

The Epistemic Mirage



Why the AI Revolution in Academia Hasn't Happened Yet



EPISTEMICA

Epistemic Intelligence Engine

www.epistematica.co.za

10 February 2026



1. The Phantom Revolution

The noise is deafening, but the needles haven't moved

The academic sector is currently suffocating under a crushing sense of FOMO (Fear of Missing Out). The narrative dictates that we are in the midst of an AI revolution. Every week brings a new champion—ChatGPT, Claude, Gemini, DeepSeek, etc—each promising to be the ultimate tool for research, or at least literature reviews, data analysis, or syllabus generation.

To many academics, the pressure is immense. You are told that if you are not actively deploying these tools, you are obsolete. Yet, when you look past the LinkedIn posts and the enthusiastic workshop demonstrations, a stark reality emerges: nothing has fundamentally changed. The AI optimists and the AI pessimists are functionally stuck at the exact same level of output. No one is suddenly publishing twice as many papers in top-tier journals. No one is setting radically unique assessments that look any different from what was already sitting on Blackboard or Moodle five years ago.

We are living through a *productivity boom-doom*, an explosion of tools that has yielded minimal (if at all any) actual academic progress. In fact, the term AI Slop was coined to capture a paradoxical phenomenon. The volume of output (like preprints) surges, but the actual quality and critical rigor stagnate or decline, and this is already [causing a crisis in computer science](#), contrary to the efficiency narrative, [it doesn't reduce work, rather it intensifies work](#), and [is destroying productivity](#)! We should not be quick to conclude on AI though. A recent article in *The Economist* (12 February 2026) asks: [Is education technology mostly useless?](#) This implies a bigger challenge at play. Technological advancement does not necessarily lead to epistemic advancement.

2. The "Compute Language" Fallacy in Research

The tech industry sells us *Compute Language*, benchmarks where one model beats another by marginal percentages in coding efficiency or logic puzzles. As an academic, you must ask: *What does this have to do with my job?* Our institutional mandates are not tied to shipping commercial software, or unmanned drones, or self-driving vehicles. We hypothesise, we research, and we publish. We design assessments and we mark them. We sit in endless committee meetings, debate policy design, supervise theses, and write exhaustive comments on student drafts. And, if we are being brutally honest, we manage a relentless avalanche of emails, which is ironically the one area where generative AI is probably doing the most actual heavy lifting for us right now!

But beyond politely clearing our inboxes, the uncomfortable truth is that an LLM cannot raise your epistemic ceiling. If a researcher was not producing the caliber of work required for a Tier-1 journal before November 30, 2022, no chatbot or any variation of it will get them there today. These tools can accelerate the drafting process, but they cannot manufacture the rigorous methodological design or the novel theoretical framework required for high-level publication.

The FOMO is misplaced. We are collecting apps to do the exact same job we always did, falsely believing the tool itself contains the secret to better research or assessment. It does not.



3. Assessment Stagnation: The Blackboard Illusion

This stagnation is most glaring in how we assess students. We were promised that Generative AI would allow us to design deeply personalized, dynamic assessments. Instead, we are still staring at the same legacy LMS platforms, setting the same multiple-choice quizzes and the same standard essays.

The proliferation of AI tools did not spark pedagogical innovation; it sparked a defensive arms race. Academics are using AI to write standard questions faster, while students use AI to answer them faster. The core mechanism of assessment has not evolved. We have not fundamentally changed *what* we assess (the student's epistemic command) because we are too distracted by *how* the text is generated. Since the industrialization of higher education, academic assessment mainly functioned as text-generation exercise where the physical generator of the text was implicitly the originator of the thought, and therefore the owner demonstrating mastery. And for over a century, this system worked. But ChatGPT irrevocably flipped that script by permanently decoupling the labor of generation from the evidence of cognition.

The genie is out of the bottle, and it is never going back in. We must confront a harsh historical reality: nowhere in the history of human progress has a disruptive technology been successfully shelved simply because a specific group of professionals was unhappy with how it complicated their workflows or threatened their established practices. One only needs to look to the Luddites.

In 1811, skilled English textile artisans began systematically destroying the newly introduced mechanized looms and wide stocking-frames. They did not do this out of an ignorant hatred for progress; they did it because the new machines fundamentally commodified their craft, allowed manufacturers to employ unskilled workers at lower wages, enabled the production of inferior "cut-up" goods that damaged their trade's reputation, and threatened their livelihoods. Their grievances regarding quality, labor standards, and fair wages were entirely valid, yet their rebellion was brutally crushed, and mechanized textile production became the permanent baseline of the global economy.

Higher education is currently staging its own intellectual Luddite rebellion. Through endless policy moratoriums, the clinging to legacy platforms, and the futile deployment of AI-detection software, academics are trying to smash the mechanized loom of the 21st century. But attempting to force-fit a cognitive efficiency engine into our traditional grading rubrics is a historically illiterate strategy. The mechanism of textual production has permanently changed. Therefore, we must stop treating assessments as mere text-generation exercises, the presence of generative AI changes absolutely nothing about the quality of university education, it merely exposes the structural fragility that AI tools have introduced to how we have always measured human mastery.

4. Policy Purgatory: "Detect but Don't Detect"

The clearest evidence of this institutional paralysis is the endless cycle of AI policy drafting. For three years, university task forces have been locked in meetings, producing iterative drafts that never seem to finalise into definite specifics.

Following the spectacular failure and proven biases of AI detection tools, universities lost their traditional enforcement mechanisms. As a result, policies have devolved into a hollow



patchwork of buzzwords: *ethical use*, *responsible AI*, and *AI literacy*. We have institutionalized a contradictory, paralyzing mandate of *detect, but don't detect*. [So where's the line? It's an absurd line](#), regardless of where or how long you draw it!

Consider the dilemma of the frontline academic who is forced to execute these top-down directives. Institutional leaders explicitly caution against using AI detection software, yet simultaneously demand that lecturers police "responsible" usage. Some policy drafts casually suggest auditing a student's AI chat history as proof of cognitive effort. This is an administrative absurdity. A student's iterative prompt chain is often pages longer than the actual 2,000-word essay. Is the academic (already drowning in administrative KPIs) supposed to spend three hours reading a chat log just to verify a single assignment? The lecturer is left staring at the rubric, entirely unsupported, wondering: *What exactly are you asking me to do? Am I detecting this or not?*

Academic publishers offer no better clarity, relying on an equally vague mandate of *use, but don't use*. Take the daily reality of a non-native English-speaking researcher. If they design a rigorous methodology, write the core manuscript by hand, and then use an LLM purely to translate and polish the text into pristine academic English, did they "use" AI? Does that constitute "responsible usage" or an ethical violation? The guidelines provide no objective answers, leaving the academic in a state of perpetual limbo.

The burden of interpreting these contradictions has been entirely offloaded onto the individual researcher and lecturer. These policies are symptoms of institutional exhaustion. They fail because they attempt to govern the tool—a rapidly evolving technology—rather than governing the epistemic integrity of the academic process itself. But if there is any consolation, academia is not alone in the doldrums of AI governance: a [KPMG partner was recently fined over using AI to pass AI test!](#)

5. The Way Forward: Frameworks Over Apps

The AI productivity boom-doom ends when we realize that the solution to an epistemic crisis is not another generative app. The solution is a return to rigorous frameworks. Academics must step off the endless treadmill of tool adoption. It does not matter if DeepSeek is marginally better at phrasing than ChatGPT. We must pick one foundational model and return to the core of the profession.

However, returning to the core does not mean ditching generative AI to blindly retreat into structurally fractured traditional methods. We already know the legacy assessment model cannot survive this new environment on its own. Rather, it means building a robust verification layer over those traditional foundations to reinforce them. As I explore in Chapter 5¹ of my forthcoming book, *Applied AI Epistemics*², the central governance problem of an intelligence-abundant world is the allocation of human agency. When cognitive capacity becomes cheap and scalable, the value shifts entirely to the human who authorizes, verifies, and defends the logic of the machine's output.

We do not need tools that do the thinking for us; we need frameworks that force us to prove we have thought at all. This is the structural mandate behind the Cognitive Growth Index

¹ Chapter 5: Cognitive Growth and the Allocation of Human Agency in a Generative Intelligence Economy

² *Applied AI Epistemics: Rethinking Cognition, Behavior and Assessment in the Age of Generative Intelligence* (publisher: Springer)



(CGI) a metric for cognitive friction, measuring the human's command over the AI-generated artifact.

When an institution adopts this mindset, its technology stack naturally evolves. We move away from generic chatbots and toward specialized epistemic engines. Platforms designed with this verification layer—such as [Epistemica's](#) *Cognitio*, *Cognitio Strat*, and *Vula*—do not exist to compete in the benchmark wars of compute language. They exist to enforce an epistemic baseline. They integrate with the academic workflow to ensure that while the AI may assist in the generation of the text, the user cannot bypass the necessary work of structural synthesis.

Ultimately, we must shift our energy away from finding the perfect AI workflow, and back toward designing assessments that require true cognitive friction, producing research built on robust frameworks, and demanding policies that govern human thought rather than machine output.

About the Author

Prof. Prince Sarpong is an Associate Professor of Finance at the School of Financial Planning Law, University of the Free State, and a post-disciplinary systems architect whose work spans corporate finance, financial psychology, AI epistemics, and institutional behaviour. He has developed several frameworks, including Brittle Financial Health, Antifragile Financial Therapy, Institutional Financial Therapy, and the Cognitive Growth Index (CGI)—a generative AI model for mapping intellectual recursion and adaptive cognition.

His broader intellectual ecosystem, Sarpong Strategia, anchors a set of transdisciplinary ventures: the Centre for Financial Psychology and Systems (CFPS) in Bloemfontein, South Africa; Financial Therapy for Men Inc. and Fortis Strategia in Sacramento, California; and Applied AI Epistemics, a research lab focused on epistemic research and reform in the age of artificial intelligence.

Prof. Sarpong currently serves on the advisory board of Epistemica, an epistemic engineering firm, and AI 2030, a global nonprofit for responsible AI. His work has been featured in Yahoo Finance, Fast Company, San Fransico Chronicle, The Australian Financial Review, RealTid (Sweden), The Psychologist (Ukraine), and MT/Sprout (Netherlands).

Prof. Sarpong's current work confronts the limits of traditional disciplinary boundaries by designing systems that metabolise pressure into structure across money, identity, capital, and corporate leadership.

