


Cognitive Depletion and Epistemic Arrogance



The Other Side of The AI in Higher Education Debate We Do Not Hear

Prince Sarpong, PhD



EPISTEMICA

Epistemic Intelligence Engine

www.epistemica.co.za

10 February 2026



The discourse surrounding artificial intelligence in higher education has settled into a comfortable, yet entirely predictable, binary. On one side, proponents herald a new era of efficiency; on the other, skeptics fiercely defend academic integrity, guarding against plagiarism and the erosion of the “human touch”.

Both sides raise valid points, but if we are to have a truly holistic debate about the future of academia, we must put all the cards on the table. Currently, a crucial card is missing from the conversation: the biological limit of the human evaluator.

When we frame the hesitation to adopt AI strictly as a noble defense of academic rigor, we inadvertently mask a profound structural inefficiency in academic labor. We do not need to pit *AI Integrationists* against *AI Precautionists*. Instead, we need to honestly examine what the traditional, purely manual system of evaluation costs us in terms of cognitive depletion and the inherent limits of human expertise.

The Biology of Mental Exhaustion

Human cognitive bandwidth is a strictly finite resource. When an academic sits down to manually parse 100 postgraduate exam scripts, tracking cascading numerical errors, verifying theoretical integrations, and writing localized feedback, they are engaging in an act of severe cognitive endurance.

Educational psychology has long understood this through Cognitive Load Theory. Working memory can only hold and manipulate a limited number of novel elements simultaneously. High-level grading carries an immense intrinsic load. When the brain is forced to process this load repeatedly, it inevitably hits a threshold of exhaustion. The resistance to delegating this task often comes from a place of deep professional pride, but it could also be much like a professor who insists on sweeping and mopping his own office because he prefers his own standard of cleanliness.

But practically, this insistence leads to decision fatigue. By the time an examiner reaches the 40th script in a stack, their brain shifts from analytical reasoning to a reliance on heuristics—mental shortcuts that require less energy but yield shallower judgments. Therefore, the depth, fairness, and rigor of the feedback a student receives on script number 85 is mathematically guaranteed to be inferior to the feedback provided on script number five. This is not a moral failing or a lack of dedication; it is a biological reality.

The Epistemic Bottleneck in Advanced Evaluation

This brings us to a second uncomfortable truth that a holistic debate must address: the systemic pressure that leads to epistemic arrogance. In philosophical terms, epistemic arrogance occurs when a knower’s confidence becomes disproportionate to their actual ability to know, often manifesting as *credentialed arrogance*.

The traditional academic system incentivizes a facade of absolute mastery. Consider the role of the external examiner on a PhD dissertation. If a finance professor is tasked with examining a thesis that utilizes a highly specific, niche econometric technique, for example, a complex machine-learning approach to selection bias, they face a severe bottleneck. Lacking the dozens of hours required to relearn the underlying mathematics, the examiner is often forced

to skim the methodology, relying on surface-level heuristics and their own credentialed authority to pass judgment.

We must weigh this reality against the current empirical landscape of machine intelligence. Frontier AI models have systematically dismantled the hardest professional benchmarks in the world. They pass the Uniform Medical Licensing Examination (USMLE), score in the 90th percentile on the Uniform Bar Exam, and, highly relevant to the finance sector, can now pass all three levels of the Chartered Financial Analyst (CFA) exam, including the notoriously difficult Level III essay questions that require strategic wealth planning.

If a machine can process and apply the codified knowledge required to pass a CFA Level III exam in seconds, we must ask ourselves a difficult question: is human intellect truly best spent manually checking basic tax calculations or theoretical definitions in a student's script?

The Cognitive Exoskeleton and Reallocation of Intellect

A complete debate requires us to move past the fear of 'handing over' education to machines. The integration of AI into evaluation is not an abdication of the professor's role; it is a strategic *reallocation* of cognitive bandwidth.

Large language models excel at the mechanical application of codified knowledge. They are flawless at identifying structural deviations, mapping facts against a rubric, and calculating baseline scores. But machines struggle with the uncoded: the nuanced context of a client's behavioral biases, the ethical gray areas of a case study, and the overarching philosophical intent behind a student's argument.

By treating AI as a *cognitive exoskeleton*, the professor allows the machine to absorb the crushing weight of data processing and algorithmic verification. This preserves the human's mental bandwidth for what the machine cannot do. The academic elevates their role from a line-by-line auditor to a high-level editor and adjudicator. They review the AI's generated critiques, spot-check complex integrations, and refine the feedback based on their own professional judgment.

Ultimately, utilizing intelligent tools for marking and evaluation is not necessarily about saving time. It is about radically improving quality and consistency. The academic who integrates these systems will maintain a level of analytical rigor across their 100th script, or their 10th PhD examination, that a purely manual worker cannot physically sustain. If we want to protect the "human touch" in education, we must first protect the human mind's capacity to deliver it.

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 Francisco 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.