Proceedings of the 5th International Conference on Art, Design and Social Sciences
DOI: 10.54254/2753-7064/2026.39294

The Synthetic Information Paradox: The Consequences of Abundant AI-Generated Content

Jiaming Chen

Veritas Collegiate Academy, Arlington, USA
michael.cjm.mcu@gmail.com

Abstract. Generative artificial intelligence (AI) has substantially reduced the cost of information production and thus accelerated this process. However, this does not guarantee that valid information becomes easier to obtain but may instead lead to the opposite. This thesis examines the mechanisms behind this change and proposes the synthetic information paradox: as the cost of producing information falls, the difficulty of identifying valid and useful information may rise. Studies have shown that AI models hallucinate at certain rates and generate fabricated, emotionally appealing, and low-utility content. This increases information noise and hence reduces the density of relevant and verified information. Therefore, users must devote more time, money, and expertise to verify collected information. Meanwhile, in a process this study calls aepistemic inheritance, unreliable content accumulates, is used in new content, and grows in negative influence. Eventually, harmful outcomes may happen in individual decision-making, scientific research, and the broader knowledge ecosystem.

Keywords: Artificial Intelligence, Valuable Information, Hallucination, Verification

1. Introduction

The amount of online content is already extremely abundant. Google estimates that its search index includes more than 100 million gigabytes of information, while another proxy, Common Crawl, reported that by August 2026 it had stored 360 tebibytes (TiB) of uncompressed content [1, 2]. Meanwhile, AI-generated information has become a substantial share of online content. About 35% of websites published by mid-2025 were classified as containing AI-generated or AI-assisted text [3]. Some people may publish large quantities of generated content to gain advertising revenue, attract search traffic, or quickly complete school and professional work, while others seek convenient explanations or emotionally supportive responses.

These figures do not guarantee the actual amount of AI-generated information encountered online because private, deleted, paywalled, and unindexed content is excluded. However, they suggest that synthetic information is no longer a minor portion of available online information. The quantity of invalid and academically useless information may also increase due to AI hallucinations, inappropriate use of AI, financial incentives to publish generated content quickly, and demand for merely emotional content. Once unreliable materials enter the public information environment, they may be reused without their flaws being recognized, potentially causing harm in consequential

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2753-7064/2026.39294

settings such as diagnosis and scientific research. It may also be incorporated into new AI-generated content, amplifying its effects.

Before proposing solutions, it is vital to understand the mechanism behind this effect. This study proposes a paradoxical theory: as the cost of producing information falls, the difficulty of identifying valid and useful information may rise. Although AI is not the only factor driving this phenomenon, it is an increasingly influential contributor.

2. Literature review

2.1. AI hallucination and misinformation

AI hallucination refers to the generation of content that appears valid while deviating from the user input, contradicting previously generated context, or misaligning with established knowledge [4]. Fluency and responsiveness cannot be treated as proof of accuracy, as AI hallucinations can obscure the inaccuracy of content in fluent language.

Models utilizing abstractive summaries have included statements that were not supported by the provided documents [5]. GPT-3, GPT-Neo/J, GPT-2, and a T5-based model tested on TruthfulQA were also reported to generate misconceptions that people commonly accept as true [6]. Retrieval and citation systems can reduce some errors, but a retrieved source may be irrelevant, incorrectly described, or fabricated [4]. Apart from technological issues, user input is also a significant drive to hallucination, as AI might prioritize user satisfaction over telling the truth [7].

Although AI-generated errors seem omnipresent among AI models, the error rate is not known for every model because it depends on the subject, language, prompt, evaluation method, and definition of error. However, even if this occurs occasionally, false information will keep contaminating the online information environment, and the effect may gradually amplify over time.

2.2. Information verification

It is generally difficult to determine the credibility of online information from fluent language alone. Instead, credibility assessment depends on the source, supporting evidence, independent corroboration, and the reader's knowledge [8]. Expertise has become more necessary for verification because it affects the details readers notice and how much weight they give to accuracy [9]. Skilled fact-checkers often incorporate lateral reading, leaving the original page to investigate its source, and comparing independent accounts before accepting a claim [10].

Verification methods vary with the information type and demand. A routine factual claim may be checked through authoritative records or corroboration from other sources. A scientific claim requires examination of methods, data, and uncertainty. A medical claim demands more caution because an incorrect claim can affect treatment decisions and theory design. In certain cases, researchers must replicate past experiments to verify their credibility, which requires time and resource.

Moreover, being aware that a claim is false does not necessarily end the influence of the information, as individuals may continue to use it after learning that it was inaccurate, especially when no better choices are provided [11]. People might not rigorously verify the content gathered or try to eliminate uncertainty due to the high cost, especially when they frequently use AI to generate information.

DOI: 10.54254/2753-7064/2026.39294

3. Growth in AI-generated content

3.1. Reduced cost of production

Generative AI greatly reduces the time and financial cost previously required for many tasks, including writing, analyzing documents, creating explanatory materials such as PowerPoints and videos, and providing guidance in daily life scenarios. In a controlled experiment on midlevel professional writing tasks, AI-assisted professionals, on average, spent 40% less time and produced results rated 18% higher in quality than the controlled group [12]. Another study on 5,172 customer-support agents measured an average 15% increase in productivity with AI assistance, with some samples demonstrating substantively better performances [13].

While making a revolutionary change to human society, AI is essentially not difficult to install and use, especially in everyday settings. In addition, many AI services offer free or low-cost models for users. These factors lead to rapid growth of AI user populations. As AI use grows and online publishing remains widely accessible, the volume of online content may expand rapidly.

3.2. Profit and demand

Convenience encourages production because customized text can be obtained without much time or money. On the other hand, profit-driven AI usage floods in. Today, media has attained an indisputable position in attracting search traffic and generating profits. AI therefore becomes a low-cost tool for commercial publishers. Independent content creators, lacking resources for human-made content, rely heavily on AI. Meanwhile, since these users are profit-oriented, the generated information may not be strictly verified or even deliberately exaggerated to incite emotional appeal, adding unreliable material to the online environment.

Non-commercial demand may also increase significantly. One example of which is demand for emotional support. Research on companion chatbots shows that users may value immediate emotional support from conversational systems [14]. This demand might originate from social problems, including loneliness, high work burden, and high academic pressure. AI usage among scholarly publications has also become substantial. In a large-scale analysis across 950,965 published papers, researchers detected signs of AI-generated text in proportions ranging from 6.3% in Mathematics papers and the Nature portfolio to 17.5% in Computer Science papers [15].

4. Decrease in valid information density

Information noise refers to materials that make relevant and reliable information more difficult to locate. As AI usage prevails, information noise may increase. Let T be the total amount of retrievable information and R the amount that is relevant and valid for a specific purpose. Valid information density is then $v = R/T$. As AI-generated information increases in abundance, verified content does not increase proportionally, resulting in lower valid information density and greater obstruction to finding reliable content.

As discussed in previous sections, AI demonstrates a statistically significant hallucination rate, and the fabricated information contaminates the online environment. Moreover, information does not necessarily need to be factually false to increase information noise. Generic explanations of products or problems, as well as vague or emotionally appealing content, can satisfy immediate demand, but the content may vary among users and thus create different arguments. This not only increases the amount of information individuals need to filter out to obtain concrete and useful data

DOI: 10.54254/2753-7064/2026.39294

but also lowers information variety. A recent study of web content, although it did not find a statistically significant decrease in factual accuracy, identified a negative relationship between AI-generated text and semantic diversity [3]. Another study on creative writing produced a similar result: individual stories improved, while the collections of stories became more alike [16].

5. Feedback loop diminishes valid information density

The unreliable and low-utility content generated by AI might be published online, intentionally or otherwise, by individual users or companies. Therefore, people using AI tools or browsing the web will all be influenced. Those who use research engines are directly exposed to contaminated information, while verification and filtering can consume time and energy. Others using AI may receive information from the sources used to train the models that are initially invalid or indirectly gather unreliable content on the internet using the online search function. Research has demonstrated a related mechanism. A recent study on GPT-2, GPT-3, and GPT-4 discovered that the models collapsed, characterized by irreversible defect—loss of original content distribution and accumulation of false information—after training on recursively generated data [17]. As AI is being used more frequently, the effect is amplified without proper restrictions, accelerating the growth of information noise.

6. Rising risk or verification

6.1. Expertise

AI can provide technical information or more emotionally supportive responses on various questions depending on users' intentions. In both cases, verification of the output is important due to the potential risk of inaccurate output. As shown in a recent study, thesis abstracts generated by AI have included fabricated data that some human reviewers found difficult to identify [18]. Expertise helps people identify various details, and it drives them to place more weight on accuracy and evidence when judging information [9]. However, the knowledge required for verification is usually not possessed by the original readers. In particular, when synthetic text enters medicine, science, or another specialized field, a nonexpert may need additional education or professional help to identify a serious error. While the barrier to enter using AI is low, verification can require substantial time or resources, and some may not even be aware of the importance of the verification process.

6.2. Source checking

When readers lack expertise in the subject information, they may need to verify sources carefully. Source checking requires readers to leave the original page, assess the publisher's credibility and the authors' credentials and interests, trace claims to primary evidence, and compare independent sources [8, 10]. Fabricated citations make this process slower because readers must first determine whether the sources exist before relying on them [19]. In other cases, when sources conflict, readers have to spend extra time determining which one is making the correct statements. From a larger scare, even a relatively low rate of error can create a substantial burden when thousands of claims require separate checking.

DOI: 10.54254/2753-7064/2026.39294

6.3. Replication

Replication can be the costliest form of verification commonly used in experimental research. As science advances, new studies, especially among the applied sciences, tend to depend more on existing findings. Using invalid information leads to misleading results, which may hinder scientific development and, in some cases, harm human lives. Scientific and medical commentaries maintain that AI assistance does not remove the need for expert oversight [20, 21]. Nevertheless, false information spreads on the internet and may be carried into later research papers.

A simple statement may be checked quickly by experts and without much cost, while an experimental or medical claim requires much longer time and substantial resource input, especially when the experiment includes complex procedures and expensive materials. This difference produces further risks. Because verification is costly, it is reasonable to expect that not every claim is checked before reusage. The unverified remainder then circulates alongside verified material, allowing uncertainty to enter the larger knowledge system.

7. Epistemic inheritance and the vulnerability of knowledge

Epistemic inheritance describes the propagation of an error into new content. This process may have three outcomes: unreliable content is detected and removed, it persists with limited reach or consequences, or it spreads, amplifies in influence, and causes serious harm. Undoubtedly, the third one possesses the greatest potential harm to human society.

As stated in the first scenario, platform moderators and fact-checkers are striving to detect and clear out false information online. A recent study on the social media platform X analyzed 82,076 marked misleading posts, 6,024 of which (7.33%) were identified as containing AI-generated content [22]. A 2026 report by Anthropic indicates a pattern of large-scale AI-assisted misinformation production: investigators identified an account using Claude to generate 70 fabricated news websites, and the content was further amplified by 70 other linked and matching social media accounts as well as more than 250 inauthentic commenting accounts [23]. These cases clearly demonstrate the influence and the accompanied risks of AI.

Regarding the second and third scenarios, it is difficult to estimate the abundance of such information. A significant portion of AI-generated misinformation has not yet been detected, and the content that leads to disastrous outcomes can also be difficult to trace back to its source.

8. Possible negative outcomes in dependent applications

As described in the third case of epistemic inheritance, undetected invalid information can lead to serious harm. Everyday decision-making may be deeply affected. Unequal digital literacy limits some users' ability to evaluate online claims [24, 25]. In addition, they may not have access to or may not be aware of how to find valid information for verification purposes. Therefore, they will believe in AI-generated false information from AI chats and other online sources, leading to failures and harm. The consequences are potentially more serious in diagnosis and treatment because the required expertise is more specialized and an inaccurate recommendation may cause more serious harm to human lives.

Published case reports have described harm linked to inaccurate AI-generated information. An article introduces a case in which a 60-year-old developed bromism after consulting the AI model ChatGPT [26]. In another case, a patient in his 30s suffering from a manic episode with psychotic features engaged in prolonged conversation with ChatGPT. As the chatbot corroborated the patient's

DOI: 10.54254/2753-7064/2026.39294

thoughts, he reinforced the belief that he was experiencing a "spiritual awakening." Only after subsequent intervention did he begin to recover [27].

Although academic research places greater emphasis on verifying information, the potential risks posed by misinformation may not be negligible. Within the scholarly ecosystem, false information fabricated by AI may leak into the public vision and later be cited as evidence in studies, particularly literature reviews that do not independently check their sources. As new studies are built on existing ones, some errors spread and persist. If researchers fail to detect these errors in experiments, later studies or applications may fail or, in high-stakes settings, cause physical harm.

9. Conclusion

Generative AI has substantially increased efficiency and reduced the time, financial cost, and expertise required to produce information. Driven by financial, psychological, and scholarly demand, individuals and organizations use AI to create content at an unprecedented scale. However, greater abundance of information does not necessarily make reliable knowledge easier to obtain but may lead to the opposite. This study proposes the synthetic information paradox to describe this phenomenon: as the cost of producing information declines, the difficulty of identifying valid and useful information may increase.

AI hallucination has been a widely recognized issue, resulting in the production of fabricated content. Depending on user input, AI also generates low-utility content, including highly repeated generic explanations and emotionally appealing materials, expanding total available information without supplementing a proportional increase of verified information. Consequently, valid information density may decline, requiring individuals to devote more time and resources, which vary from online source checking to replicating experiments, to verification.

The problem becomes more severe when unreliable information is spread onto the internet and repeatedly reused. Some errors may be detected and removed, while others, highly driven by the increasing cost of verification, may persist unnoticed in the information environment. The rest, through the process described in this study as epistemic inheritance, may be incorporated into later publications, AI outputs, research, and routine decision-making. Eventually, this process causes harm to human society, as misinformation may lead to disastrous results in research experiments, medical treatment, and individual problem-solving.

Generative AI has created the challenge of expanding inaccurate content, accompanied by increasingly costly information verification. Future research should pay more attention to investigating how to reduce AI-generated errors and make verification more efficient.

References

[1] Google Search. (2026). Google how Google Search organizes information. *Google Search*. Retrieved September 9, 2026, from https://www.google.com/intl/en_us/search/howsearchworks/how-search-works/organizing-information/

[2] Paris, M., Nagel, S., Foppiano, L., & Çelikkanat, H. (2026). August 2026 crawl archive now available. *Common Crawl*. Retrieved September 9, 2026, from https://commoncrawl.org/blog/august-2026-crawl-archive-now-available

[3] Dolezal, J., Alam, S., Graham, M., & Bohacek, M. (2026). The impact of AI-generated text on the internet. *arXiv*. https://arxiv.org/abs/2604.26965 doi: 10.48550/arXiv.2604.26965

[4] Zhang, Y., Li, Y., Cui, L., et al. (2025). Siren's song in the AI ocean: A survey on hallucination in large language models. *Computational Linguistics, 51*(4): 1373–418. doi: 10.1162/COLI.a.16

[5] Maynez, J., Narayan, S., Bohnet, B., & McDonald, R. (2020). On faithfulness and factuality in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp.

DOI: 10.54254/2753-7064/2026.39294

1906–1919). Association for Computational Linguistics. https://aclanthology.org/2020.acl-main.173/ doi: 10.18653/v1/2020.acl-main.173

[6] Lin, S., Hilton, J., & Evans, O. (2022). TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 3214–3252). Association for Computational Linguistics. https://aclanthology.org/2022.acl-long.229/ doi: 10.18653/v1/2022.acl-long.229

[7] Sharma, M., Tong, M., Korbak, T., et al. (2024). Towards understanding sycophancy in language models. *International Conference on Learning Representations*. http://arxiv.org/abs/2310.13548 doi: 10.48550/arXiv.2310.13548

[8] Metzger, M. J. (2007). Making sense of credibility on the web: Models for evaluating online information and recommendations for future research. *Journal of the American Society for Information Science and Technology, 58*(13), 2078-2091. doi: 10.1002/asi.20672

[9] Lucassen, T., & Schraagen, J. M. (2011). Factual accuracy and trust in information: The role of expertise. *Journal of the American Society for Information Science and Technology, 62*(7), 1232–42. doi: 10.1002/asi.21545

[10] Wineburg, S., McGrew, S. (2019). Lateral reading and the nature of expertise: Reading less and learning more when evaluating digital information. *Teach Coll Rec, 121*(11), 1–40. doi: 10.1177/016146811912101102

[11] Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest, 13*(3), 106–31. doi: 10.1177/1529100612451018

[12] Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187–92. doi: 10.1126/science.adh2586

[13] Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Q J Econ, 140*(2), 889–942. doi: 10.1093/qje/qjae044

[14] Ta, V., Griffith, C., Boatfield, C., et al. (2020). User experiences of social support from companion chatbots in everyday contexts: Thematic analysis. *J Med Internet Res, 22*(3), e16235. doi: 10.2196/16235

[15] Liang, W., Zhang, Y., Wu, Z., et al. (2024). Mapping the increasing use of LLMs in scientific papers. *Conference on Language Modeling*. https://openreview.net/pdf?id=YX7QnhxESU doi: 10.48550/arXiv.2404.01268

[16] Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances, 10*(28), eadn5290. doi: 10.1126/sciadv.adn5290

[17] Shumailov, I., Shumaylov, Z., Zhao, Y., et al. (2024). AI models collapse when trained on recursively generated data. *Nature, 631*(8022), 755–759. doi: 10.1038/s41586-024-07566-y

[18] Gao, C. A., Howard, F. M., Markov, N. S., et al. (2023). Comparing scientific abstracts generated by ChatGPT to real abstracts with detectors and blinded human reviewers. *NPJ Digit Med, 6*, 75. doi: 10.1038/s41746-023-00819-6

[19] Walters, W. H., & Wilder, E. I. (2023). Fabrication and errors in the bibliographic citations generated by ChatGPT. *Scientific Reports, 13*, 14045. https://doi.org/10.1038/s41598-023-41032-5

[20] Salvagno, M., Taccone, F. S., & Gerli, A. G. (2023). Artificial intelligence hallucinations. *Critical Care, 27*(1), 180. https://doi.org/10.1186/s13054-023-04473-y

[21] van Dis, E. A. M., Bollen, J., Zuidema, W., van Rooij, R., & Bockting, C. L. (2023). ChatGPT: Five priorities for research. *Nature, 614*(7947), 224–226. https://doi.org/10.1038/d41586-023-00288-7

[22] Drolsbach, C., Demirel, E., & Pröllochs, N. (2026). Characterizing AI-generated misinformation on social media. *arXiv*. http://arxiv.org/abs/2505.10266 https://doi.org/10.48550/arXiv.2505.10266

[23] Anthropic. (2026). Detecting and countering misuse of AI: September 2026. Retrieved September 22, 2026, from https://www.anthropic.com/threat-intelligence-report-september-2026

[24] Misra, S., & Stokols, D. (2012). Psychological and health outcomes of perceived information overload. *Environment and Behavior, 44*(6), 737–759. https://doi.org/10.1177/0013916511404408

[25] Kozyreva, A., Lewandowsky, S., & Hertwig, R. (2020). Citizens versus the internet: Confronting digital challenges with cognitive tools. *Psychological Science in the Public Interest, 21*(3), 103-156. doi: 10.1177/1529100620946707

[26] Eichenberger, A., Thielke, S., & Van Buskirk, A. (2025). A case of bromism influenced by use of artificial intelligence. *Annals of Internal Medicine Clinical Cases, 4*(8). doi: 10.7326/aimcc.2024.1260

[27] Shah, S., & Morrin, H. (2026). Substance-induced manic psychosis in which delusions were corroborated by a chatbot: Case report. *BMC Psychiatry, 26*(1). doi: 10.1186/s12888-026-08137-3 PubMed PMID: 42243814.