



Research Note

GPT-fabricated scientific papers on Google Scholar: Key features, spread, and implications for preempting evidence manipulation

Academic journals, archives, and repositories are seeing an increasing number of questionable research papers clearly produced using generative AI. They are often created with widely available, general-purpose AI applications, most likely ChatGPT, and mimic scientific writing. Google Scholar easily locates and lists these questionable papers alongside reputable, quality-controlled research. Our analysis of a selection of questionable GPT-fabricated scientific papers found in Google Scholar shows that many are about applied, often controversial topics susceptible to disinformation: the environment, health, and computing. The resulting enhanced potential for malicious manipulation of society's evidence base, particularly in politically divisive domains, is a growing concern.

Authors: Jutta Haider (1), Kristofer Rolf Söderström (2), Björn Ekström (1), Malte Rödl (3)

Affiliations: (1) Swedish School of Library and Information Science, University of Borås, Sweden, (2) Department of Arts and Cultural Sciences, Lund University, Sweden, (3) Division of Environmental Communication, Swedish University of Agricultural Sciences, Sweden

How to cite: Haider, J., Söderström, K. R., Ekström, B., & Rödl, M. (2024). GPT-fabricated scientific papers on Google Scholar: Key features, spread, and implications for preempting evidence manipulation. *Harvard Kennedy School (HKS) Misinformation Review*, 5(5).

Received: May 20th, 2024. Accepted: August 14th, 2024. Published: September 3rd, 2024.

Research questions

- Where are questionable publications produced with generative pre-trained transformers (GPTs) that can be found via Google Scholar published or deposited?
- What are the main characteristics of these publications in relation to predominant subject categories?
- How are these publications spread in the research infrastructure for scholarly communication?
- How is the role of the scholarly communication infrastructure challenged in maintaining public trust in science and evidence through inappropriate use of generative AI?

Research note summary

- A sample of scientific papers with signs of GPT-use found on Google Scholar was retrieved, downloaded, and analyzed using a combination of qualitative coding and descriptive statistics. All papers contained at least one of two common phrases returned by conversational agents that use

¹ A publication of the Shorenstein Center on Media, Politics and Public Policy at Harvard University, John F. Kennedy School of Government.

large language models (LLM) like OpenAI's ChatGPT. Google Search was then used to determine the extent to which copies of questionable, GPT-fabricated papers were available in various repositories, archives, citation databases, and social media platforms.

- Roughly two-thirds of the retrieved papers were found to have been produced, at least in part, through undisclosed, potentially deceptive use of GPT. The majority (57%) of these questionable papers dealt with policy-relevant subjects (i.e., environment, health, computing), susceptible to influence operations. Most were available in several copies on different domains (e.g., social media, archives, and repositories).
- Two main risks arise from the increasingly common use of GPT to (mass-)produce fake, scientific publications. First, the abundance of fabricated “studies” seeping into all areas of the research infrastructure threatens to overwhelm the scholarly communication system and jeopardize the integrity of the scientific record. A second risk lies in the increased possibility that convincingly scientific-looking content was in fact deceitfully created with AI tools and is also optimized to be retrieved by publicly available academic search engines, particularly Google Scholar. However small, this possibility and awareness of it risks undermining the basis for trust in scientific knowledge and poses serious societal risks.

Implications

The use of ChatGPT to generate text for academic papers has raised concerns about research integrity. Discussion of this phenomenon is ongoing in editorials, commentaries, opinion pieces, and on social media (Bom, 2023; Stokel-Walker, 2024; Thorp, 2023). There are now several lists of papers suspected of GPT misuse, and new papers are constantly being added.² While many legitimate uses of GPT for research and academic writing exist (Huang & Tan, 2023; Kitamura, 2023; Lund et al., 2023), its undeclared use—beyond proofreading—has potentially far-reaching implications for both science and society, but especially for their relationship. It, therefore, seems important to extend the discussion to one of the most accessible and well-known intermediaries between science, but also certain types of misinformation, and the public, namely Google Scholar, also in response to the legitimate concerns that the discussion of generative AI and misinformation needs to be more nuanced and empirically substantiated (Simon et al., 2023).

Google Scholar, <https://scholar.google.com>, is an easy-to-use academic search engine. It is available for free, and its index is extensive (Gusenbauer & Haddaway, 2020). It is also often touted as a credible source for academic literature and even recommended in library guides, by media and information literacy initiatives, and fact checkers (Tripodi et al., 2023). However, Google Scholar lacks the transparency and adherence to standards that usually characterize citation databases. Instead, Google Scholar uses automated crawlers, like Google’s web search engine (Martín-Martín et al., 2021), and the inclusion criteria are based on primarily technical standards, allowing any individual author—with or without scientific affiliation—to upload papers to be indexed (Google Scholar Help, n.d.). It has been shown that Google Scholar is susceptible to manipulation through citation exploits (Antkare, 2020) and by providing access to fake scientific papers (Dadkhah et al., 2017). A large part of Google Scholar’s index consists of publications from established scientific journals or other forms of quality-controlled, scholarly literature. However, the index also contains a large amount of gray literature, including student papers, working papers, reports, preprint servers, and academic networking sites, as well as material from so-called “questionable” academic journals, including paper mills. The search interface does not offer the possibility

² See for example Academ-AI, <https://www.academ-ai.info/>, and Retraction Watch, <https://retractionwatch.com/papers-and-peer-reviews-with-evidence-of-chatgpt-writing/>.

to filter the results meaningfully by material type, publication status, or form of quality control, such as limiting the search to peer-reviewed material.

To understand the occurrence of ChatGPT (co-)authored work in Google Scholar's index, we scraped it for publications, including one of two common ChatGPT responses (see Appendix A) that we encountered on social media and in media reports (DeGeurin, 2024). The results of our descriptive statistical analyses showed that around 62% did not declare the use of GPTs. Most of these GPT-fabricated papers were found in non-indexed journals and working papers, but some cases included research published in mainstream scientific journals and conference proceedings.³ More than half (57%) of these GPT-fabricated papers concerned policy-relevant subject areas susceptible to influence operations. To avoid increasing the visibility of these publications, we abstained from referencing them in this research note. However, we have made the data available in the Harvard Dataverse repository.

The publications were related to three issue areas—health (14.5%), environment (19.5%) and computing (23%)—with key terms such “healthcare,” “COVID-19,” or “infection” for health-related papers, and “analysis,” “sustainable,” and “global” for environment-related papers. In several cases, the papers had titles that strung together general keywords and buzzwords, thus alluding to very broad and current research. These terms included “biology,” “telehealth,” “climate policy,” “diversity,” and “disrupting,” to name just a few. While the study's scope and design did not include a detailed analysis of which parts of the articles included fabricated text, our dataset did contain the surrounding sentences for each occurrence of the suspicious phrases that formed the basis for our search and subsequent selection. Based on that, we can say that the phrases occurred in most sections typically found in scientific publications, including the literature review, methods, conceptual and theoretical frameworks, background, motivation or societal relevance, and even discussion. This was confirmed during the joint coding, where we read and discussed all articles. It became clear that not just the text related to the telltale phrases was created by GPT, but that almost all articles in our sample of questionable articles likely contained traces of GPT-fabricated text everywhere.

Evidence hacking and backfiring effects

Generative pre-trained transformers (GPTs) can be used to produce texts that mimic scientific writing. These texts, when made available online—as we demonstrate—leak into the databases of academic search engines and other parts of the research infrastructure for scholarly communication. This development exacerbates problems that were already present with less sophisticated text generators (Antkare, 2020; Cabanac & Labbé, 2021). Yet, the public release of ChatGPT in 2022, together with the way Google Scholar works, has increased the likelihood of lay people (e.g., media, politicians, patients, students) coming across questionable (or even entirely GPT-fabricated) papers and other problematic research findings. Previous research has emphasized that the ability to determine the value and status of scientific publications for lay people is at stake when misleading articles are passed off as reputable (Haider & Åström, 2017) and that systematic literature reviews risk being compromised (Dadkhah et al., 2017). It has also been highlighted that Google Scholar, in particular, can be and has been exploited for manipulating the evidence base for politically charged issues and to fuel conspiracy narratives (Tripodi et al., 2023). Both concerns are likely to be magnified in the future, increasing the risk of what we suggest calling *evidence hacking*—the strategic and coordinated malicious manipulation of society's evidence base.

The authority of quality-controlled research as evidence to support legislation, policy, politics, and other forms of decision-making is undermined by the presence of undeclared GPT-fabricated content in

³ Indexed journals mean scholarly journals indexed by abstract and citation databases such as Scopus and Web of Science, where the indexation implies journals with high scientific quality. Non-indexed journals are journals that fall outside of this indexation.

publications professing to be scientific. Due to the large number of archives, repositories, mirror sites, and shadow libraries to which they spread, there is a clear risk that GPT-fabricated, questionable papers will reach audiences even after a possible retraction. There are considerable technical difficulties involved in identifying and tracing computer-fabricated papers (Cabanac & Labbé, 2021; Dadkhah et al., 2023; Jones, 2024), not to mention preventing and curbing their spread and uptake.

However, as the rise of the so-called anti-vaxx movement during the COVID-19 pandemic and the ongoing obstruction and denial of climate change show, retracting erroneous publications often fuels conspiracies and increases the following of these movements rather than stopping them. To illustrate this mechanism, climate deniers frequently question established scientific consensus by pointing to other, supposedly scientific, studies that support their claims. Usually, these are poorly executed, not peer-reviewed, based on obsolete data, or even fraudulent (Dunlap & Brulle, 2020). A similar strategy is successful in the alternative epistemic world of the global anti-vaccination movement (Carrion, 2018) and the persistence of flawed and questionable publications in the scientific record already poses significant problems for health research, policy, and lawmakers, and thus for society as a whole (Littell et al., 2024). Considering that a person's support for "doing your own research" is associated with increased mistrust in scientific institutions (Chinn & Hasell, 2023), it will be of utmost importance to anticipate and consider such backfiring effects already when designing a technical solution, when suggesting industry or legal regulation, and in the planning of educational measures.

Recommendations

Solutions should be based on simultaneous considerations of technical, educational, and regulatory approaches, as well as incentives, including social ones, across the entire research infrastructure. Paying attention to how these approaches and incentives relate to each other can help identify points and mechanisms for disruption. Recognizing fraudulent academic papers must happen alongside understanding how they reach their audiences and what reasons there might be for some of these papers successfully "sticking around." A possible way to mitigate some of the risks associated with GPT-fabricated scholarly texts finding their way into academic search engine results would be to provide filtering options for facets such as indexed journals, gray literature, peer-review, and similar on the interface of publicly available academic search engines. Furthermore, evaluation tools for indexed journals⁴ could be integrated into the graphical user interfaces and the crawlers of these academic search engines. To enable accountability, it is important that the index (database) of such a search engine is populated according to criteria that are transparent, open to scrutiny, and appropriate to the workings of science and other forms of academic research. Moreover, considering that Google Scholar has no real competitor, there is a strong case for establishing a freely accessible, non-specialized academic search engine that is not run for commercial reasons but for reasons of public interest. Such measures, together with educational initiatives aimed particularly at policymakers, science communicators, journalists, and other media workers, will be crucial to reducing the possibilities for and effects of malicious manipulation or evidence hacking. It is important not to present this as a technical problem that exists only because of AI text generators but to relate it to the wider concerns in which it is embedded. These range from a largely dysfunctional scholarly publishing system (Haider & Åström, 2017) and academia's "publish or perish" paradigm to Google's near-monopoly and ideological battles over the control of information and ultimately knowledge. Any intervention is likely to have systemic effects; these effects need to be considered and assessed in advance and, ideally, followed up on.

Our study focused on a selection of papers that were easily recognizable as fraudulent. We used this relatively small sample as a magnifying glass to examine, delineate, and understand a problem that goes

⁴ Such as LiU Journal CheckUp, <https://ep.liu.se/JournalCheckup/default.aspx?lang=eng>.

beyond the scope of the sample itself, which however points towards larger concerns that require further investigation. The work of ongoing whistleblowing initiatives⁵, recent media reports of journal closures (Subbaraman, 2024), or GPT-related changes in word use and writing style (Cabanac et al., 2021; Stokel-Walker, 2024) suggest that we only see the tip of the iceberg. There are already more sophisticated cases (Dadkhah et al., 2023) as well as cases involving fabricated images (Gu et al., 2022). Our analysis shows that questionable and potentially manipulative GPT-fabricated papers permeate the research infrastructure and are likely to become a widespread phenomenon. Our findings underline that the risk of fake scientific papers being used to maliciously manipulate evidence (see Dadkhah et al., 2017) must be taken seriously. Manipulation may involve undeclared automatic summaries of texts, inclusion in literature reviews, explicit scientific claims, or the concealment of errors in studies so that they are difficult to detect in peer review. However, the mere possibility of these things happening is a significant risk in its own right that can be strategically exploited and will have ramifications for trust in and perception of science. Society's methods of evaluating sources and the foundations of media and information literacy are under threat and public trust in science is at risk of further erosion, with far-reaching consequences for society in dealing with information disorders. To address this multifaceted problem, we first need to understand why it exists and proliferates.

Findings

Finding 1: 139 GPT-fabricated, questionable papers were found and listed as regular results on the Google Scholar results page. Non-indexed journals dominate.

Most questionable papers we found were in non-indexed journals or were working papers, but we did also find some in established journals, publications, conferences, and repositories. We found a total of 139 papers with a suspected deceptive use of ChatGPT or similar LLM applications (see Table 1). Out of these, 19 were in indexed journals, 89 were in non-indexed journals, 19 were student papers found in university databases, and 12 were working papers (mostly in preprint databases). Table 1 divides these papers into categories. Health and environment papers made up around 34% (47) of the sample. Of these, 66% were present in non-indexed journals.

Table 1. Number of papers across topics and venues using ChatGPT fraudulently or undeclared.

Paper category	Computing	Environment	Health	Others	Total
Indexed journals*	5	3	4	7	19
Non-indexed journals	18	18	13	40	89
Student papers	4	3	1	11	19
Working papers	5	3	2	2	12
Total	32	27	20	60	139

* Indexed by Scopus, Norwegian register for scientific journals, series and publishers, WoS and/or DOAJ.

⁵ Such as Academ-AI, <https://www.academ-ai.info/>, and Retraction Watch, <https://retractionwatch.com/papers-and-peer-reviews-with-evidence-of-chatgpt-writing/>.

Finding 2: GPT-fabricated, questionable papers are disseminated online, permeating the research infrastructure for scholarly communication, often in multiple copies. Applied topics with practical implications dominate.

The 20 papers concerning health-related issues are distributed across 20 unique domains, accounting for 46 URLs. The 27 papers dealing with environmental issues can be found across 26 unique domains, accounting for 56 URLs. Most of the identified papers exist in multiple copies and have already spread to several archives, repositories, and social media. It would be difficult, or impossible, to remove them from the scientific record.

As apparent from Table 2, GPT-fabricated, questionable papers are seeping into most parts of the online research infrastructure for scholarly communication. Platforms on which identified papers have appeared include ResearchGate, ORCID, Journal of Population Therapeutics and Clinical Pharmacology (JPTCP), EasyChair, Frontiers, the Institute of Electrical and Electronics Engineer (IEEE), and X/Twitter. Thus, even if they are retracted from their original source, it will prove very difficult to track, remove, or even just mark them up on other platforms. Moreover, unless regulated, Google Scholar will enable their continued and most likely unlabeled discoverability.

Table 2. Top domains by subject.

Subject	1	2	3	4	5
Environment	researchgate.net (13)	orcid.org (4)	easychair.org (3)	ijope.com* (3)	publikasiindonesia.id (3)
Health	researchgate.net (15)	ieee.org (4)	twitter.com (3)	jptcp.com** (2)	frontiersin.org (2)

* *International Journal of Open Publication and Exploration (ISSN: 3006-2853)*

** *The Journal of Population Therapeutics and Clinical Pharmacology (ISSN 2561-8741)*

Note: We removed the original publication URL to avoid double counting.

A word rain visualization (Centre for Digital Humanities Uppsala, 2023), which combines word prominences through TF-IDF⁶ scores with semantic similarity of the full texts of our sample of GPT-generated articles that fall into the “Environment” and “Health” categories, reflects the two categories in question. However, as can be seen in Figure 1, it also reveals overlap and sub-areas. The y-axis shows word prominences through word positions and font sizes, while the x-axis indicates semantic similarity. In addition to a certain amount of overlap, this reveals sub-areas, which are best described as two distinct events within the word rain. The event on the left bundles terms related to the development and management of health and healthcare with “challenges,” “impact,” and “potential of artificial intelligence” emerging as semantically related terms. Terms related to research infrastructures, environmental, epistemic, and technological concepts are arranged further down in the same event (e.g., “system,” “climate,” “understanding,” “knowledge,” “learning,” “education,” “sustainable”). A second distinct event further to the right bundles terms associated with fish farming and aquatic medicinal plants, highlighting the presence of an aquaculture cluster. Here, the prominence of groups of terms such as “used,” “model,” “-based,” and “traditional” suggests the presence of applied research on these topics.

⁶ *Term frequency–inverse document frequency*, a method for measuring the significance of a word in a document compared to its frequency across all documents in a collection.

The two events making up the word rain visualization, are linked by a less dominant but overlapping cluster of terms related to “energy” and “water.”

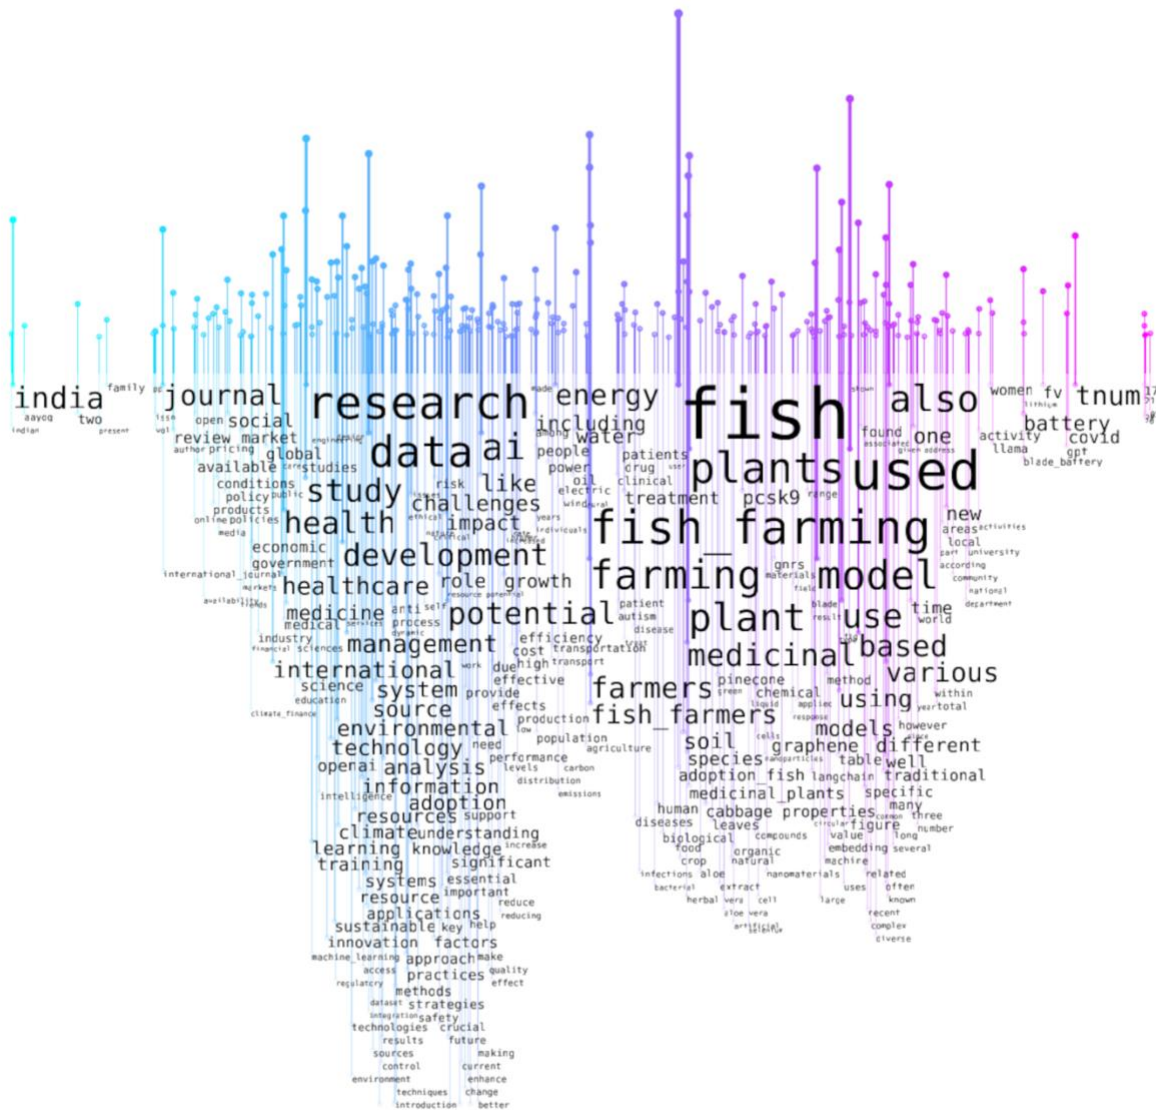


Figure 1. Word rain of environment- and health-related GPT-fabricated, questionable full-text papers.

The bar chart of the terms in the paper subset (see Figure 2) complements the word rain visualization by depicting the most prominent terms in the full texts along the y-axis. Here, word prominences across health and environment papers are arranged descendingly, where values outside parentheses are TF-IDF values (relative frequencies) and values inside parentheses are raw term frequencies (absolute frequencies).




Figure 2. Most prominent terms in environment- and health-related GPT-fabricated, questionable full-text papers.

Finding 3: Google Scholar presents results from quality-controlled and non-controlled citation databases on the same interface, providing unfiltered access to GPT-fabricated questionable papers.

Google Scholar's central position in the publicly accessible scholarly communication infrastructure, as well as its lack of standards, transparency, and accountability in terms of inclusion criteria, has potentially serious implications for public trust in science. This is likely to exacerbate the already-known potential to exploit Google Scholar for evidence hacking (Tripodi et al., 2023) and will have implications for any attempts to retract or remove fraudulent papers from their original publication venues. Any solution must consider the entirety of the research infrastructure for scholarly communication and the interplay of different actors, interests, and incentives.

Methods

We searched and scraped Google Scholar using the Python library Scholarly (Cholewiak et al., 2023) for papers that included specific phrases known to be common responses from ChatGPT and similar applications with the same underlying model (GPT3.5 or GPT4): “as of my last knowledge update” and/or “I don’t have access to real-time data” (see Appendix A). This facilitated the identification of papers that likely used generative AI to produce text, resulting in 227 retrieved papers. The papers’ bibliographic information was automatically added to a spreadsheet and downloaded into Zotero.⁷

We employed multiple coding (Barbour, 2001) to classify the papers based on their content. First, we jointly assessed whether the paper was suspected of fraudulent use of ChatGPT (or similar) based on how

⁷ An open-source reference manager, <https://zotero.org>.

the text was integrated into the papers and whether the paper was presented as original research output or the AI tool's role was acknowledged. Second, in analyzing the content of the papers, we continued the multiple coding by classifying the fraudulent papers into four categories identified during an initial round of analysis—health, environment, computing, and others—and then determining which subjects were most affected by this issue (see Table 1). Out of the 227 retrieved papers, 88 papers were written with legitimate and/or declared use of GPTs (i.e., false positives, which were excluded from further analysis), and 139 papers were written with undeclared and/or fraudulent use (i.e., true positives, which were included in further analysis). The multiple coding was conducted jointly by all authors of the present article, who collaboratively coded and cross-checked each other's interpretation of the data simultaneously in a shared spreadsheet file. This was done to single out coding discrepancies and settle coding disagreements, which in turn ensured methodological thoroughness and analytical consensus (see Barbour, 2001). Redoing the category coding later based on our established coding schedule, we achieved an intercoder reliability (Cohen's kappa) of 0.806 after eradicating obvious differences.

The ranking algorithm of Google Scholar prioritizes highly cited and older publications (Martín-Martín et al., 2016). Therefore, the position of the articles on the search engine results pages was not particularly informative, considering the relatively small number of results in combination with the recency of the publications. Only the query “as of my last knowledge update” had more than two search engine result pages. On those, questionable articles with undeclared use of GPTs were evenly distributed across all result pages (min: 4, max: 9, mode: 8), with the proportion of undeclared use being slightly higher on average on later search result pages.

To understand how the papers making fraudulent use of generative AI were disseminated online, we programmatically searched for the paper titles (with exact string matching) in Google Search from our local IP address (see Appendix B) using the `googlesearch-python` library (Vikramaditya, 2020). We manually verified each search result to filter out false positives—results that were not related to the paper—and then compiled the most prominent URLs by field. This enabled the identification of other platforms through which the papers had been spread. We did not, however, investigate whether copies had spread into SciHub or other shadow libraries, or if they were referenced in Wikipedia.

We used descriptive statistics to count the prevalence of the number of GPT-fabricated papers across topics and venues and top domains by subject. The `pandas` software library for the Python programming language (The pandas development team, 2024) was used for this part of the analysis. Based on the multiple coding, paper occurrences were counted in relation to their categories, divided into indexed journals, non-indexed journals, student papers, and working papers. The schemes, subdomains, and subdirectories of the URL strings were filtered out while top-level domains and second-level domains were kept, which led to normalizing domain names. This, in turn, allowed the counting of domain frequencies in the environment and health categories. To distinguish word prominences and meanings in the environment and health-related GPT-fabricated questionable papers, a semantically-aware word cloud visualization was produced through the use of a word rain (Centre for Digital Humanities Uppsala, 2023) for full-text versions of the papers. Font size and y-axis positions indicate word prominences through TF-IDF scores for the environment and health papers (also visualized in a separate bar chart with raw term frequencies in parentheses), and words are positioned along the x-axis to reflect semantic similarity (Skeppstedt et al., 2024), with an English Word2vec skip gram model space (Fares et al., 2017). An English stop word list was used, along with a manually produced list including terms such as “https,” “volume,” or “years.”

Bibliography

- Antkare, I. (2020). Ike Antkare, his publications, and those of his disciples. In M. Biagioli & A. Lippman (Eds.), *Gaming the metrics* (pp. 177–200). The MIT Press.
<https://doi.org/10.7551/mitpress/11087.003.0018>
- Barbour, R. S. (2001). Checklists for improving rigour in qualitative research: A case of the tail wagging the dog? *BMJ*, 322(7294), 1115–1117. <https://doi.org/10.1136/bmj.322.7294.1115>
- Bom, H.-S. H. (2023). Exploring the opportunities and challenges of ChatGPT in academic writing: A roundtable discussion. *Nuclear Medicine and Molecular Imaging*, 57(4), 165–167.
<https://doi.org/10.1007/s13139-023-00809-2>
- Cabanac, G., & Labbé, C. (2021). Prevalence of nonsensical algorithmically generated papers in the scientific literature. *Journal of the Association for Information Science and Technology*, 72(12), 1461–1476. <https://doi.org/10.1002/asi.24495>
- Cabanac, G., Labbé, C., & Magazinov, A. (2021). *Tortured phrases: A dubious writing style emerging in science. Evidence of critical issues affecting established journals*. arXiv.
<https://doi.org/10.48550/arXiv.2107.06751>
- Carrion, M. L. (2018). “You need to do your research”: Vaccines, contestable science, and maternal epistemology. *Public Understanding of Science*, 27(3), 310–324.
<https://doi.org/10.1177/0963662517728024>
- Centre for Digital Humanities Uppsala (2023). CDHUppsala/word-rain [Computer software].
<https://github.com/CDHUppsala/word-rain>
- Chinn, S., & Hasell, A. (2023). Support for “doing your own research” is associated with COVID-19 misperceptions and scientific mistrust. *Harvard Kennedy School (HSK) Misinformation Review*, 4(3). <https://doi.org/10.37016/mr-2020-117>
- Cholewiak, S. A., Ipeirotis, P., Silva, V., & Kannawadi, A. (2023). SCHOLARLY: Simple access to Google Scholar authors and citation using Python (1.5.0) [Computer software].
<https://doi.org/10.5281/zenodo.5764801>
- Dadkhah, M., Lagzian, M., & Borchardt, G. (2017). Questionable papers in citation databases as an issue for literature review. *Journal of Cell Communication and Signaling*, 11(2), 181–185.
<https://doi.org/10.1007/s12079-016-0370-6>
- Dadkhah, M., Oermann, M. H., Hegedüs, M., Raman, R., & Dávid, L. D. (2023). Detection of fake papers in the era of artificial intelligence. *Diagnosis*, 10(4), 390–397. <https://doi.org/10.1515/dx-2023-0090>
- DeGeurin, M. (2024, March 19). *AI-generated nonsense is leaking into scientific journals*. Popular Science. <https://www.popsci.com/technology/ai-generated-text-scientific-journals/>
- Dunlap, R. E., & Brulle, R. J. (2020). Sources and amplifiers of climate change denial. In D.C. Holmes & L. M. Richardson (Eds.), *Research handbook on communicating climate change* (pp. 49–61). Edward Elgar Publishing. <https://doi.org/10.4337/9781789900408.00013>
- Fares, M., Kutuzov, A., Oepen, S., & Velldal, E. (2017). Word vectors, reuse, and replicability: Towards a community repository of large-text resources. In J. Tiedemann & N. Tahmasebi (Eds.), *Proceedings of the 21st Nordic Conference on Computational Linguistics* (pp. 271–276). Association for Computational Linguistics. <https://aclanthology.org/W17-0237>
- Google Scholar Help. (n.d.). *Inclusion guidelines for webmasters*.
<https://scholar.google.com/intl/en/scholar/inclusion.html>
- Gu, J., Wang, X., Li, C., Zhao, J., Fu, W., Liang, G., & Qiu, J. (2022). AI-enabled image fraud in scientific publications. *Patterns*, 3(7), 100511. <https://doi.org/10.1016/j.patter.2022.100511>

- Gusenbauer, M., & Haddaway, N. R. (2020). Which academic search systems are suitable for systematic reviews or meta-analyses? Evaluating retrieval qualities of Google Scholar, PubMed, and 26 other resources. *Research Synthesis Methods*, 11(2), 181–217. <https://doi.org/10.1002/jrsm.1378>
- Haider, J., & Åström, F. (2017). Dimensions of trust in scholarly communication: Problematizing peer review in the aftermath of John Bohannon's "Sting" in science. *Journal of the Association for Information Science and Technology*, 68(2), 450–467. <https://doi.org/10.1002/asi.23669>
- Huang, J., & Tan, M. (2023). The role of ChatGPT in scientific communication: Writing better scientific review articles. *American Journal of Cancer Research*, 13(4), 1148–1154. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10164801/>
- Jones, N. (2024). How journals are fighting back against a wave of questionable images. *Nature*, 626(8000), 697–698. <https://doi.org/10.1038/d41586-024-00372-6>
- Kitamura, F. C. (2023). ChatGPT is shaping the future of medical writing but still requires human judgment. *Radiology*, 307(2), e230171. <https://doi.org/10.1148/radiol.230171>
- Littell, J. H., Abel, K. M., Biggs, M. A., Blum, R. W., Foster, D. G., Haddad, L. B., Major, B., Munk-Olsen, T., Polis, C. B., Robinson, G. E., Rocca, C. H., Russo, N. F., Steinberg, J. R., Stewart, D. E., Stotland, N. L., Upadhyay, U. D., & Ditzhuijzen, J. van. (2024). Correcting the scientific record on abortion and mental health outcomes. *BMJ*, 384, e076518. <https://doi.org/10.1136/bmj-2023-076518>
- Lund, B. D., Wang, T., Mannuru, N. R., Nie, B., Shimray, S., & Wang, Z. (2023). ChatGPT and a new academic reality: Artificial Intelligence-written research papers and the ethics of the large language models in scholarly publishing. *Journal of the Association for Information Science and Technology*, 74(5), 570–581. <https://doi.org/10.1002/asi.24750>
- Martín-Martín, A., Orduna-Malea, E., Ayllón, J. M., & Delgado López-Cózar, E. (2016). Back to the past: On the shoulders of an academic search engine giant. *Scientometrics*, 107, 1477–1487. <https://doi.org/10.1007/s11192-016-1917-2>
- Martín-Martín, A., Thelwall, M., Orduna-Malea, E., & Delgado López-Cózar, E. (2021). Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations' COCI: A multidisciplinary comparison of coverage via citations. *Scientometrics*, 126(1), 871–906. <https://doi.org/10.1007/s11192-020-03690-4>
- Simon, F. M., Altay, S., & Mercier, H. (2023). Misinformation reloaded? Fears about the impact of generative AI on misinformation are overblown. *Harvard Kennedy School (HKS) Misinformation Review*, 4(5). <https://doi.org/10.37016/mr-2020-127>
- Skeppstedt, M., Ahltorp, M., Kucher, K., & Lindström, M. (2024). From word clouds to Word Rain: Revisiting the classic word cloud to visualize climate change texts. *Information Visualization*, 23(3), 217–238. <https://doi.org/10.1177/14738716241236188>
- Swedish Research Council. (2017). *Good research practice*. Vetenskapsrådet.
- Stokel-Walker, C. (2024, May 1.). *AI Chatbots Have Thoroughly Infiltrated Scientific Publishing*. Scientific American. <https://www.scientificamerican.com/article/chatbots-have-thoroughly-infiltrated-scientific-publishing/>
- Subbaraman, N. (2024, May 14). Flood of fake science forces multiple journal closures: Wiley to shutter 19 more journals, some tainted by fraud. *The Wall Street Journal*. <https://www.wsj.com/science/academic-studies-research-paper-mills-journals-publishing-f5a3d4bc>
- The pandas development team. (2024). pandas-dev/pandas: Pandas (v2.2.2) [Computer software]. Zenodo. <https://doi.org/10.5281/zenodo.10957263>
- Thorp, H. H. (2023). ChatGPT is fun, but not an author. *Science*, 379(6630), 313–313. <https://doi.org/10.1126/science.adg7879>

Tripodi, F. B., Garcia, L. C., & Marwick, A. E. (2023). 'Do your own research': Affordance activation and disinformation spread. *Information, Communication & Society*, 27(6), 1212–1228.

<https://doi.org/10.1080/1369118X.2023.2245869>

Vikramaditya, N. (2020). Nv7-GitHub/googlesearch [Computer software]. <https://github.com/Nv7-GitHub/googlesearch>

Acknowledgements

The authors wish to thank two anonymous reviewers for their valuable comments on the article manuscript as well as the editorial group of *Harvard Kennedy School (HKS) Misinformation Review* for their thoughtful feedback and input.

Funding

This research has been supported by Mistra, the Swedish Foundation for Strategic Environmental Research, through the research program Mistra Environmental Communication (Haider, Ekström, Rödl) and the Marcus and Amalia Wallenberg Foundation [2020.0004] (Söderström).

Competing interests

The authors declare no competing interests.

Ethics

The research described in this article was carried out under Swedish legislation. According to the relevant EU and Swedish legislation (2003:460) on the ethical review of research involving humans (“Ethical Review Act”), the research reported on here is not subject to authorization by the Swedish Ethical Review Authority (“etikprövningsmyndigheten”) (SRC, 2017).

Copyright

This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided that the original author and source are properly credited.

Data availability

All data needed to replicate this study are available at the Harvard Dataverse:
<https://doi.org/10.7910/DVN/WUVD8X>

Appendix A: Google Scholar queries

Google Scholar queries for retrieving papers with two cases of specific responses from ChatGPT (and similar applications):

Query 1

"as of my last knowledge update"

Query 2

"I don't have access to real-time data"

Query 3

"as of my last knowledge update" AND "I don't have access to real-time data"

Appendix B: Google Scholar Search query script

```
import pandas as pd
import time
from datetime import date
from scholarly import scholarly
from tqdm import tqdm

# Start the timer
start_time = time.time()
today = date.today()

queries = [
    "as of my last knowledge update",
    "I don't have access to real-time data",
    "as of my last knowledge update" AND "I don't have access to real-time data",
]

for idx,query in enumerate(queries):
    print(query[1:-1])
    search_query = scholarly.search_pubs(query)
    #print(next(search_query))
    # List to store paper data
    papers_data = []
    urls = []
    flag = []

    # Loop over the results
    for i in range(250): # set the number of papers to retrieve
        try:
            # Attempt to fetch a paper
            paper = next(search_query)
            papers_data.append(paper['bib'])
            urls.append(paper['pub_url']) # Add the paper's bibliographic info to the list
            flag.append(0)
            time.sleep(1)
        except KeyError as e:
            # Check what key is missing and decide the action
            if 'eprint_url' in paper:
                urls.append(paper['eprint_url'])
            else:
                urls.append('na')
            flag.append(1)
        except StopIteration:
            # If there are no more papers, break
            break
    # Print out the progress along with how much time has passed
```

```
    elapsed_time = time.time() - start_time
    print(f"Fetch paper {i + 1}, elapsed time: {elapsed_time:.2f} seconds")

# Convert the list of paper data into a pandas DataFrame
df = pd.DataFrame(papers_data)
df['pub_url'] = urls
df['query'] = query
# Save the DataFrame to CSV and Excel formats
df.to_csv('data/scholarly_papers_{}.csv'.format(idx), index=False)
df.to_excel('data/scholarly_papers_{}.xlsx'.format(idx), engine='xlsxwriter', index=False)
# Print the total elapsed time
total_time = time.time() - start_time
print(f"Total elapsed time: {total_time:.2f} seconds")
```