■ Research Article

Navigating The Destructive Impact of Generative AI on Academia and Education

Johan EDDEBO	Uppsala University
Dennis RICHES	Seijo University

FUTUREd, Volume 1, Issue 1 (2026)

Pages: 11 - 22

ISSN: XXXX-XXXX (print)
ISSN 2760-8271 (Online)

Keywords:

artificial intelligence, technology and education, AI and science

Corresponding author(s):

Johan Eddebo
johan.eddebo@crs.uu.se

Dennis Riches
dennis.riches@proton.me

Abstract: This paper investigates how generative AI's ability to produce knowledge-like outputs reshapes the ways knowledge, evidence, and truth claims are created and evaluated and what challenges this creates for education and research. Addressing these issues is critical in that the basic function of academia depends on reliable methods for producing and evaluating knowledge, and generative AI may change how evidence is created, interpreted, and trusted within these institutions and thus for the broader society.

We also note a tendency for the centralization of epistemic authority due to these technologies, which in tandem with increasing political and ideological concerns over maintaining orthodoxies of belief portend unique challenges to the very idea of the university. The emerging modalities of AI can be seen as both a reflection and continuation of a reductive instrumentalization of knowledge and a centralized control of opinion formation, and we argue that the entrenchment of contemporary generative AI in our institutions of knowledge-production is likely to reproduce such structural objectives or intentionalities.

We finally reflect on some possible approaches towards navigating the situation with special regard to educational reform and innovation.

1. INTRODUCTION

Current debates around the impact of generative AI on the academic environment are extensive, and cover a broad set of issues pertaining to educational outcomes and institutional disruption. There is, however, a lack of nuance in these discussions, as well as clear and coherent definitions of the relevant forms of AI which also hampers a meaningful discussion of their impact (see e.g. Bearman, M., Ryan, J., & Ajjawi, R. 2023).

This article intends to address this issue by taking a step back and approaching the emerging modalities of AI within an academic, educational and scientific context by focusing on how the technologies in question impact knowledge-production. What we are emphasizing is the sense in which AI, and specifically its generative modality, produces truth-claims and evidence and is involved in generating meaningful cultural content (Alvarado, R. (2023). This narrower approach is intended to be helpful in underscoring

the types of causal impact of the technology that is especially relevant for science, academia and higher education as knowledge-producing institutions.

The approach of the paper is to start from this focus on AI, and then to outline a set of synergizing challenges for educational and research-related institutions that are likely to emerge from the entrenchment of such epistemic functionalities pertaining to AI. From that overall assessment, the paper also explores strategies towards addressing these challenges through intervening in the educational and methodological foundations of academia and science.

2. THE SENSITIVITY OF SCIENTIFIC AND ACADEMIC ECOSYSTEMS

Science and academia operate as a system for knowledge production and theory formation anchored in empirical reality, and are guided by a certain set of established methodological principles (Cetina 1999, Merton 1973). While this is a vague and broad description, it implies a depth and complexity of these institutions that cannot readily be instrumentalized or reduced to a simple ruleset. If we take a step back and consider the characteristics of science as a set of traditions of knowledge, their complex epistemic substructure is immediately apparent. First of all, science involves crucial and fundamental theories about what the world is like and how it works, such as the principles of causation and the basic structure of logic. It also operates on a foundation of crucial values and objectives that hold the enterprise of science together and which are often unstated, and there is a wealth of hands-on knowledge that has to be learned in the field, and which cannot easily be reduced to a simple set of rules to follow.

And it's especially these tacit, silent or hidden skills that are most sensitive to disruptions in the basic methods and theories that a tradition of knowledge is built on, partly because they're generally preserved through custom, habit or person-to-person instruction, so if the tradition is significantly modified, these skills are often lost (Cf. Polanyi 1966).

One clear example of such displacement can be seen in the disappearance of riveting in construction and engineering, in spite of the current alternatives actually being inferior (cf. Braverman 1974).

Riveting is where we fuse together a structure, generally one of metal, with a kind of peg that in the process gets permanently deformed and cannot be removed – you push a softer metal peg through a perforation, and then you flatten the peg on the other side of the metal sheet, and this creates a solid, fixed connection. The rivet has to be destroyed to be removed.

Structural steel connections used to be either welded or riveted, with riveting sometimes being the preferable option even above welding, depending on the kinds of forces that are likely to impact the structure. However, riveting these connections normally takes four skilled workers with different specialties, and this increases the labor costs and time needed to train the workers (Giddings 1950). This has resulted in riveting nowadays mostly being replaced by bolting, which while much cheaper is not vibration resistant and has a number of additional disadvantages. In specialized sub-fields such as aircraft construction, riveting is still indispensible, since here you need full vibration resistance.

But this lowering of labor costs has also led to the requisite skillset for riveting being eroded, so in such areas where it's an absolute necessity, there's now a shortage of this skillset, very few people knows how to do it properly anymore, and this arguably also leads a lower average level of competence across the board, because if almost nobody knows how to rivet properly, you will have fewer teachers able to convey the skillset to other people.

Science as a whole is likely to be impacted by AI technology in similar ways. There are many kinds of tacit or hidden skills and knowledge that are absolutely crucial for science to function properly, yet which are taken for granted and in the short term may be displaced by cheaper and less time-consuming solutions that may not be perfect, but like bolting, as simulations of the original, seem

generally adequate.

We have already mentioned values and fundamental philosophical principles like causation and the framework of logic that are tacit, or unspoken, in this way, and the comprehension of which cannot even in principle be reproduced by the probabilistic weighting of a large language model. The simulation of such comprehension can only ever be a limited reflection of the outward appearance or trappings of genuine understanding.

But what's perhaps more obvious are the equally often unspoken rules of observation and interpretation of data that play a crucial role in how facts are described within different academic and scientific disciplines.

> Well, when I did not see what I had expected to see, I complained to my biology teacher. He calmed me down saying that this difficulty was experienced by everyone and that I had to learn to see (Feyerabend 1991 p. 106).
>
> Moreover, I have the curious feeling that the phenomenon was there all the time, but that I was too dense to notice it. Like after-images which intermingle with our vision at all times but which must be brought to our attention by special methods (Ibid p. 101).

And such rules or discreet methods of observation also tend to be deeply intertwined with a hands-on experience and skills that cannot even in principle be strictly formalized. That cannot even be written down as rules without undermining their flexibility and their adaptability and openness to new and unforeseen situations. So if we codify approximations or simulations of them, they become inflexible and can no longer be adapted to changing circumstances. It's like telling a boxer that there are only three types of punches and that he can never use anything that doesn't strictly fall into one of these categories – but that's too narrow for the dynamic and complex reality of the sport. There's much more to boxing well than what you can fit within a simple set of rules.

There's likewise a crucial domain of science that relies on this tacit and hidden experience and form of knowledge and which can't be codified once and for all, but which is rather a certain kind of sensibility that must be learned from within the field. It's similar to how you learn to improvise in music. If you put down hard and stringent rules for this, it's no longer improvisation and cannot provide the associated goods.

> Where serious music satisfies its own concept, every detail gets its concrete meaning from the total course, and this totality in turn receives it from the living interrelation of details that oppose and continue one another, pass into each other, and recur. Where the form is dictated from outside, on the other hand, one will, as Wagner put it, "hear the harness rattle." (Adorno 1962 p. 28)

And we can especially see the need for such intuitive and tentative approaches in critical situations when new and important discoveries must be integrated in spite of challenging established theoretical frameworks. This happened, for instance, when Max Planck's anomalous findings on black body radiation, explained through is model of the action quantum, disrupted even basic principles for how facts in physics were assembled and theories confirmed (Planck 1900). In such circumstances, there are no established rules to go by, so you have to fall back on this kind of skilled and intuitive improvisation that cannot be translated into anything else, and simply play it by ear.

If we lose the complex and often hidden skills of human cognition that not only science but really all human traditions of knowledge depend upon, such adaptations, crucial for meaningful progress, may be rendered impossible.

From much work in the philosophy of technology, we know there's a significant impact of the measuring apparatus or of any devices or techniques that enter into the methodological framework that particularly disrupts or displaces this non-codified, tacit or soft knowledge, including implicit values, objectives and metaphysical principles (Cf. Latour & Woolgar

1979). A prominent example is described by Jonathan Crary in his *Techniques of the Observer* (1990), showing that the experimental setup of the Camera Obscura basically constitutes modern Cartesian philosophy of mind and the seminal Kantian separation between the observer and the thing-in-itself, ideas which were originally mapped on this influential and quite tangible distinction between the darkened chamber and the precise yet synthetic representations of the world outside.

This soft knowledge and tacit principles can be defined as an epistemic substructure beneath such implementation and reproduction of information that takes place within traditions of knowledge like modern science. An emphasis on the importance of these modes of knowledge is found already with Aristotle (2009), but the modern debate around their intersection with technology and automation and how this impacts knowledge production probably originates with Merleau-Ponty's *Phenomenology of Perception* (1945) and is first clearly outlined in relation to digital technologies in Dreyfus (1972).

A key sensitivity of scientific and academic knowledge production lies precisely in the crucial role of these softer modalities of knowledge, that are also uniquely difficult to adequately simulate and reduce to instrumental functions. Systems or methodologies that exclusively operate from such simulation or instrumentalization, however useful in other respects, must then be considered inadequate substitutes for these tacit modalities of knowledge.

For this reason, the introduction of generative AI in the frameworks of scientific and academic knowledge-production may produce severe disruptions in these traditions.

3. THE CHARACTERISTICS OF THE DISRUPTIONS

Generative AI simulates and instrumentalizes knowledge, and cannot adequately represent the rooted and embodied comprehension that characterizes tacit and soft knowledge (Merleau-Ponty 1945). This implies a *decontextualization* of knowledge, an uprooting of knowledge and information from their original embeddedness.

A progressive decontextualization can be seen in from the earlier history of digital information technology. We first observe a basic technological centralization that makes users and developers dependent on a few proprietary platforms and standard frameworks (van Dijck, Poell & de Waal 2018). This situation has been much exacerbated over the last 20 years, with de facto monopolies for access to content and services emerged, which impaired the autonomy of scientific disciplines for managing their own data and methodological resources, but also impaired market access to locally created content and applications – the types of software and output that reflected the needs and objectives of local communities, whether among the public or in traditions of knowledge like science (Couldry & Meijas 2019).

So for instance, when IRC (Internet Relay Chat) was developed back in 1988 by Jarkko Oikarinen in a collaboration between the universities of Oulo and Helsinki, the application was designed to cater to the specific needs of communication, privacy and interaction that were obvious in this particular context (Oikarinen & Reed 1993). But as IRC was pushed out by commercial platforms such as Facebook, the objectives of advertising and customer intelligence or surveillance replaced these locally expressed goals and needs, and soon, there was not really any practical alternative to the major proprietary solutions.

The same thing is happening with AI in relation to science on a much larger scale, and in a way that not only supersedes locally developed objectives and approaches to knowledge. It also disrupts the contextually critical and unique skillsets for the different academic and scientific disciplines with a one-size-fits-all-solution reflecting the objectives of the platforms (Srnicek 2017).

3.1. Generative AI and information-curating algorithms

Generative AI has recently emerged as a transformative set of toolkits that potentially influence the entirety of human culture. It simulates more or less every kind of human

cognitive output that can be objectively described. The intersection between generative AI and the broader algorithmic curation of information further implies that the knowledge repositories upon which complex human activity builds will be impacted by the inherent tendencies of the AI tools.

The decay of complex skills

We know that AI use undercuts basic human cognitive abilities. This can be related to the clear evidence of a cognitive impact even with the comparatively primitive digital technologies of more than 20 years ago. One prominent example was of how memory and effective recalling of information was impaired by Google use – one can describe it as how the brain doesn't prioritize actually retaining the information if we know that it's going to be easily and immediately available through the search engine (Spitzer 2012). In the same sense, in 20 years from now, few people will likely learn a foreign language in-depth assuming that translation devices become more widespread and readily available.

But these kinds of developments can with AI in principle go much further than just the erosion of basic cognitive abilities. With widespread AI use, there is a thousand "invisible" skills that potentially get lost along the way. There is something you learn when you are forced to do these boring, repetitive tasks, such as sifting through the pages of a dusty old dictionary and memorizing the words, which gets lost when you have a search engine constantly at your fingertips. There is something intangible and uniquely valuable that is gained just through suffering the boredom and the tediousness that is normally associated with acquiring the basic foundations of our higher skills, whether it's about learning grammar, figuring out through trial and error how to interact with the opposite sex, or developing the muscle-memory connection of practicing scales on a musical instrument. We gain some kind of pain tolerance, some kind of cognitive endurance or robustness that's necessary for really internalizing cognitive abilities and making them reflexive or "second nature". If and when such a skill decay meaningfully impacts structurally important expert populations, we are very likely to see aggregate effects on the institutions that they operate. (See also MacNamara et al 2024).

This situation becomes particularly significant when we consider how lower-order "invisible" skills are crucial for the development of true mastery in most disciplines and traditions of knowledge (Anderson 1982).

The risks of entrenching generative AI in research methodologies

Insofar as technology reproduces ideology, an information architecture that structures research methodologies is likely to influence the values and objectives of associated research output (Winner 1980).

There is a long tradition of critique uncovering the ideological impact of technologies, from Mumford to Foucault, and contemporary AI is in no way immune to similar problems. *Artificial Whiteness* by Katz (2020) or *Algorithms of Oppression* (Noble 2018) are prominent investigations in more recent literature in terms of the consequences of the intersection between digital and social systems. In terms of knowledge-producing traditions and institutions, we should expect even more significant downstream effects if and when the basic conditions for producing and evaluating evidence are impacted. Contemporary AI can be considered a uniquely "epistemic technology" in that its purpose is to produce meaningful cultural content and truth-claims (Alvarado 2023).

Crucially, self-learning algorithms are also adaptive in a sense that no other technologies have arguably ever approached, and are thus capable of "strategically" adapting towards more effectively reproducing certain inherent values, ideologies and objectives. This is likely to entrench the ideological effects of contemporary AI much more robustly in the social systems that come to integrate them.

Even Plato, arguably the father of written rational reflection in the West, warns that something intangible is being lost when writing predominates our thought and communication, not only in terms of people's capacity to remember things by themselves, but also in terms of a certain flexibility of critical thought that he argued can only be

properly practiced in a real-life dialogue. Moreover – reading and writing augments and changes a relatively narrow aspect of the human experience. The textual medium is restricted to the sphere of abstract thought and communication, and does not replace or interfere with most normal aspects of human behavior or activity. The digital technologies, however, have a much more intrusive character. They penetrate into almost every aspect of human culture, industry and activity. They even enter into the structure and formation of intimate human relationships in various ways.

This implies likely aggregate effects of the ideology inherent to AI technology emerging in the output of any and all traditions of knowledge that come to depend on it, including the institutions of modern science. (Hirvonen et al 2023). One clear path to such aggregate effects is the issue of the massive increase of AI-generated material in the datasets on which scientific and academic research rest. First of all, we have the publish-or-perish framework of research where output often trumps quality – and since AI provides numerous ways to cut corners in this process there is a significant incentive for researchers at all levels to inflate their output using generative AI. At this moment, there are also dozens of tailored AI tools specifically marketed to researchers, scholars and students for producing academic and scholarly articles- Research indicates that tens of thousands of peer-reviewed articles produced with generative AI were published only in 2023 at a quite early stage of the technology (Brainard 2024), and there are stylistic signs that theoretical output has been influenced by generative AI in terms of its content (Gray 2024). This overall process inevitably causes an inflation of low-grade output that crowds out high-quality, time-consuming research since it can easily be made to look just good enough and just passes the criteria for inclusion while not necessarily providing any meaningful contributions to the actual field in question.

And the researchers who refuse to cheat in this way are of course at a disadvantage – their output will be more limited in relation to that of researchers who through specialized genAI-tools can massively increase quantity of output. And while these may be sub-par and derivative, the bibliometric framework is mainly about numbers, so the researchers who avoid using these technologies will be disadvantaged in the competition for tenure or research positions, and are thus more likely to leave academia altogether.

Of course, such an inflation of low-grade output will also add increased pressure to the peer-review process and may exacerbate weaknesses of the current model – which is based in voluntary contributions from reviewers and editors. With a significant increase in material produced, the capacity of the current peer-review "workforce", which is made up of researchers who are willing to give some of their valuable time to do painstaking work with little to no compensation, this capacity is likely to be severely challenged (Ebadi et al. 2025).

The path of least resistance, due to the increased workload of editors and scholars providing peer review, is towards an increasing utilization of generative AI, implying that LLMs are already becoming entrenched in academic knowledge production both upstream and downstream.

In relation to an increasing proportion of AI-generated research output, the well-known issue of "hallucinations" becomes potentially much more destructive. A related issue is the problem of misleading or spurious correlations emerging in analyses based in automated pattern recognition, where the system picks up on inconsequential correlations and masks them as ostensibly real signals yet that are seemingly to weak to be picked up by humans (Geirhos et al. 2020). This is due to LLMs being optimized for statistical associations and incapable of comprehending actual causal relationships – another example of how structural intentionalities become reproduced at scale.

The entrenchment of such correlations in datasets, including of hallucinations proper, is also prone to cause a variety of problematic feedback loops apart from the full-blown model collapse (where output just becomes noise). Less catastrophic destabilizations errors implanted deep within knowledge repositories, undermining the structural foundations of knowledge production. One can

envision how this could work out on a longer timescale where historical knowledge gets filtered through the weighted averages of an LLM- First it gets trained through authentic historical knowledge, which over time becomes displaced by genAI output that subtly skews the probability weighting of the discursive nodes in the model. Over time, the simulated knowledge output may subtly and slowly detach from authentic historical data, reproducing falsehood due to minor destabilizing hallucinations akin to an entropic process.

4. REFLECTIONS ON THE EDUCATIONAL IMPACT

The effects of AI on education are perhaps self-evident given this overall context. They are obvious and they have been thoroughly discussed in casual discourse, mass media, social media and academic discourse since the emergence of LLMs several years ago (Întorsureanu et al. 2025). Parents and teachers see the effects with dismay while some of them use AI reluctantly or enthusiastically, admitting that there are serious concerns but insisting that it will have too many advantages and be too unstoppable to be dismissed in a modern "Luddite" reaction.

The obvious distinction that is lost on many observers is the distinction between AI use by those who have already acquired particular skills and AI use by young learners to solve problems for which they lack the requisite skills. Habitual use of AI will very likely lead to them never acquiring skills that previous generations took for granted (cf. section 3.1.1).

Most adults watching children use AI are aware of this hazard, but they feel powerless to stop it. Parents are too economically stressed to intervene constantly. The same goes for educators who are overwhelmed by the changes and constrained by needing to fill the short-term demands of their institutions. One cannot simply ban access to technology because we are all compromised by the need we have for some aspects of it. Public telephones and landline telephones have disappeared as the standard ways for people to contact each other. Educational institutions now expect students and parents to be online in order to receive important information, make payments, contact personnel, etc.

Something that has forced many people in recent years to revise their optimistic view of human nature is the realization of how many students of all ages are eager to take shortcuts with AI In spite of all the stern warnings about academic cheating and the harmful effects on their development, they take the lazy route because "everyone's doing it." The "arms race" between AI users and AI use detectors (students vs. teachers) escalates endlessly. A college student writes an essay in the foreign language that he is studying by finding an AI tool that will insert a few errors that are typical of people who speak his native language. Now, the educator must exhaust her time and patience in efforts to detect work done by AI. This distraction inevitably takes the place of thoughtful evaluation and feedback.

Furthermore, since the pandemic lockdowns forced students to go online and use digital technology, handwriting and hard copy reading and writing have gone by the wayside, and there is no incentive to bring them back. There is no valuation of what could be called "slow learning" (not to be confused semantically with the label "slow learner" assigned to some students). What we mean by "slow learning" is doing things slowly with pencil and paper, learning math by "showing the work," or learning photography with the manual settings. The list of such skills could be very long.

Ironically, it is the technology billionaires who boast about not letting their children have access to digital devices. They knew long ago, before the advent of LLMs, that technological dependence would impede the acquisition of skills that adults take for granted. Their children receive the new luxury of learning cursive writing, reading books on paper, and having tutors and teachers who can make sure they stay off of digital devices and social media. This is, metaphorically, the difference between a diet of processed food and a diet of freshly cooked organic food. The children of the rich can go slowly because they need not be too concerned about the competition to enter higher education. They know already that they are growing up with great advantages and will have all

the educational opportunities they wish.

One may see, however, at the time of this writing, a certain turning of the tide in the enthusiasm for LLMs and AI of the 2020-2025 period. People are starting to understand that the miraculous textual output of these systems is not evidence of an underlying intelligence or an emergent sentient being. They are glorified search engines and not actually intelligent in any meaningful sense of the word (cf. Eddebo 2025). This awareness is emerging also as it becomes apparent that the massive investment in AI data centers is a financial bubble that risks collapse as it hits the wall of real-world constraints—limited supplies of computer components, energy, water, land, labor, hospitable communities, and so on.

Parents and teachers are noting the AI solution hitting a wall. The young are not doing anything constructive. They are often depressed and demotivated, and it's difficult for parents to think of a reason why they should be motivated to go to school to just copy and paste some answers taken from ChatGPT. As a possible solution, we will attempt a drastic rejection of AI and technological devices. In our own teaching, we have decided to make some changes next semester. We will cut students' access to technology during class time. Their smartphones will go into a box during the lesson. They will have no computers or internet access. All assignments will be written by hand during class time or spoken aloud. They can have all the digital access they want outside of class, but the work that counts for their grades will be done in the classroom. We expect resistance during a difficult period of withdrawal, but this is what has to happen. This method, something inspired by the "elite education" that the children of billionaires receive, seems to be the only way out of the technological trap that has been laid over the past decade.

The most promising approaches to dealing with the detrimental impact of generative AI in the educational context may very well be ones that can effectively leverage the inherent capabilities and unique tendencies of these technologies towards fostering entirely new modalities of critical thinking and autonomous reasoning. Just as literacy undermined traditional forms of thought, experience and knowledge while opening up spaces of discourse and reasoning entirely impossible without it, there's at least the prima facie possibility that today's AI tools, if perhaps only jury-rigged, open-source, hacked or redesigned, can enable new forms of thought and hitherto unfathomed complexes of meaning – much like how the counterpoint in music or the rediscovery of perspective in 14th century art enabled previously unthinkable aesthetic meaning.

5. CONCLUSIONS: A NEW EPISTEMIC CRISIS IN SCIENCE AND ACADEMIA: ROOTS AND REMEDIES

Contemporary AI is arguably in and of itself a major epistemic paradigm shift for the production of knowledge. It provides a complex set of challenges much like earlier periods of crisis in science, except that such crises have generally been internal transformations driven by theoretical developments and discoveries of seemingly inexplicable phenomena, rather than imposed from the outside through an autonomous machinery for simulating knowledge and truth-claims.

While it's certainly possible that we may yet see adaptive innovations that leverage these new technologies in surprising ways, bringing about new forms of genuine understanding or even modes of rational reflection, their disruptive potential is quite daunting. They key aspects of this potential arguably lie in the tendency towards wholesale cognitive outsourcing and the potential disconnect between AI output and actual reality, the quintessential "hyperreal" character of AI as purely a representation of representations. Crucially, today's AI expresses separation and simulation rather than a genuine interaction with reality in much the same way that the early-modern interpretation of the Camera Obscura renewed and seemed to confirm the Cartesian estrangement of the self from the world.

If the contemporary situation is rather to be characterized as an external disruption of an ecosystem of traditions of knowledge than something akin to a fruitful clash of opposing paradigms or theoretical rivalries within such

traditions or even a mutually supportive ecosystem (Lakatos 1978, p. 176), the outlook becomes a more pessimistic one. An imposition of a foreign system of knowledge-production upon the domain of another isn't necessarily going to be a fruitful prospect for the recipient, which the history of human cultural interaction clearly evinces. Moreover, if the intruding tradition is essentially alien in terms of its function and objectives in relation to the recipient, a fruitful cross-pollination may not even be possible.

Since today's predominant forms of AI does not actually produce knowledge or genuinely intentional truth-claims, it's really only a mimic in relation to human agents participating in actual traditions of knowledge. In this sense, AI can perhaps be likened to the trans fats that mimic dietary lipids to the extent that they can integrate into cellular structures, but due to their foreign structure cannot be processed in the human metabolism and cause persistent inflammation and oxidative stress (Oteng & Kersten 2000).

Pushing this analogy further, a critically informed and experimental education may be the most straightforward approach to both find ways towards a safe and controlled use of AI tools, and for adapting societies to the impact of these emerging technologies. In a sense, the reductionist simulation of knowledge that AI reproduces has many deep connections to the Fordist tendencies of the authoritarian schooling that Althusser identified as a repressive ideological state apparatus (1970). On this interpretation, the integration of generative AI in modern schooling and education is simply the technological refinement and extension of the institution's underlying ideological function and a perfection of its inherent tendencies.

In the final analysis, education is neither more nor less than the handing over of culture to the next generation, the reproduction of a society's worldview, way of life, and privileged forms of knowledge. On this outlook, the character of contemporary AI, and the unique reception it receives, is ultimately a product of education in this more general sense. Of industrial modernity's reductionist, instrumentalist and reifying system of knowledge-production that in a sense created the contemporary LLM in its own image.

But just like Francisco Ferrer and John Taylor Gatto in their day saw in education not only a potentially repressive set of limitations on the human condition, but also a key for unlocking hitherto unimagined human potential and for reshaping the very foundations of our social institutions, we may look to education for possible solutions to the underlying imbalances of which the darker side of today's AI are really only a symptom.

The root of the instrumentalizing mimic of reason reproduced by generative AI namely lies in an ideological trajectory of reductionism and materialism that has been part and parcel of (Western) modernity at least since the Enlightenment. In this sense, the artificial agent of the large language model functions as modernity's quintessential metaphor of the human being on an ontological backdrop of materialism, where truth, meaning, values and intentions are merely "helpful fictions" and really reduce to weighted probability calculations optimizing the survival of Dawkins' "selfish genes", mirroring the institutional logic of industrial civilization.

And precisely because education as a reproduction of culture ultimately depends on genuine relationality, it shouldn't be surprising that the unique contradictions of today's AI modalities are so strongly expressed in the context of schooling and academia. An organic transmission of a way of life is namely impossible without a rootedness and embeddedness, without the real relation that the LLMs reductive and simulative representation of knowing and being ultimately eschews.

But just like there are many other approaches to education that foster radically different and non-reductionist modes of knowing and being, there are also many different functional approaches to artificial intelligence – many of which express and reproduce ontologies that arguably foster genuine relationality and which are arguably much more compatible with actual embedded human knowing.

Rodney Brooks (1991) explores approaches to AI which are embodied and *enactive*, where

knowledge isn't reduced to representation or simulation. In this family of AI models, knowledge is situated and performative, and emerges spontaneously from a process of responses to a real-world environment and avoids emanating into detached representations.

Likewise, Maas, Natschläger & Markram (2002) exemplifies work on the category of AI-models that function as biological analogues, and thus approach other forms of genuinely emergent intelligence rather than merely simulations thereof. This modality of AI potentially expresses a direct embeddedness since it builds on analog computation expressed in physical systems (rather than a digital emulation of natural intelligence).

There are other worlds than this. There are other ways of knowing and being, and it seems that a critical task of future education is to strategically challenge the mainstream modalities of artificial intelligence which so obviously carry significant risk for the integrity of our traditions of knowledge. But this work must not stop short of a more comprehensive interrogation of the fundamental ideological principles and values of the society that has given birth to the LLM and today's generative AI.

The contradictions inherent to these systems may paradoxically provide us with a unique opportunity to rethink the deeper roots of the worldviews and ways of life that have produced many of the contradictions and challenges that are now emerging within the framework of generative AI and the algorithmic curation of information.

And while the Luddite option of outright rejecting these technologies indeed has much merit, there is nonetheless a wealth of unexplored approaches and perspectives even within the field of artificial intelligence that may provide ample support for such a task.

> "There are more things in heaven and earth, Horatio,
> Than are dreamt of in your philosophy." (*Hamlet,* Act 1, Scene 5)

ACKNOWLEDGEMENTS

This article has been prepared with the support of the Marianne and Marcus Wallenberg Foundation and is part of the output of the WASP-HS framework.

We would also like to acknowledge the kind support of the Sweden-Japan Foundation and the John Björkhem Memorial Foundation.

REFERENCES

Adorno, T. (1962). Introduction to the Sociology of Music. Seabury Press.

Althusser, L. (1970). "Ideology and Ideological State Apparatuses". La Pensée.

Alvarado, R. (2023) "AI as an Epistemic Technology". Science and Engineering Ethics 29 (5): 32. https://doi.org/10.1007/s11948-023-00451-3.

Anderson, J. R. (1982). "Acquisition of cognitive skill". Psychological Review, 89(4), 369–406. https://doi.org/10.1037/0033-295X.89.4.369

Aristotle. (2009). Nicomachean ethics. Oxford University Press.

Bearman, M., Ryan, J., & Ajjawi, R. (2023). Discourses of artificial intelligence in higher education: A critical literature review. Higher Education, 86(369–385). https://doi.org/10.1007/s10734-022-00937-2

Brainard, J. (2024). "One-fifth of computer science papers may include AI-generated content". Science. https://www.science.org/content/article/one-fifth-computer-science-papers-may-include-ai-content

Braverman, Harry. (1974). Labor and monopoly capital: The degradation of work in the twentieth century. Monthly Review Press.

Brooks, R. A. (1991). "Intelligence without representation". Artificial Intelligence, 47(1–3), 139–159. https://doi.org/10.1016/0004-3702(91)90053-M

Coan, T. G., Boussalis, C., Cook, J., & Nanko, M. O. (2021) "Computer-assisted classification of contrarian claims about climate change". Scientific Reports (Vol. 11, Issue 1). Springer Science and Business Media LLC. https://doi.org/10.1038/s41598-021-01714-4

Couldry, N., & Mejias, U. A. (2019). The costs of connection. Stanford University Press.

Crary, J. (1990) Techniques of the Observer: On Vision and Modernity in the Nineteenth Century. Cambridge: MIT Press.

Dreyfus, Hubert L. (1972). What computers still can't do. MIT Press.

Van Dijck, J., Poell, T., & de Waal, M. (2018). The platform society. Oxford University Press.

Ebadi, S., Nejadghanbar, H., Salman, A. R., & Khosravi, H. (2025). "Exploring the impact of generative AI on peer review: Insights from journal reviewers". Journal of Academic Ethics. https://doi.org/10.1007/s10805-025-09604-4

Eddebo, J. (2025). "The myth of artificial creative agency". In H. P. Olsen, H. P., Slosser, J. L., Ravn, S. A., Eddebo, J. & Rosenberg, J. H. (Eds.), Artificial intelligence, humans and the law. Routledge.

Feyerabend, P. (1991).

Three dialogues on knowledge. Blackwell.

Geirhos, R., Jacobsen, J.-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., & Wichmann, F. A. (2020). "Shortcut learning in deep neural networks". Nature Machine Intelligence, 2, 665–673. https://doi.org/10.1038/s42256-020-00257-z

Giddings, H. (1950/2016 reprint).

"Aircraft riveting". The Aeronautical Journal Vol. 54, 480. https://doi.org/10.1017/S0368393100116049

Gray, A. (2024) "ChatGPT 'contamination': estimating the prevalence of LLMs in the scholarly literature" (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2403.16887

Hirvonen, N., Jylhä, V., Lao, Y., & Larsson, S. (2023) "Artificial intelligence in the information ecosystem: Affordances for everyday information seeking". Journal of the Association for Information Science and Technology. Wiley. https://doi.org/10.1002/asi.24860

Întorsureanu, I., Oprea, S. V., Bâra, A., & Vespan, D. (2025). Generative AI in education: Perspectives through an academic lens. Electronics, 14(5), 1053. https://doi.org/10.3390/electronics14051053

Katz, Y. (2020). Artificial Whiteness. Columbia University Press.

Lakatos, I. (1978) The Methodology of Scientific Research Programmes. CUP.

Latour, B, & Woolgar, S. (1979/1986).

Laboratory life: The construction of scientific facts (2nd ed.). Princeton University Press.

Maass, W., Natschläger, T., & Markram, H. (2002). "Real-time computing without stable states: A new framework for neural computation based on perturbations". Neural Computation, 14(11), 2531–2560. https://doi.org/10.1162/089976602760407955

Macnamara, B. N., Berber, I., Çavuşoğlu, M. C., Krupinski, E. A., Nallapareddy, N., Nelson, N. E., Smith, P. J., Wilson-Delfosse, A. L., & Ray, S. (2024) "Does using artificial intelligence assistance accelerate skill decay and hinder skill development without performers' awareness?" Cognitive Research: Principles and Implications (Vol. 9, Issue 1). Springer Science and Business Media LLC. https://doi.org/10.1186/s41235-024-00572-8

Merleau-Ponty, Maurice. (1945). Phenomenology of perception. Routledge.

Merton, R. K. (1973). The sociology of science: Theoretical and empirical investigations. University of Chicago Press.

Noble, S. U. (2018). Algorithms of oppression: How search engines reinforce racism. NYU Press.

Oikarinen, Jarkko, & Reed, Darren. (1993). Internet Relay Chat Protocol (RFC 1459). Internet Engineering Task Force. https://www.rfc-editor.org/rfc/rfc1459

Oteng, A.-B., & Kersten, S. (2020). "Mechanisms of Action of trans Fatty Acids". Advances in Nutrition. 11(3), 697–708. https://doi.org/10.1093/advances/nmz125

Polanyi, M. (1966) The Tacit Dimension. The University of Chicago Press.

Portitu, T. (2024) "14 Best AI Writing Software Tools In 2024" (Top Picks). AuthorityHacker. https://www.authorityhacker.com/best-ai-writing-software/

Planck M. (1900): "Zur Theorie des Gesetzes der Energieverteilung im Normalspektrum". Verhandlungen der Deutschen physikalischen Gesellschaft Bd.2:237-245

Shakespeare. (ca. 1600). Hamlet.

Spitzer, M. (2012). Digitale Demenz: Wie wir uns und unsere Kinder um den Verstand bringen. Droemer Verlag.

Srnicek, N. (2017). Platform Capitalism. Polity Press.

Winner, L. (1980) “Do artifacts have politics?” Daedalus, 109(1), 121–136.