

## Editorial

### Epistemic decay: Generative artificial intelligence and the recombination of culture

Luke Tredinnick<sup>1</sup> and Claire Laybats<sup>2</sup>

---

<sup>1</sup>London Metropolitan University, London, UK

<sup>2</sup>Freelance Consultant, UK

Generative Artificial Intelligence is a transformational technology that augurs profound socio-cultural change on a scale that may ultimately surpass the impact of the Internet and the World Wide Web. But although offering clear benefits and opportunities, its rise has also been met with anxiety about its near and long term effects. We have previously addressed in *Business Information Review* for example the impact of generative technologies on professional roles ([Tredinnick, 2017](#)) and the ethical implications of artificial intelligence ([Laybats and Tredinnick, 2024](#)). There has also been widespread alarm at the growing use of AI in the creative industries ([Amankwah-Amoah et al., 2024](#); [Bender, 2025](#)) particularly advertising, publishing and the media. In addition, apocalyptic fears attend to the anxiety of a coming technological singularity, the point at which machines will surpass humans intelligence, initiating a snowball effect of every increasing machine capabilities and ultimately dominance ([Shanahan, 2015](#)).

Some of these perceived risks are no doubt overstated; while significant challenges and some structural transformation will accompany the wider use of generative technologies there will also be new opportunities and emerging markets. However, one potential risk has garnered less attention despite being perhaps the most immediate of them all. Generative artificial intelligence may be contributing to a gradual erosion of the epistemic foundations of our technologically and scientifically dependent culture. This possibility arises not from their apparent ability to create new knowledge, nor from the quality and reliability of the outputs that they produce, but from the ways in which generative applications have become implicated in a progressive recirculation of material culture. Successive generations of generative technologies may bring improved accuracy and fewer hallucinations, but these iterative improvements may have little or no impact of the problem of epistemic decay. This editorial explores the profound threat posed by generative artificial intelligence to our long-term understanding of what we believe we know, and what steps we can take to mitigate those risks.

### Epistemic decay

Epistemic decay describes the deterioration of the capacity for individual and societies to differentiate between what is true and what is false. It therefore describes an erosion of the foundations of our rational culture implying an abandonment of enlightenment ideals of sustained intellectual progress via the aggregation of knowledge. The threat represented by epistemic decay is fundamental in nature, an unpicking of the anchors of our scientific and technological culture by the gradual erasure or corruption of the shared knowledge based. This may ultimately have significant social and economic consequences.

Epistemic decay has frequently been associated with the effects of both social media and postmodernism. This implies that epistemic decay is a social phenomenon, arising from the ways in which truths, shared knowledge and foundational beliefs emerge in the form of a social epistemology. The relativism of postmodern narratives, combined with the post-truth nature of social media discourse arguably undermines both claims to objective truth and the basis on which public opinion is formed, creating reality filters ([Tredinnick and Laybats, 2019](#)) and disrupting Habermas' public sphere ([1974](#)). This threat is often been framed by an appeal to the traditional standards of truth and integrity that are perceived to be in decline. This means that epistemic decay can sometimes be seen as an expression of technophobic concerns about social change.

Generative artificial intelligence presents a very different challenge. The grounding of the epistemic decay for which it may be responsible is not social but systemic. It arises not from the changing ways in which people come to understand the world around them or from what constitutes truth or knowledge, nor from assumptions concerning the quality of information and knowledge disseminated through the public discourse, but from the systemic effects of iterative generative content creation itself. In particular it arises from the consequences of a progressive recombination of culture, or the tendency for generative technologies to feed off their own outputs. This is a different order of epistemic challenge, undermining not the evaluation and ratification of collective knowledge, but what we mean by collective knowledge in the first place.

That threat arises because of a fundamental connection between the material products of human culture and the ways in which large language models, reasoning engines, and other generative technologies are trained. Contemporary generative models are trained on truly vast quantities of data derived from digital collections. ChatGPT-4 for example was trained on 1 Petabyte of information ([Laizure, 2024](#)); that is roughly 25 times as much information as contained in the Library of Congress. While training uses curated collections, the volume of data required means that their curation is relatively light-touch. Generative models are left to gorge on the digital surrogates of the products of human culture. However this leads to a fundamental problem because the product of human culture are increasingly created with the help of the very generative technologies they are used to train. As over time more and more generative content is produced and

circulated the proportion of generative content used in training future generative models will inevitably grow.

Over the past year the volume of generative content produced and consumed has risen sharply; synthetic content now contributes a significant proportion of online materials. For example, a recent study showed that over half of all webpages are comprised of machine translated content ([Thompson et al., 2024](#)). This surprising statistic does not mean that half of the content that we consult online is machine generated; the majority of this content represents low-traffic web pages that consists of translations of content categorised as conversation and opinion and scraped to generate ad revenue ([Thompson et al., 2024](#)). Nevertheless it demonstrates the tendency towards the predominance of machined generated content. Indeed, recent studies suggest that a very high proportion new web pages contain some generative artificial intelligence content, with the majority of those showing complete, dominant or substantial use of generated content ([Law et al., 2025](#)). Generative content has also begun to predominate in the social media sphere. A growing proportion of content housed by the major social media platforms has been generated wholly or particularly using generative artificial intelligence technologies, including long and short-form content on video platforms like YouTube and TikTok, and marketing and promotional materials. Generative technologies have rapidly become an integral part of the digital content creation process.

Perhaps more worryingly, synthetic content is also infiltrating traditional sites of cultural production, particularly publishing, and the mainstream media. A growing proportion of academic papers are now produced entirely or substantially through the use of generative artificial intelligence. A recent study for example suggests that between 13% and 40% of papers in the biomedical field show evidence of generative artificial intelligence use ([Kodak et al., 2025](#)). This represents a significant risk, a gradual pollution of the well of collective knowledge, rendering not only those individual papers potentially unreliable, but also the meta-analysis and systematic reviews that draw on them. Generative content is also becoming more common in book publishing, particularly in self-publishing and small press markets. These synthetic products are frequently undifferentiated from the products of human culture, and their synthetic origins is frequently concealed. As we progress, generative artificial intelligence is likely to find a growing role in traditional sites of cultural production.

While generative tools are often used to produce entirely synthetic cultural texts, they are more commonly used to revise, refine and correct human cultural production, supplementing and complementing the creative process. These finer and more delicate contributions are invisibly woven into the cultural record, the faintest tracing of digital thread binding much of the information that is now produced and consumed. This kind of Integrated use of generative technologies represents a predominantly silent interventions in the cultural record. While it is still generally possible to differentiate

synthetic content on the level of the document, video, or image, the threading of those influences through the products of human cultural production are less easy to identify, untangle, and undo. This arguably represents a major structural transformation not only in cultural and intellectual production but also to the composition of the infosphere, one taking place with unprecedented rapidity and with comparatively little regulatory oversight.

While the rise of synthetic content may be of concern for the quality of public discourse, it is not in itself a threat to our capacity to differentiate between truth and lies, or information and misinformation. Low-quality, low-traffic content merely changes the noise to signal ratio, an irritation rather than a profound change. The problem with the gradual pollution of the infosphere is the systemic nature of error that arises. While both humans and machines are prone to error, the mistakes that humans make are generally not systemic, and therefore can cancel out ([Surowiecki, 2004](#)). There are exceptions to this reflecting the role of shared values, traditions and beliefs, and the cognitive biases to which everyone is prone (cf. [Tredinnick & Laybats, 2017](#); [Tredinnick, 2023](#)). However human error tends not to aggregate. Machine learning is different. The mistakes of machine learning technologies are often systemic, reflecting subtle biases in the ways in which they process information that result in small differences in probabilities associated with individual tokens. This kind of systemic bias can be self-reinforcing over time. Thus as the same content is repeatedly put through the mill of generative systems the outputs can begin to diverge significantly over time, and be subject to declining accuracy and reliability, and increased semantic drift.

Epistemic decay reflects the fact that today's synthetic cultural products are tomorrow's training materials. At present synthetic content represents only a tiny fraction of training data, but as our use of artificial intelligence begins to predominate in cultural production, that proportion will grow. This ouroboric loop arguably folds the epistemic base of generative models back onto themselves, eroding the diversity and originality of information over successive iterations, and recycling patterns, errors, and biases in a potentially catastrophic cycle of decline. Each time the products of generative technologies are folded-back into the training of future models, there is a narrowing of diversity and a tendency towards homogeneity. Like a technological Soylent Green, the dead synthetic products of prior generative technologies nurture and sustain each successive generation of model. As the use of generative technologies accelerates so does the gradual pollution of the infosphere.

Epistemic decay is a risk to generative technologies themselves as their training materials become progressively polluted with degraded or derivative content leading to a gradual decline in the quality of their outputs. This may become visible in gradual semantic drift, rising hallucinations, or a decline in factual accuracy. However improvement in modelling might obscure or delay the visible effect of that decay in the

medium-term. By the time epistemic decay becomes visible in the performance of generative technologies, it may be too late to reverse their effects. The result may be not a gradual degradation of generative outputs, but a sudden collapse of their generative capacity, a failure of meaning rather than of computation. Recent studies have shown for example that feeding the results of generative processes back into successive iterations of generative models can result in just this kind of sudden model collapse ([Shumailov et al., 2024](#)).

But more importantly epistemic decay may also conceivably present an existential threat. We risk stumbling past a tipping point after which it becomes essentially impossible or impracticable to recover epistemic certainty from the corrosive effects of the ouroboric nature of cultural recombination, to separate the products of human culture from their synthetic rivals spun through the cultural record like mycelium. As synthetic content is processed through successive generative processes it becomes gradually more difficult to untangle their influence. At that point questions of truth, value and meaning may be hard to articulate, let alone decide. It may conceivably be possible to reset the cultural record to a point before the rise of generative artificial intelligence, like a kind of grand cultural reboot, but as we become increasingly dependent on the integration of artificial intelligence, the social, political and economic consequences of that may themselves be catastrophic.

### **Risk and mitigation**

Epistemic decay is currently a speculative risk, and if realised its effects are likely to play out over the long term. There is time to mitigate the risk. There are three main drivers of the threat. The most obvious is the rapid rise of generative content in all spheres of cultural production. The second is our current inability to reliably differentiate between the products of artificial intelligence and the product of human cultural production. The third is the tendency for generative models to exhibit systematic bias and error. However, epistemic decay is not inevitable, and all three main drivers can be addressed.

Warding the epistemic decay is partly a matter of reliably differentiating between generative and human content, and acting on that knowledge. Digital watermarking of generative content, by for example embedding subtle biases in their generative algorithms that allow for subsequent detection without impacting on the reception of those outputs, may allow us to make this differentiation. OpenAI for example is exploring using embedded patterns to watermark the outputs of ChatGPT to allow future detection in principle ([OpenAI, 2024](#)). Learning algorithms may also present a solution. One of the strengths of neural network based artificial intelligence system is their pattern recognition abilities, and therefore artificial intelligence applications may themselves come to our aid in differentiating between human and synthetic content. Reliably differentiating between human and synthetic content will enable future models

to be trained on materials that do not imply a threat of epistemic decay. The use of human-in-the-loop reinforcement training may allow generative models to avoid the epistemic consequences of cultural recombination.

There is also a very urgent need for proper regulation of the industry and international co-operation on this, which may help mitigate these risks. That regulation should be predicated on an understanding of the subtleties of generative production, rather than rooted in issues such as intellectual property that reflect an entirely different model of cultural production.

Whether mitigation is a long term solution to epistemic decay, however, is unclear. At present generative technologies are dominated by a handful of major tech companies with significant resources. However, growing competition is likely to result in diversification, and perhaps a technological arms race in the production and detection of synthetic products. There are clear economic drivers for producing content that is as indistinguishable from that of human cultural production as it is possible to achieve, and few disincentives in doing so. As the cost of developing and running generative applications declines, those drivers may dominate and overwhelm the currently more cautious and socially responsible approaches.

### **September's issue of business information review**

September's issue of *Business Information* contains the traditional mix of professional and research articles focussed on issues of concern to information and knowledge professionals in the commercial sector. Our first research paper highlights issues in Knowledge Management framework development and implementation, with a particular focus on the Not-for-Profit sector. Entitled "Towards a Unified Knowledge Management Framework in Non-Profit Sector: The Case of Canada" the paper develops a framework named CANVAS-KM that contains a unique combination of Knowledge Management. Our second research paper also focusses on Knowledge Management. Entitled "Enhancing Open Innovation in University: Collaborative Frameworks and Knowledge Management Systems 4.0," the paper seeks to identify critical dimensions in building collaborative frameworks intended to enhance open innovation in universities, integrate this with Knowledge Management System 4.0, and identify critical success factors and barriers to open innovation implementation. The theme of Knowledge Management also informs our first professional paper by returning contributor Mostafa Sayyadi, "The Rise of New Pillars for Effective Knowledge Management Leadership."

The issue is completed by two papers of general and special concern. "The Artificial Intelligence (AI) Revolution: Evolving Business Decision-Making in the Digital Age" continues the theme of artificial intelligence in business contexts that has become a major concern in recent years. "Gamification for Digital Preservation Awareness: Encouraging Public Engagement with Digital Archives Through Game Design" returns to

the concept of gamification. This is an issue that Business Information has addressed before, and that continues to inform service contexts including libraries, museums and archives.

## **Funding**

The authors received no financial support for the research, authorship, and/or publication of this article.

## **Declaration of conflicting interests**

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

## **References**

Amankwah-Amoah J, Abdalla S, Mogaji E, et al. (2024) The impending disruption of creative industries by generative AI: opportunities, challenges, and research agenda. *International Journal of Information Management* 79: 102759. [Crossref](#). [Web of Science](#).

Bender S (2025) Generative-AI, the media industries, and the disappearance of human creative labour. *Media Practice and Education* 26(2): 200–217. [Crossref](#).

Habermas J, Lennox S, Lennox F (1974) The public sphere: an encyclopedia article (1964). *New German Critique* (3): 49–55, Originally published in Fischer-Lexikon Staat und Politik, Frankfurt am Main, 1964. [https://psi412.cankaya.edu.tr/uploads/files/Habermas,20Public20Sphere20\(1974\).pdf](https://psi412.cankaya.edu.tr/uploads/files/Habermas,20Public20Sphere20(1974).pdf) [Crossref](#).

Kobak D, González-Márquez R, Horvát E-Á, et al. (2025) Delving into LLM-assisted writing in biomedical publications through excess vocabulary. *Science Advances* 11(27): eadt3813. (Accessed 1 Aug 2025). [Crossref](#). [PubMed](#). [Web of Science](#).

Laizure SC (2024) Caution: ChatGPT doesn't know what you are asking and doesn't know what it is saying. *Journal of Pediatric Pharmacology and Therapeutics: JPPT: the Official Journal of PPAG* 29(5): 558–560, Available at. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11472406/> (Accessed 1 Aug 2025). [PubMed](#).

Law R, Guan X, Soulo T (2025) 74% of New Webpages Include AI Content (Study of 900k Pages). Ahrefs Blog. Available at. <https://ahrefs.com/blog/what-percentage-of-new-content-is-ai-generated/> (Accessed 1 Aug 2025).

Laybats C, Tredinnick L (2024) Ethics and artificial intelligence—the rise of the deep fake. *Business Information Review* 41(2): 46–47. [Crossref](#).



OpenAI. (2024) Understanding the source of what we see and hear online. 7 May 2024 (updated 4 August 2024). Available at: <https://openai.com/index/understanding-the-source-of-what-we-see-and-hear-online/> (Accessed: 9 August 2025).

Shanahan M (2015) *The Technological Singularity*. Cambridge, MA: MIT Press. [Crossref](#).

Shumailov I, Shumaylov Z, Zhao Y, et al. (2024) AI models collapse when trained on recursively generated data. *Nature* 631(8022): 755–759. [Crossref](#). [PubMed](#). [Web of Science](#).

Surowiecki J. (2004) *The wisdom of crowds: Why the many are smarter than the few*. Little, Brown.

Thompson B, Dhaliwal MP, Frisch P, et al. (2024) A shocking amount of the web is machine translated: insights from multi-way parallelism. arXiv preprint arXiv:2401.05749 [cs.CL], revised June 5 2024. Available at: <https://arxiv.org/abs/2401.05749> (Accessed 1 Aug 2025). [Crossref](#).

Tredinnick L (2017) Artificial intelligence and professional roles. *Business Information Review* 34(1): 37–41. [Crossref](#). [Web of Science](#).

Tredinnick Luke (2023) Dangerous data: analytics and information behaviour in the commercial world. *Business Information Review* 40(1): 10–20. [Crossref](#).

Tredinnick Luke, Laybats Claire (2017) Evaluating digital sources: trust, truth and lies. *Business Information Review* 34(4): 172–174. [Crossref](#).

Tredinnick Luke, Laybats Claire (2019) Reality filters: Disinformation and fake news. *Business Information Review* 36(3): 92–94. [Crossref](#).