publications

MDPI

Viewpoint

When AI Writes the Letters: Recognizing Synthetic Authorship Patterns in Medical Publishing

Elise Lupon * and Grégoire Micicoi

Institut Universitaire Locomoteur et du Sport, Pasteur 2 Hospital, University Côte d'Azur, 06000 Nice, France; micicoi.g@chu-nice.fr
* Correspondence: lupon.e@chu-nice.fr

Abstract

The rapid integration of generative artificial intelligence into scientific publishing is reshaping how academic text can be produced, revised, and scaled. While transparent and limited use of AI for language support may be acceptable, a new structural vulnerability may be emerging in medical publishing: the large-scale production of short, plausible, and weakly individualized correspondence across multiple specialties. In this viewpoint, we describe and conceptualize a pattern that may be termed *synthetic authorship*, defined not as undisclosed AI use alone, but as a reproducible mode of scholarly output structurally facilitated by automation. We focus particularly on letters to the editor, a format that combines brevity, rapid editorial handling, and formal indexation, and may therefore be especially exposed to this phenomenon. Based on recurring patterns observed in PubMed-indexed literature, including unusually high publication velocity, abrupt thematic dispersion, and stylistic uniformity across unrelated domains, we argue that such outputs may challenge the authenticity, epistemic value, and editorial function of scientific correspondence. We do not present empirical proof of misconduct, but rather outline a conceptual framework for understanding this emerging risk and propose proportionate editorial safeguards, including cross-domain pattern detection and contextual assessment of authorship coherence. As AI lowers the threshold for generating domain-plausible commentary at scale, scientific publishing must adapt its integrity frameworks accordingly. In this context, vigilance toward synthetic authorship may become an essential component of editorial responsibility and post-publication quality control.

Keywords: artificial intelligence; ethical responsibility; scientific publishing; research integrity; authorship ethics; synthetic authorship; editorial oversight

Check for updates

Academic Editor: Eungi Kim

Received: 4 January 2026
Revised: 17 March 2026
Accepted: 19 March 2026
Published: 25 March 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Hyperproductivity in scientific publishing is not a new phenomenon. Long before the emergence of large language models, bibliometric studies had already identified subsets of researchers publishing at extraordinary rates across multiple domains, reflecting the structural pressures of metric-driven academic systems rather than a uniform acceleration of scientific creativity (Trimble, 2010; Ioannidis et al., 2018; Abramo & D'Angelo, 2025). In this context, generative AI may represent less a rupture than an accelerant: a tool capable of intensifying pre-existing dynamics of output maximization, especially in formats that are short, rapid, and lightly scrutinized.

Within this broader landscape, one particularly exposed format is the letter to the editor. In recent years, an intriguing pattern has begun to appear within PubMed-indexed literature: a growing number of letters to editors and comments on published articles

are authored by the same small groups of individuals, often from a single institution, and published across a wide range of medical fields, from neurology and cardiology to dermatology, ophthalmology, endocrinology, oncology, surgery, and psychiatry. Many of these papers acknowledge the use of AI assistance only for grammatical correction. In several instances, close to one hundred such publications have appeared within a single year, and in some cases within as little as six months. An investigative report once described a single author publishing over 500 letters to the editor within one year, illustrating how such patterns can escape editorial scrutiny (Ioannidis et al., 2018).

At first glance, these letters seem legitimate: they reference genuine papers, follow accepted academic structure, and appear in reputable journals. Yet their sheer volume, thematic incoherence, and uniform linguistic style raise questions about the process of their creation. They are consistently devoid of personal clinical experience, illustrative cases, or references to the authors' prior work, elements that usually lend authenticity and depth to scientific correspondence. It would be difficult to reconcile such breadth and velocity with conventional expectations of domain-specific expertise, particularly when outputs span dozens of unrelated or loosely related fields within compressed timeframes (Conroy, 2024).

This pattern may signal a new form of automation in scientific publishing, one that could be termed synthetic authorship. By synthetic authorship, we refer not to concealed AI use per se, but to reproducible authorship patterns structurally facilitated by automation. In this model, large language models possess sufficient medical fluency to generate brief, plausible critiques or comments on recently published papers across multiple specialties. When lightly edited and rapidly submitted, such texts may evade detection, particularly in correspondence sections where brevity and rapid review are standard. Rather than a single act of undisclosed assistance, synthetic authorship describes a scalable pattern of output in which automation becomes embedded in the production of scholarly voice.

The rapid maturation of large language models has demonstrated that AI-generated scientific prose can be highly fluent and difficult to distinguish from human-authored text, particularly in short formats (van Dis et al., 2023; Kobak et al., 2025). Editorial analyses have further emphasized the structural implications of generative AI for peer review and scientific publishing ecosystems (Stokel-Walker & Van Noorden, 2023). Within this technological landscape, the threshold for producing domain-plausible commentary across multiple specialties has dramatically lowered.

It is important to clarify that the concern is not the ethical use of AI as a drafting assistant, an increasingly accepted practice under transparent disclosure, but rather the possibility of scaled automation of correspondence outputs. The concern is therefore both structural and epistemic. The epistemic influence of such letters remains difficult to quantify empirically; however, their cumulative presence may reshape visibility dynamics within indexing systems even if their direct citation impact remains modest.

The implications extend beyond mere inconvenience or metric inflation. Letters to the editor occupy a distinctive place in scientific communication: they contribute to post-publication dialogue, allow readers to challenge interpretations, raise methodological concerns, refine conclusions, and sustain the critical exchange through which scientific knowledge remains self-correcting. When correspondence becomes saturated with potentially synthetic or low-substance contributions, the issue is not only administrative but scientific. The signal-to-noise ratio of scholarly dialogue may deteriorate, meaningful critique may be diluted by formulaic commentary, and the credibility of post-publication exchange may be weakened.

This is precisely why letters to the editor deserve specific attention. Their structural characteristics, namely brevity, rapid editorial turnover, often limited peer review, and formal indexation, may make them particularly susceptible to automation. These features

are not problematic in themselves and indeed support timely scientific discussion. However, in the current technological environment, they may also lower the practical barriers to producing high volumes of plausible but low-authenticity commentary at scale.

Previous investigations into integrity vulnerabilities in scientific publishing have shown how certain formats and editorial workflows may be more susceptible to high-throughput practices, particularly when rapid turnaround and low word counts are prioritized (Siegel et al., 2018; Kozlov, 2024). Correspondence sections, by design, combine brevity, speed, and limited methodological scrutiny, potentially creating structural exposure to automation at scale.

Recognizing and addressing this pattern requires both ethical awareness and procedural innovation. Journals could implement cross-domain pattern detection to flag sudden surges of multi-specialty output from the same authors. Editorial boards might systematically review coherence between authors' expertise, institutional affiliation, and thematic scope, particularly when publication profiles appear inconsistent or implausibly broad. Such steps would mirror emerging integrity checks in other domains of scientific publishing, but with a focus on authorship authenticity rather than image manipulation or plagiarism.

Concrete editorial "triggers" could include: (i) unusually high publication velocity within short timeframes; (ii) disproportionate reliance on specific short-form formats such as letters or comments; (iii) abrupt thematic dispersion across unrelated specialties; and (iv) stylistic or structural uniformity across otherwise domain-specific texts. Importantly, these indicators should not be interpreted as proof of misconduct, nor should they be applied in a rigid or punitive manner. Rather, they may serve as contextual safeguards prompting proportionate editorial assessment when several unusual patterns occur together. Our intention is not to penalize legitimate interdisciplinarity or productivity, especially in original research, but to encourage calibrated vigilance in situations where extreme output, cross-specialty dispersion, and repeated short-format publication converge.

This phenomenon illustrates a central paradox of the AI era: the same technology that accelerates discovery can also generate convincing noise at unprecedented scale (Gao & Harrison, 2025). The boundary between human and synthetic scientific voice is no longer defined by intent to deceive, but by automation's quiet efficiency. As medicine embraces AI for diagnosis and therapy, scientific publishing must adapt with equal vigilance. Early recognition of repetitive phrasing, exaggerated thematic drift, and stylistic uniformity may become as essential to editorial hygiene as rigor is to clinical practice (Prillaman, 2023; Hosseini & Resnik, 2025). Provenance, transparency, and authenticity will define the next era of academic integrity. The first step, as in any clinical scenario, is awareness.

The present reflection does not seek to empirically quantify the scope of the phenomenon but rather to delimit and conceptualize a potential structural vulnerability in editorial ecosystems. The patterns discussed here are based on observations within PubMed-indexed literature and are intended as a conceptual reflection on structural vulnerability rather than as empirical evidence of misconduct in specific cases. Future work employing bibliometric and scientometric methodologies could help determine whether these patterns represent isolated anomalies or emerging systemic trends. In the interim, calibrated vigilance, grounded in proportionality rather than alarmism, appears warranted.

Author Contributions: Conceptualization, E.L.; validation, G.M.; formal analysis, E.L. and G.M.; investigation, E.L.; resources, E.L.; data curation, E.L.; writing—original draft preparation, E.L.; writing—review and editing, E.L. and G.M.; visualization, E.L.; supervision, G.M.; project administration, E.L. and G.M.; funding acquisition, E.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No new data were created.

Acknowledgments: No generative AI was used in the writing of this article beyond minor grammar correction.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial intelligence
PubMed	Public Medline

References

Abramo, G., & D'Angelo, C. A. (2025). Hyperprolific authorship: Unveiling the extent of extreme publishing in the 'publish or perish' era. *Journal of Informetrics*, *19*(2), 101658. [CrossRef]

Conroy, G. (2024). Surge in number of 'extremely productive' authors concerns scientists. *Nature*, *625*(7993), 14–15. [CrossRef] [PubMed]

Gao, J., & Harrison, E. M. (2025). The promise and perils of autonomous AI in science. *NEJM AI*, *2*(1), AIe2401073. [CrossRef]

Hosseini, M., & Resnik, D. B. (2025). Guidance needed for using artificial intelligence to screen journal submissions for misconduct. *Research Ethics*, *21*, 1–8. [CrossRef] [PubMed]

Ioannidis, J. P. A., Klavans, R., & Boyack, K. W. (2018). Thousands of scientists publish a paper every five days. *Nature*, *561*(7722), 167–169. [CrossRef] [PubMed]

Kobak, D., González-Márquez, R., Horvát, E., & Lause, J. (2025). Delving into LLM-assisted writing in biomedical publications through excess vocabulary. *Science Advances*, *11*(27), eadt3813. [CrossRef] [PubMed]

Kozlov, M. (2024). 'Publish or perish' is now a card game—Not just an academic's life. *Nature*, *632*(8025), 483. [CrossRef] [PubMed]

Prillaman, M. (2023). 'ChatGPT detector' catches AI-generated papers with unprecedented accuracy. *Nature*. [CrossRef] [PubMed]

Siegel, M. G., Brand, J. C., Rossi, M. J., & Lubowitz, J. H. (2018). "Publish or perish" Promotes medical literature quantity over quality. *Arthroscopy*, *34*(11), 2941–2942. [CrossRef] [PubMed]

Stokel-Walker, C., & Van Noorden, R. (2023). What ChatGPT and generative AI mean for science. *Nature*, *614*(7947), 214–216. [CrossRef] [PubMed]

Trimble, S. W. (2010). Reward quality not quantity. *Nature*, *467*(7317), 789. [CrossRef] [PubMed]

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. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.