

FORGING THE FUTURE

How AI is rewriting
the rules
of knowledge,
expertise, and
practice.

● **AI is reshaping how humans create knowledge and learn. While it can accelerate understanding, bypassing the need to internalise facts could result in the erosion of critical thinking, originality, and deep engagement.**

●● **Large Language Models (LLMs) cannot create output beyond the data on which they are trained. Effective use of AI depends on asking the right questions based on contextual awareness, ethical judgement, and intellectual curiosity.**

●●● **Universities should redesign curricula for thoughtful AI use and reshape assessments to value dimensions that resist automation – originality, justification, and interpretation.**

Artificial Intelligence (AI) is not just automating tasks – it is transforming how we learn, how we create knowledge, and how we make decisions. The implications are especially profound for the social sciences, which lie at the intersection of data, behaviour, and meaning. For leaders in business, government, and education, understanding this transformation is no longer optional. It is essential.

Imagine a university student beginning a term paper or a faculty member starting a research article. They launch ChatGPT to “just get started”, and within seconds, a structured essay or draft appears. What once required hours of reading and thinking is now made possible in minutes. But what is being lost when the machine thinks before they do? This quiet shift – from human synthesis to machine prompting – is emblematic of a much broader change now underway across society.

We find ourselves at an inflection point. The rise of AI, particularly generative tools like LLMs, compels us to rethink the foundations of knowledge transmission, research, and application. This is a pivotal moment for all of us as social scientists, experts, and managers navigating a future shaped by intelligent machines.

AI is reshaping how we teach, learn, generate knowledge, and make decisions. In education, it challenges traditional pedagogy by offering tools that can explain concepts, summarise arguments, and even draft essays – changing not just what students do, but how they think. In research, AI accelerates tasks such as literature reviews and data analysis, but also raises concerns about bias, transparency, and over-reliance on machine judgement. Beyond academia, AI is now being widely adopted in sectors such as business, law, and public policy, where it is used to generate reports, forecast outcomes, and support complex decisions, often without fully appreciating the human purpose or social context. While AI excels at aggregation and prediction, the human ability to ask new questions, interpret meaning in context, and challenge assumptions remains essential for real insight.

AI'S DUAL ROLE IN LEARNING: CATALYST AND CRUTCH

For generations, education has centred on the transmission of curated knowledge and the cultivation of judgement. Professors lecture, students analyse, and learning occurs through synthesis and debate. Generative AI (or genAI) is transforming this classroom arrangement – not by replacing it, but by reshaping how students engage with knowledge. With genAI tools, students can now ask complex questions and receive instant answers, draft essays from a single prompt, and summarise dense readings in seconds. This introduces a new dynamic in education: AI can accelerate learning, but also bypass the very processes through which deep understanding is traditionally developed.

Take prompting, for example. In the context of genAI, prompting refers to the act of crafting instructions or queries that guide the model's response. More than a technical input, prompting is becoming a cognitive skill in its own right. Well-constructed prompts – those that specify goals, provide context, and impose constraints – tend to produce outputs that are more accurate, nuanced, and relevant. In this way, effective prompting encourages students to clarify their intent, which can deepen their reasoning and reflection. Recent studies suggest that students trained in prompt engineering show higher engagement and comprehension when using AI tools.¹ Yet this also presents a pedagogical dilemma: when prompting replaces writing,

are students still learning to build an argument, or merely outsourcing it?

One of the most immediate effects of AI on learning is a shift from recall to recognition. Students no longer need to internalise facts in order to access them. Whether this shift is harmful or beneficial depends on what kind of learning we value. Some cognitive scientists argue that freeing learners from rote memorisation can allow for more conceptual and applied learning. But others caution that the loss of foundational knowledge weakens long-term retention and critical synthesis. Without internal scaffolding, students may find it harder to judge the quality of what an AI prompt returns, making them more vulnerable to plausible-sounding errors or hallucinations.

Similarly, when students rely on AI to produce polished outputs, the emphasis may move from process to product. Learning becomes transactional, focussing on achieving a desired outcome, such as an essay, a summary, or a solution, rather than grappling with ambiguity or constructing understanding through effort. In the short term, this may appear efficient. But over time, it risks eroding the intellectual qualities that education is meant to cultivate – curiosity, perseverance, and the ability to tolerate complexity.

Educators must therefore recalibrate not just their assessments but their entire pedagogical approach. Rather than retaining their traditional roles as gatekeepers of knowledge, teachers are becoming guides to its navigation. This means designing assignments that

resist automation: those that ask students to justify, reflect, or apply concepts in novel settings. It means teaching students to evaluate AI outputs critically, cross-check sources, and revise generative responses, rather than submit them wholesale. For instance, Columbia University's 'University Writing' course has integrated ChatGPT into the curriculum to enhance students' metacognitive skills.² Students engage with AI-generated content, analysing and reflecting on the outputs to deepen their understanding of both the subject matter, and the capabilities and limitations of AI tools. Such approaches embrace the opportunity that AI provides to teach meta-cognition – not just *what* to learn, but *how* to think in a world where information is abundant but understanding is scarce.

WHO CREATES KNOWLEDGE NOW?

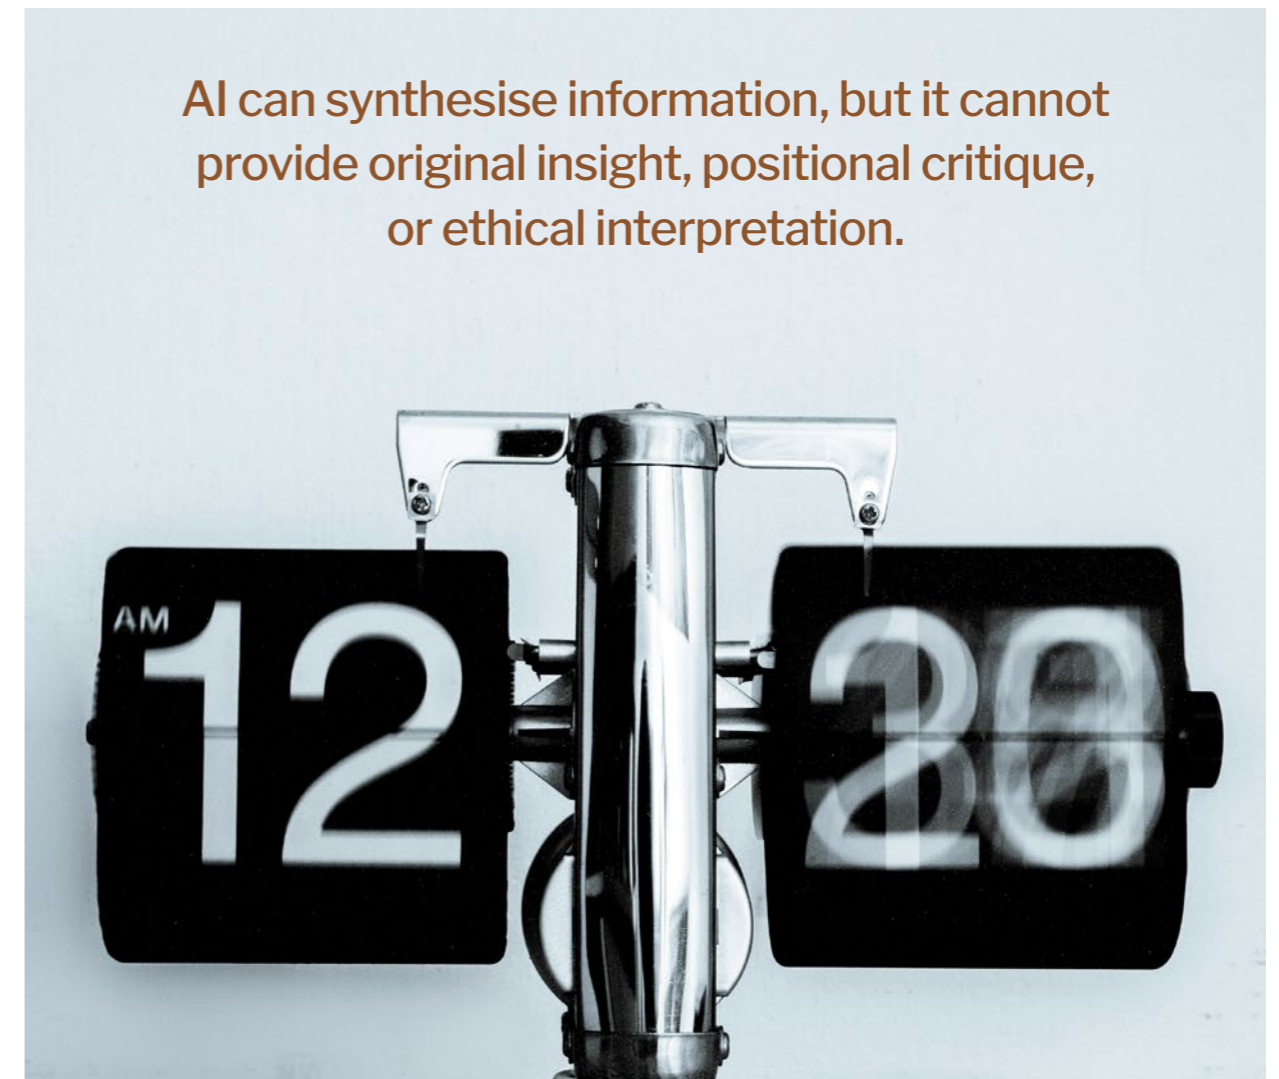
AI is rapidly reshaping the academic research landscape. In fields from economics to linguistics, researchers are using AI to search literature, summarise content, analyse text, and even generate hypotheses. A Stanford University study published in 2024 found that the number of peer-reviewed publications engaging with AI across 20 scientific disciplines, including political science and psychology, quadrupled between 2020 and 2022.³ This growth reflects AI's power not only to accelerate traditional workflows, but also to enhance the precision of analytical methods and uncover patterns at scale that were previously invisible to researchers.

One clear benefit is in predictive analytics. AI models, especially LLMs and transformer-based architectures, can ingest vast quantities of data and generate forecasts, identify latent structures, and refine classification models. In my own field of computational social science, tools like quanteda, a library of open-source software for the "quantitative analysis of textual data" that I have spent over a decade developing, have enabled large-scale text analysis in political science, marketing, communication, and law.⁴ While quanteda predates the rise of genAI, it shares with these newer tools the capacity to automate formerly labour-intensive tasks such as content coding, sentiment analysis, and topic modelling. In many ways, genAI has now moved

the state of the art well beyond what quanteda's earlier generation of machine learning techniques could accomplish. But it is also enabling further innovation; we now use AI-powered coding assistants to accelerate the development of quanteda itself. In this way, AI is not only transforming research workflows but also helping to address the technical challenges it has created.

However, not all research domains are equally affected or enhanced by AI. In more quantitative and model-driven disciplines like economics, AI tends to be an evolutionary tool. It speeds up simulation, improves estimation, and reduces the time taken from idea to result. The same is true in computational political science, where algorithmic approaches

AI can synthesise information, but it cannot provide original insight, positional critique, or ethical interpretation.



already dominate parts of the methodological toolkit.

By contrast, in more interpretative fields such as anthropology, history, or critical sociology, the role of AI can feel more disruptive. These disciplines emphasise reflexivity, meaning-making, and cultural context. When AI tools are used to summarise ethnographies or simulate theoretical arguments, they risk bypassing the deeply situated reasoning processes that are central to qualitative research. Rather than enhancing existing workflows, they may impose a logic of fluency and efficiency that runs counter to the epistemic values of the discipline itself.

This divergence raises a broader question about what constitutes legitimate knowledge production. AI can synthesise information, but it cannot provide original insight, positional critique, or ethical interpretation. For social scientists, this is not just a technical concern; it also goes to the heart of how we understand truth, meaning, and scholarly contribution.

IS AI IMITATING LIFE, OR IS LIFE IMITATING AI?

Now that AI can generate text that appears indistinguishable from that written by experts, it raises the question: If the machine can synthesise existing knowledge faster and more fluently than we can, what value remains in human scholarship and leadership?

The answer lies in what AI cannot do.

Social scientists – and, by extension, informed leaders –

derive value from three capacities: deep expertise, critical thinking, and originality. While AI can replicate the outputs associated with these capacities, it lacks the ability to carry out the underlying processes that make them meaningful. As authors McKendrick and Thurai argue in a recent *Harvard Business Review* article, AI’s limitations in capturing intangible human factors underscore the necessity for human judgement in real-life decisions.⁵ These limitations become especially evident when we consider the context in which judgement is exercised, which is through the ability to ask original questions and interpret information within its proper context.

Asking the right questions

AI is trained to answer questions, but it cannot ask them – not in the way that humans can. True progress in science, policy, or business comes not from generating faster answers, but from asking better, more original questions. Consider this: a model like DALL-E might generate a visually compelling image if asked to “paint a 16th-century European woman in the style of Leonardo da Vinci”, but would it ever decide to invent the Mona Lisa? genAI can imitate, extrapolate, and remix existing styles, but it cannot originate the kind of cultural moment that reframes how the world is seen. Originality lies in perspective, that is, the ability to frame familiar phenomena in unfamiliar ways, and that depends on context, dissent, and the willingness to challenge dominant narratives.

This is precisely where the limitations of AI are most visible. Because it learns from the past, it tends to reinforce it. LLMs are trained on vast archives of existing content, meaning that their outputs reflect historical consensus, not emergent insight. This conservatism risks crowding out minority perspectives and stifling innovation. In research, policymaking, and even art, the ability to ask questions that break away from inherited assumptions remains fundamentally human.

Understanding meaning in context

Clifford Geertz, one of the 20th century’s most influential anthropologists, famously distinguished between a wink and a twitch: both are physically identical movements, but only one carries intentional meaning.⁶ The distinction, drawn from philosopher Gilbert Ryle, lies in the unobservable yet essential layer of social intention.⁷ A wink signals, implies, and invites interpretation; a twitch does not. As Sidnell and Enfield note, this difference hinges on social accountability: a wink can be misread or challenged, while a twitch simply happens.⁸ Meaning, in this sense, arises from context: cultural, interpersonal, and interpretative. For Geertz, this was the essence of “thick description” – not just recording behaviour, but understanding why it matters. This idea underpins much of qualitative research: insights are embedded in context that resist reduction. AI, in contrast, can mimic the surface of expression but cannot access its depth. It produces what might be

called thin description or fluent summaries without the capacity to interpret significance.


This leads to a deeper concern: if AI can mimic the outward expression of human thinking, what becomes of the intellectual space once claimed as uniquely human? Jean-Paul Sartre, in his 1944 existentialist play *No Exit*, explored the unsettling idea of selfhood as mediated through the gaze of others, where hell is not a place of torment, but is being constantly seen and defined by someone else. Today, that gaze is machinic. When a language model replicates our tone, structure, or style, it reflects a version of thought stripped of context, intent, and doubt. The machine exhibits fluency without understanding, yet its output can appear authoritative, even insightful.

For leaders and educators, this poses a unique risk – mistaking imitation for comprehension, or worse, allowing it to shape decisions without human interrogation. We trained the AI on our thinking, but through over-reliance and what some call

intellectual deskilling, we now risk training future generations on AI. The way to break this circle is not to reject AI, but to clearly understand its limitations.

WHAT CAN UNIVERSITIES DO?

Universities face a significant challenge of adapting institutional policies and pedagogical practices to a technology that is evolving faster than most curricula, assessment regimes, or academic codes of conduct. genAI is not a passing trend but a foundational shift in how we access, apply, and even construct knowledge. This means rethinking curriculum design, assessment methods, and academic policies to ensure that the core objectives of higher education, such as critical thinking, intellectual integrity, and the pursuit of understanding, are not only preserved, but strengthened. Courses should be updated to explicitly teach students how to work with AI systems thoughtfully – how to craft prompts, assess the reliability of outputs, and understand the limitations of models.



Academic integrity frameworks
must evolve to reflect the fact that it is
not only students who are using AI;
faculty are as well.

At the same time, assignments and assessments should be redesigned to value originality, justification, and interpretation – dimensions of student work that resist automation. Universities must continue to train students not only to use AI, but to think around it, with the same creativity, scepticism, and contextual awareness that makes human inquiry indispensable.

Universities are beginning to grapple with these tensions. Some are integrating AI literacy into research training programmes, while others are revising ethical guidelines to address questions of transparency, accountability, and authorship. Research integrity offices and funding agencies are now issuing guidance on how AI can – and cannot – be used in scholarly work, signalling that the norms of knowledge production are being actively renegotiated.

Academic integrity frameworks must evolve to reflect the fact that it is not only students who are using AI; faculty are as well. Instructors increasingly turn to generative tools for drafting lecture slides, writing recommendation letters, preparing course materials, providing feedback or grading, and increasingly, conducting and writing up research. A growing industry of AI-powered tools now caters to academics under intensifying pressure to produce and publish. Whether these tools serve as a catalyst for scholarly productivity or a crutch that undermines critical engagement depends on how institutions guide their use.

Mustafa Suleyman, AI policy thinker and co-founder of DeepMind, argues that transformative technologies like AI pass through three phases: they become inevitable, then ubiquitous, and eventually invisible.⁹ Precisely because AI will operate beneath the surface of our systems and decisions, it is critical that students and educators that train them develop the skills to use it well, and cultivate the ethical awareness to question how it is being deployed around them.

This makes it essential for universities to establish clear and consistent standards of appropriate use for everyone involved in education and research. Integrity guidelines should not merely police student misconduct, but foster a shared culture of transparency and responsible engagement with AI across the academic community. That means supporting faculty not only with technical training, but also with opportunities to reflect on how AI intersects with their pedagogical values and scholarly responsibilities.

At a broader level, universities must model responsible innovation. This includes supporting interdisciplinary research on AI’s societal impacts, experimenting with AI-enhanced pedagogy, and collaborating with industry and policymakers to help shape ethical standards. But it also means preparing students, educators, and researchers to live and work in a world where AI will be deeply embedded in economic and institutional life.

Beyond academia, the age of AI confronts other sectors with parallel challenges.

WHAT CAN THE WORLD OF PRACTICE DO?

AI is no longer confined to research labs or tech start-ups; it is being rapidly adopted in corporate boardrooms, government agencies, media organisations, and others. Yet many leaders outside academia misunderstand what AI systems actually do. There is a risk that AI-generated outputs are treated as authoritative or insightful, when in fact such systems are simply predicting future outcomes based on patterns in past data. This misperception risks delegating strategic decisions to tools that possess no judgement of their own.

Consider a policymaker using AI to score funding applications. If the system is trained on past decisions, it may favour familiar formats and conventional topics, inadvertently penalising novel or dissenting ideas. Or take a company screening job candidates using an AI model shaped by prior hiring patterns – it might efficiently reproduce old biases while appearing objective. In both cases, AI is not offering insight. It is extrapolating from precedent, thereby reinforcing the status quo rather than challenging it.

Used responsibly, however, AI can augment a skilled workforce. It can accelerate research, support strategic decision-making, and make complex analysis more accessible. But it must not become a substitute for human judgement. Over-reliance risks producing a ‘deskilled’ workforce that defers to fluent machine outputs without understanding their provenance or limitations. The role of leadership is

not simply to adopt AI, but to ensure it enhances insight without replacing thoughtful decision-making.

Singapore offers a compelling case study on how this can be done. With its strong institutions, culture of anticipatory governance, and emphasis on education and innovation, Singapore is well-positioned to model how AI can be integrated responsibly into public and economic life. Already, AI is being applied in areas such as urban planning, public health, and financial regulation. What distinguishes Singapore’s approach is not just the pace of adoption, but the emphasis on aligning technology with societal goals – through public engagement, regulatory clarity, and cross-sector collaboration.

Initiatives like the Model AI Governance Framework for Generative AI, developed by the Infocomm Media Development Authority (IMDA) and the AI Verify Foundation, offer one example of how national policy can help guide the responsible deployment of AI at scale.¹⁰ Singapore’s approach reflects an understanding of how a society can harness AI not only to boost productivity and efficiency, but also to deepen trust, empower citizens, and expand access to opportunity. That requires strong public institutions, a commitment to ethics and inclusion, and a clear understanding that AI is not an end in itself. It is a tool, and as with any technology applied to public life, its application must be guided and held accountable by human judgement and anchored in democratic values.

THE WAY FORWARD: AI AS CRUTCH OR CATALYST?

The decades to come will be defined not just by the spread of AI, but by how we respond to it. Will we treat AI as a crutch, relying on it for abilities we no longer cultivate, or as a catalyst for greater productivity and a tool for deeper insight? Will we prioritise speed and scale, or context and care?

Two emerging frontiers underscore the urgency of these questions. First, advances in hardware, especially specialised chips and the potential of quantum computing, promise to vastly increase the power of AI. This raises profound questions about control, access, and governance. Second, as AI systems grow more fluent and interactive, they increasingly raise philosophical and ethical questions once confined to science fiction: Can machines be sentient? Should they have moral standing? What obligations do we have to entities that appear to think, feel, or reason?

These are no longer abstract inquiries. They demand the insights of philosophers, ethicists, and social scientists who can explore not only what AI can do, but what it should do, and how it redefines what it means to be human.

As thought leaders and stewards of educational and institutional systems that shape both the present and future generations, we have an opportunity to lead by example. This means investing in research that bridges disciplines, education that fosters critical thinking, and policies that guide the development of AI responsibly.

The future of knowledge is being rewritten. The question is: Who holds the pen?

In a world increasingly shaped by machines, the answer must still be: humans, who are curious, creative, and courageous enough to ask, lead, and imagine what no machine ever could. ■



DR KENNETH BENOIT

is Professor of Computational Social Science and Dean, School of Social Sciences at Singapore Management University. This is an adaptation of the inaugural lecture “*AI and the Future of Social Science*” that he delivered on April 22, 2025

For a list of references to this article, please visit <https://smu.sg/SMUAMIJuly2025> or scan the QR code below.

