knowledge, the authority to produce it, and the measures to safeguard its integrity in a time dominated by automated, probabilistic truth.

The impact of AI in research is unmistakable. For instance, DeepMind's AlphaFold used AI to predict the 3D structures of over 200 million proteins, resolving a longstanding biological challenge that had puzzled scientists for decades (Jumper et al., 2021). This progress shifted protein folding from a costly, labour-intensive experimental process to a quick computational approach, ushering in a new era of AI-powered discovery. Generative AI expands this potential across nearly all academic disciplines, acting as a digital partner that can help or even automate many research tasks (Korngiebel & Mooney, 2021). In medicine and the life sciences, GenAI can analyse millions of research papers to perform systematic reviews quickly, identify subtle connections between studies to generate new hypotheses, or design novel molecular structures for drug discovery (Mbakwe et al., 2023). In social sciences, GenAI assists in analysing large

qualitative datasets, like interview transcripts or open-ended surveys, by consistently identifying emerging themes that support human coding. In computer science, it improves code creation, debugging, and refinement, making advanced computational techniques more accessible to researchers in other fields (Mhlanga, 2023). In digital humanities, GenAI explores extensive textual and visual collections to uncover hidden literary patterns, stylistic influences, or visual trends. It also functions as a creative collaborator, producing new forms of digital art or text and expanding the scope of humanistic research (Manjavacas & Fonteyn, 2022).

This acceleration could serve as a significant democratising force. Tasks like complex computational modelling, data analysis, or systematic reviews — once requiring specialised expertise, substantial funding, or large research teams — may become more accessible to a broader range of scholars. This shift could lower entry barriers for researchers from resource-limited institutions or the Global South, fostering increased global participation in science (Biswas, 2023). However, this optimistic view relies heavily on equitable access to the latest models, which are often proprietary and expensive to operate, risking a new AI divide. This rapid advancement also raises critical questions about research ethics, particularly regarding authorship and academic integrity. Using GenAI to draft or revise major parts of a manuscript challenges traditional ideas of authorship and who —or what —can be considered an author.

In response, major organisations focused on publication ethics, and journal editors quickly implemented policies. The International Committee of Medical Journal Editors (ICMJE, 2024) emphasises that authorship must meet four criteria: (1) making significant contributions, (2) critically revising or drafting the work, (3) giving final approval of the publication, and (4) agreeing to take responsibility for all aspects of the work. As Eke (2023) and Thorp (2023) convincingly argue, AI tools cannot fulfil these criteria, particularly the fourth, since an AI cannot be held accountable, take responsibility for the work, verify its integrity, or participate in the peer-review process. Therefore, organisations like the Committee on Publication Ethics (COPE) and top journals such as Science and Nature agree that AI cannot be listed as an author (Nature, 2023).

However, this policy introduces a significant grey area. If AI cannot be regarded as an author, its use still needs to be recognised. However, where is the line for acknowledgement? Is it merely using AI to correct grammar, like a word processor? Or to brainstorm ideas? Or to draft a preliminary literature review? Or even to write the entire methods section? The lack of clear, field-specific standards on when and how to disclose Genai use leads to a crisis of transparency

(Stokel-Walker, 2023). Without such guidelines, the integrity of the scholarly record is at risk. An AI-assisted paper that omits disclosure of the tool's use effectively amounts to a new form of plagiarism- passing off the model's work, trained on millions of uncredited sources, as the human authors. Among the most dangerous threats posed by GenAI in research is its potential to magnify systemic bias. This issue goes beyond authorship debates, challenging the validity and fairness of the knowledge produced.

GenAI models are not inherently objective sources of truth. They rely on large-scale internet datasets that reflect a simplified and biased version of human knowledge (Manjavacas & Fonteyn, 2022, p. 7). As Safiya Noble (2018) showed in *Algorithms of Oppression*, this data contains numerous historical and systemic biases related to race, gender, and culture. There is a well-documented risk that GenAI models will not only reproduce but also amplify these biases in their research applications. For example:

- An AI model used in medical research, trained mainly on data from white, male populations, might produce results, or develop drugs that are less accurate or effective for underrepresented groups, thereby worsening health disparities (Mhlanga, 2023).
- An AI tool used for systematic reviews may hallucinate or prioritise citations from high-prestige, English-language journals, further marginalising scholarship from the Global South, non-English-language sources, and critical or niche fields (Biswas, 2023).
- An AI used in social science modelling, trained on biased historical data, may produce predictive models that reinforce discriminatory policing, hiring, or loan practices.

This results in a "Matthew Effect" for knowledge, reinforcing dominant perspectives while marginalising others (Agrawal et al., 2024). Additionally, the black box problem exacerbates this issue. The decision-making processes within large language models are so intricate that they remain obscure, even to their developers. This lack of transparency directly conflicts with the principles of the scientific method, which relies on openness, reproducibility, and the capacity to explain how specific results are obtained.

Moreover, these models can hallucinate, meaning they may confidently present plausible yet false information, often with fabricated citations (Adesina et al., 2023). During research, such hallucinations are especially risky. Suppose a fabricated fact or ghost citation in an AI-generated literature review goes unnoticed by the human reviewer. In that case, it can become part of the scholarly record, be cited by others, and propagate, polluting the knowledge ecosystem with AI-produced misinformation that is hard to remove.

Therefore, incorporating GenAI into research is more than just using a tool; it represents an epistemological intervention. It urges the academic community to create new transparency measures, vigilantly check for algorithmic bias, and uphold human accountability. In the twenty-first century, the researcher's role is expanding; it is no longer solely about discovering knowledge but also about critically validating and ethically managing the knowledge produced by our advanced yet imperfect computational counterparts.

### **Strategic Alignment: Preparing Graduates for the 4IR Workforce.**

The third disruption stems from the need for universities to align their curricula with the changing demands of a 4IR-driven workforce. As AI rapidly automates routine cognitive and analytical tasks, the labour market is undergoing a fundamental transformation, generating widespread concerns about a skills gap and the relevance of traditional university degrees (Schwab, 2017). This puts significant pressure on higher education to become more flexible, responsive to market needs, and focused on graduating job-ready students. The World Economic Forum's (2023) Future of Jobs Report highlights this shift, noting a meaningful change in the most sought-after skills. While technical and AI-specific skills continue to grow in importance, the report emphasises the increasing value of human-centric skills. Creative thinking is now seen as the most crucial, followed by resilience, flexibility, and agility, motivation and self-awareness, and curiosity and lifelong learning. This creates a paradox: as technology advances, so does the need for soft, uniquely human skills. While universities must adapt, relying solely on market-driven solutions would be a mistake. First, we will critically analyse the skills gap narrative as a neoliberal appropriation of the university's mission. Then, we will assess solutions like micro-credentialing and conclude that the key strategic response to the 4IR is not to abandon the university's core mission but to reinvigorate it, focusing on cultivating deep, critical, and ethical human potential.

The debate over the skills gap is extensive, with industry leaders and policymakers claiming that universities do not sufficiently prepare students with the precise technical skills needed today. They see universities mainly as sources of human capital development. The proposed fix is to make universities more flexible so that curricula can quickly adapt to current job market demands. However, this view is both myopic and driven by ideological motives. Critical higher education scholars have long argued that this standpoint reduces universities to neoliberal vocational training, prioritising civic and humanistic functions over the immediate, often

transient, needs of the market (Brown, 2015; Giroux, 2014). For assorted reasons, trying to close the skills gap is ultimately futile.

The rapid pace of technological change in the 4IR means that specific technical skills, like proficiency in a particular coding language or AI platform, learned by a first-year student can become outdated by graduation. If a university focuses its curriculum on current job-specific skills, it is essentially preparing students for jobs that no longer exist. This view subtly and openly devalues disciplines that are not directly tied to specific careers, such as the humanities and social sciences. It creates a false divide between useful STEM degrees and useless philosophy or history degrees, ignoring the fact that essential skills valued by industry leaders—critical thinking, creativity, and ethical reasoning—are at the core of these useless fields (Aoun, 2017). A university is more than a job training centre; it should be a vital part of a democratic society, a space where students develop critical citizenship, grapple with complexity, question authority, and deal with uncertainty (Giroux, 2014). An exclusively vocational approach neglects this civic role, resulting in technically skilled but critically disengaged individuals, mere cogs in the machine.

Universities must recognise the significance of the labour market in their strategic planning for the Fourth Industrial Revolution (4IR). Their responses should be thoughtful and initiative-taking rather than passive, emphasising a long-term vision focused on human well-being instead of short-term financial gains like quarterly profits. A widely suggested solution to the skills gap is adopting more flexible credentialing methods, such as micro-credentials- short, targeted, skills-based courses that can be completed on demand and combined into larger qualifications. Proponents see this as a way to dismantle traditional, monolithic degrees, enabling just-in-time curriculum updates that keep pace with rapid technological changes (Gallagher, 2018). This approach aims to provide learners with flexibility and continuous learning pathways, fostering lifelong education. The movement has gained momentum, with private ed-tech firms and established universities investing heavily in micro-credential platforms. While offering clear benefits in flexibility and accessibility, scholars highlight significant risks (Lytle & Lorn, 2023). Fragmenting education into small 'skills' can detach knowledge from its theoretical and ethical context, which is essential for understanding. For example, a student might learn how to operate a data analysis model but not grasp the underlying statistical theory, assumptions, or ethical considerations. This model could also promote a gig economy for education, reducing faculty to 'content providers' for standardised modules and training learners in functions rather than disciplines. A curriculum composed

solely of stacked micro-credentials may lack the coherence, scaffolding, and deep theoretical foundation of a comprehensive degree program. This approach risks producing technicians skilled in tools but lacking the essential knowledge for true innovation and adaptation as tools evolve.

While micro-credentials can be a helpful addition to a traditional degree—offering pathways for continuous professional development or serving as on-ramps to higher education—they do not replace it. Fully embracing the unbundling concept risks turning universities from institutions of deep, transformative learning into superficial, just-in-time skills markets. If chasing specific technical skills is ultimately unproductive, and unbundling threatens to fragment knowledge, what is the best strategic approach for the 4IR? The answer, as the World Economic Forum (2023) report indicates, lies in a paradox: as AI automates routine cognitive tasks, the most valuable, lasting, and irreplaceable skills are those that are distinctly human. Consequently, the role of university education should focus on intentionally and rigorously cultivating these higher-order human skills. This is not a new or radical idea; rather, it signifies a resurgence of the timeless, humanistic purpose of a university.

Curricula, across all disciplines, must be explicitly redesigned to cultivate advanced critical thinking and complex problem-solving. This includes the ability to ask pertinent questions, define novel problems, synthesise information from diverse sources, and navigate ambiguity and uncertainty. In an era of AI-generated answers, individuals who can frame better questions become invaluable. Additionally, curricula should promote creativity, curiosity, and originality. Since GenAI is inherently probabilistic and derivative—remixing existing data—true creativity—the capacity to envision innovative ideas, connect seemingly unrelated concepts, and make intuitive leaps—remains a distinctly human trait. Education must encourage this, moving beyond standardised testing and correct answers. It must also cultivate social and emotional intelligence. Skills such as communicating, empathising, collaborating, and leading human teams cannot be outsourced to algorithms. As work becomes more collaborative and project-based, these social competencies are increasingly critical for success. Finally, curricula must ensure the development of ethical and moral reasoning (phronesis), which is arguably the most vital skill of all. The Fourth Industrial Revolution (4IR) is characterised by powerful technologies that generate new ethical dilemmas—regarding data privacy, algorithmic bias, and genetic engineering, for example. Universities have a moral obligation to graduate individuals capable of making principled judgments about applying these technologies. This entails moving beyond superficial tech ethics modules and embedding ethical reasoning—a

form of Aristotelian phronesis or practical wisdom—throughout the entire curriculum (Aoun, 2017).

This strategic realignment emphasises blending technical literacy with the humanities instead of separating them. It includes interdisciplinary courses like Tech Ethics, Computational Social Science, or AI and Philosophy to create graduates who can both develop technology and critically assess it. The ideal 4IR graduate is not just a STEM-trained engineer or a humanities-trained critic; they are a humanistic technologist who can bridge the gap between the two cultures and guide innovation rooted in human values.

#### Conclusion: From Reactive Posture to Proactive Stewardship

Generative AI is transforming higher education in the 4IR, representing a sociotechnical shift that requires proactive engagement rather than defensive reactions. In pedagogy, the assessment crisis calls for process-oriented, dialogic, and AI-inclusive methods, underpinned by critical AI literacy. In research, GenAI accelerates discovery but demands safeguards against authorship issues, hallucinations, and bias. For workforce readiness, universities should reject neoliberal vocationalism and prioritise enduring human skills like critical thinking and ethical reasoning. To progress, institutions must: (1) establish adaptable ethical governance; (2) invest in faculty and student upskilling for AI integration; and (3) reaffirm the university's humanistic mission to foster irreplaceable human potential. Given GenAI's rapid evolution, key research priorities include longitudinal studies on learning outcomes, equity audits of AI tools, epistemological analyses across disciplines, and comparative evaluations of governance models. Through human-centric stewardship, universities can harness the 4IR to enhance, rather than diminish, human capabilities.

#### References

Adesina, I., Gbadamosi, S. O., Foster, C., & Gbadamosi, A. (2023). 'Hallucinations' in large language models: A systematic review of a potentially catastrophic barrier to AI-driven progress. arXiv preprint arXiv:2311.13010. <https://doi.org/10.48550/arXiv.2311.13010>

Agrawal, A., Gans, J., & Goldfarb, A. (2024). *The economics of artificial intelligence: An agenda*. University of Chicago Press.

Aoun, J. E. (2017). *Robot-proof: Higher education in the age of artificial intelligence*. MIT Press.

- Baidoo-Anu, D., & Ansah, L. O. (2023). Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI*, 7(1), 52–62. <https://doi.org/10.61969/jai.v7i1.303>
- Biswas, S. (2023). Generative AI and the Global South: A new colonial frontier? *Development*, 66(1), 1–8. <https://doi.org/10.1057/s41301-023-00388-9>
- Boud, D., & Molloy, E. (2013). Rethinking models of feedback for learning: The challenge of design. *Assessment & Evaluation in Higher Education*, 38(6), 698–712. <https://doi.org/10.1080/02602938.2012.691462>
- Bozkurt, A. (2023). Generative artificial intelligence (GenAI) in education: The potential of AI, the hype of the 21st century, and the new era of education. *Journal of Learning for Development*, 10(3), i-ix.
- Brown, W. (2015). *Undoing the demos: Neoliberalism's stealth revolution*. Zone Books.
- Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2023). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 1–12. <https://doi.org/10.1080/14703297.2023.2190148>
- Eke, D. O. (2023). AI and the authorship conundrum: A case for accountability. *Journal of Medical Ethics*, 49(12), 856-857. <https://doi.org/10.1136/jme-2023-109015>
- Gallagher, S. (2018). *Educational credentials come of age: A developmental and systems view of micro-credentialing*. Northeastern University. <https://files.eric.ed.gov/fulltext/ED590924.pdf>
- Gao, C. A., Howard, F. M., Markov, N. B., Dyer, E. C., Ramesh, S., Luo, Y., & Pearson, A. T. (2023). Comparing AI-generated and human-authored text for web content, news, and online messaging. arXiv preprint arXiv:2307.09288. <https://doi.org/10.48550/arXiv.2307.09288>
- García-Peñalvo, F. J. (2023). The assessment of the impact of artificial intelligence in education. *Education in the Knowledge Society*, 24, e312. <https://doi.org/10.14201/eks.312assessment>
- Giroux, H. A. (2014). *Neoliberalism's war on higher education*. Haymarket Books.
- Holmes, B. (2023, April 19). In the age of AI, we need to bring back the 'viva'. Times Higher Education. <https://www.timeshighereducation.com/campus/age-ai-we-need-bring-back-viva>

International Committee of Medical Journal Editors. (2024). *Recommendations for the conduct, reporting, editing, and publication of scholarly work in medical journals*. ICMJE. <https://www.icmje.org/icmje-recommendations.pdf>

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., ... & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>

Korngiebel, D. M., & Mooney, S. D. (2021). Considering the possibilities and pitfalls of artificial intelligence in clinical and translational research. *Clinical and Translational Science*, 14(4), 1184–1188. <https://doi.org/10.1111/cts.13007>

Latour, B. (2005). *Reassembling the social: An introduction to actor-network theory*. Oxford University Press.

Liang, W., Yuksekogonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. arXiv preprint arXiv:2304.02819. <https://doi.org/10.48550/arXiv.2304.02819>

Lodge, J. M., & Bonsanquet, A. (2024). *Assessing in the age of artificial intelligence*. Routledge.

Lodge, J. M., Hases, G., & D'Souza, L. (2023). Assessing in the age of AI. *Journal of University Teaching & Learning Practice*, 20(4), 1–6. <https://doi.org/10.53761/1.20.4.01>

Lytle, R. J., & Lorn, C. (2023). Unbundled and unmoored: The potentials and pitfalls of microcredentials. *New America*. <https://www.newamerica.org/education-policy/reports/unbundled-and-unmoored/>

Manjavacas, E., & Fonteyn, L. (2022). A computational perspective on humanistic text analysis. arXiv preprint arXiv:2211.02523. <https://doi.org/10.48550/arXiv.2211.02523>

Marche, S. (2022, December 6). The college essay is dead. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2022/12/chatgpt-ai-writing-college-student-essay/672361/>

Mbakwe, A. B., Lourentzou, I., Celi, L. A., Mechanic, O. J., & Dagan, A. (2023). The future of academic medical publishing: How generative artificial intelligence will reshape the authoring, editing, and peer review landscape. *JMIR Medical Education*, 9(1), e48630. <https://doi.org/10.2196/48630>

- Mhlanga, D. (2023). Open AI in higher education: A comparative analysis of ChatGPT and its potential impact on the university's teaching and learning. *Education and Information Technologies*, 28(11), 15009–15031.
- Mollick, E. R. (2024). *Co-intelligence: Living and working with AI*. Harvard Business Review Press.
- Nature. (2023). Tools such as ChatGPT threaten transparent science; here are our ground rules for their use. *Nature*, 613(7945), 612. <https://doi.org/10.1038/d41586-023-00191-1>
- Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
- OpenAI. (2023). GPT-4 technical report. arXiv preprint arXiv:2303.08774. <https://doi.org/10.48550/arXiv.2303.08774>
- Schwab, K. (2017). *The Fourth Industrial Revolution*. Portfolio.
- Selwyn, N. (2022). The postdigital university: A personal re-appraisal. *Postdigital Science and Education*, 4(3), 553-571. <https://doi.org/10.1007/s42438-022-00331-5>
- Southworth, J., Zaidi, M., & Sallit, S. (2023). A guide to generative AI in education. The Institute for Ethical AI in Education. <https://www.e-a-i.org/reports/a-guide-to-generative-ai-in-education>
- Stokel-Walker, C. (2023). AI and the future of scientific publishing. *Nature*, 615(7954), 785–787. <https://doi.org/10.1038/d41586-023-00816-5>
- Thorp, H. H. (2023). ChatGPT is fun, but not an author. *Science*, 379(6628), 119. <https://doi.org/10.1126/science.adg7879>
- UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386618>
- World Economic Forum. (2023). *Future of Jobs Report 2023*. WEF. <https://www.weforum.org/publications/future-of-jobs-report-2023/>