Artificial intelligence

VitaHash. Vol 1. No 1. (2026)
DOI https://doi.org/10.68919/vitahash.2026.00003
On-chain: STAMP-2026-0910-5NCUY2J7

Editorial

Toward Transparent AI Disclosure in Medical Publishing: A Proposed Clinical AI Score Framework

Chayakrit Krittanawong, MD, FACC, FAHA, FSCAI

AFFILIATIONS

HumanX, Delaware, DE, United States

Artificial intelligence (AI) is rapidly transforming scientific publishing, yet current disclosure practices provide limited information about the extent of AI involvement. The proposed Clinical AI Score (CAIS) is a standardized framework for quantifying AI assistance across manuscript preparation, analysis, and visualization. CAIS is intended to promote transparency and human accountability rather than restrict AI use. A preliminary threshold of <10 is proposed to represent limited-to-moderate AI involvement, pending future validation. Transparent AI disclosure, combined with digital provenance and human oversight, may provide a practical framework for maintaining scientific integrity in the evolving era of AI-assisted medical publishing.

KEYWORDS: Artificial intelligence, Clinical AI Score, CAIS, Scientific publishing, AI disclosure, Medical Publication

For decades, scientific manuscripts were predominantly conceived, written, and edited by human authors without the assistance of artificial intelligence (AI). However, generative AI and large language models (LLMs) have rapidly advanced and entered medical research and scientific communication. AI-assisted writing, literature synthesis, systematic reviews, meta-analyses, data analysis, figure creation, and other forms of research support are likely to remain integral components of academic publishing. The central challenge is no longer whether AI will be used, but rather how its use can be transparently disclosed, ethically governed, and appropriately integrated into scholarly work. Recent studies suggest that AI-assisted writing may already be common in biomedical publications, highlighting the urgent need for practical frameworks for AI disclosure and accountability.[1,2]

As a clinician and researcher working across multiple roles, including author, reviewer, and editor, as well as a mentor to medical students, residents, and fellows, I have observed increasing use of AI in academic work. In my own mentorship, I have encouraged trainees to fully disclose their use of AI

© 2026 VitaHash LLC. All rights reserved.

in manuscript preparation. Although these observations are anecdotal, they are consistent with the broader emergence of AI-assisted scientific writing and research support across multiple stages of manuscript preparation.

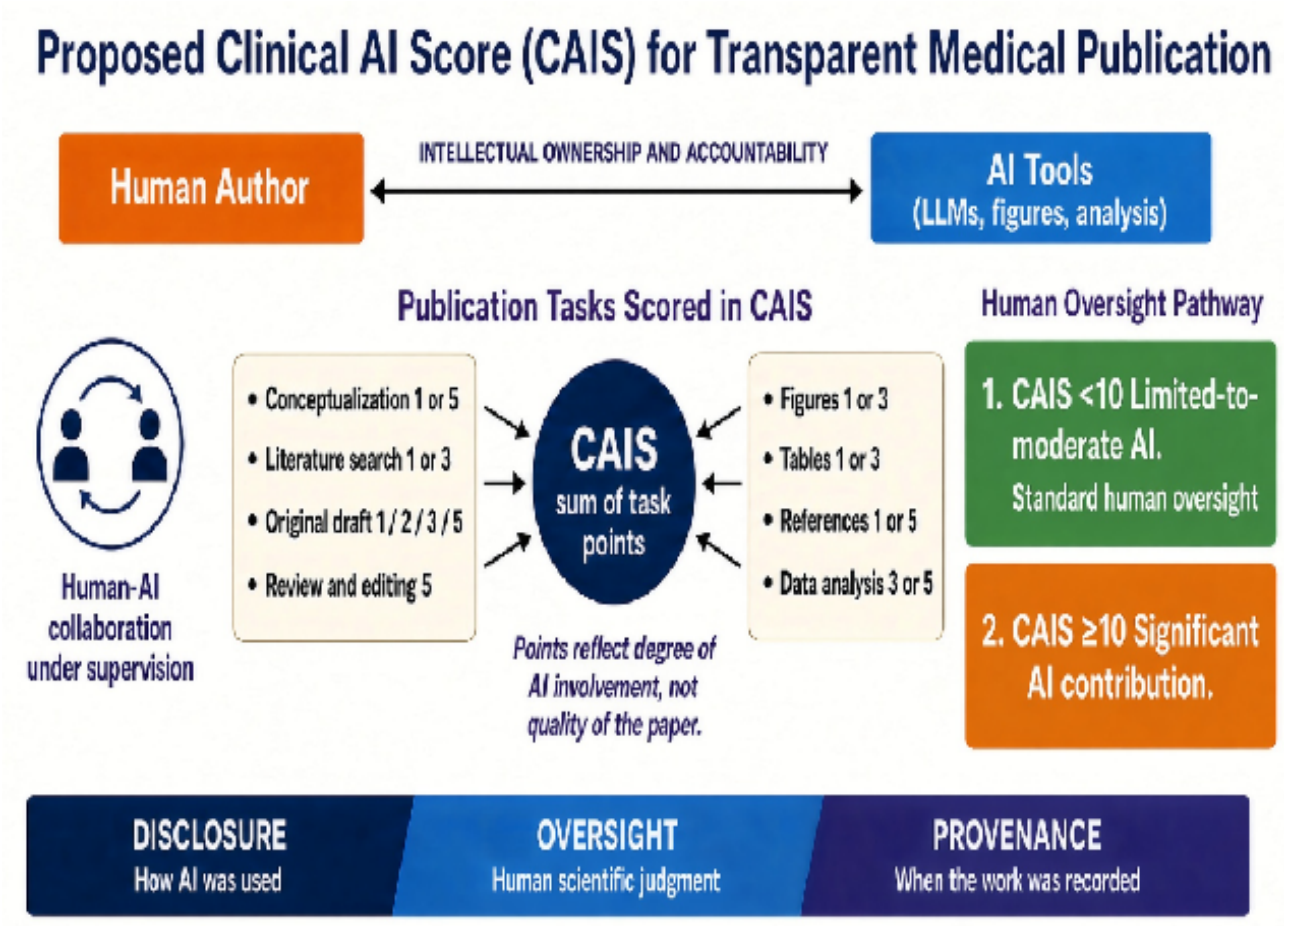
This changing landscape highlights the need for technological approaches that enable transparent documentation of AI involvement, including emerging platforms such as VitaHash.org that use blockchain-based verification methods.3 Standardized AI disclosure should become an important component of scientific publishing, allowing readers, editors, and reviewers to understand the extent and nature of AI contributions.

Recently, updated guidance from JAMA has addressed the appropriate use of AI by authors in scientific publications.4 While these recommendations provide important principles, additional quantitative methods may be needed to describe the degree and type of AI task involvement. Therefore, I propose the Clinical AI Score (CAIS) for Medical Publication, a standardized framework designed to quantify AI assistance throughout the publication process. (Figure)

The Clinical AI Score is not intended to function as a clinical prediction model or as a mechanism to prohibit AI use. Rather, it is a publication transparency framework designed to communicate the extent of AI contribution and promote responsible human oversight.

A high-quality scientific publication should not necessarily be defined as one created with zero AI involvement. The future of scientific publishing may involve effective collaboration between human expertise and AI-assisted capabilities when appropriately implemented. The critical distinction is not whether AI was used, but whether human authors maintain intellectual ownership, scientific judgment, accountability, and appropriate oversight.

I propose that a preliminary Clinical AI Score threshold of <10 may represent limited-to-moderate AI involvement with appropriate human oversight. Manuscripts exceeding this threshold may warrant enhanced human verification and revision. This threshold is not intended to represent a definitive regulatory boundary but rather an initial hypothesis requiring future validation.



AI-Assisted Medical Illustration and Visualization

AI-generated images, such as figures and graphical abstracts, represent an area of particular interest. Modern AI tools can produce sophisticated visual content, and AI-assisted medical illustration under expert human supervision may improve efficiency, accessibility, and cost-effectiveness. Historically, some researchers have already relied on both commercial and noncommercial digital illustration software and computer-assisted design tools for scientific visualization.

However, highly detailed and complex anatomical illustrations remain challenging. Current evidence suggests that AI-generated medical illustrations may contain anatomical inaccuracies and cannot yet consistently replace expert medical illustrators.5 Therefore, AI should be viewed as a supportive tool rather than a replacement for specialized human expertise.

When AI improves efficiency while maintaining human oversight, the time saved may allow researchers and clinicians to focus on higher-value activities, including hypothesis development, clinical interpretation, statistical analysis, and scientific reasoning.

The Challenge of AI Detection and Disclosure

Attempts to prohibit AI use in scientific publishing are unlikely to be successful. Even with restrictive policies, AI-assisted manuscripts may remain undisclosed. Furthermore, current AI detection methods remain imperfect and should not replace transparent disclosure practices. Although commercial AI detection systems have self-reported high accuracy, independent evaluations have identified false-positive results in certain contexts.6 Even a small false-positive rate may have significant consequences when evaluating authorship integrity or potential academic misconduct.

In my observation, text may appear AI-generated despite being written primarily by humans, while AI-generated or heavily AI-assisted text may appear human-like. Most importantly, post hoc AI detection attempts to infer the origin of a final text without access to the process by which that text was produced. The final output alone, without audit logs, may therefore be insufficient to reliably distinguish between human-authored, AI-assisted, and AI-generated content.

Future advances in AI detection may improve performance; however, detecting the provenance of a final text without audit logs remains conceptually different from documenting the process by which it was created. Transparent disclosure therefore provides information that post hoc detection alone cannot reliably provide.

Readers should be able to understand not only whether AI was used, but also the degree, purpose, and specific areas of AI contribution.

In fact, emerging evidence suggests that human-AI collaboration may outperform either humans or AI alone in certain tasks.7,8 The future of scientific publishing should focus on responsible collaboration rather than artificial separation between human and AI contributions.

AI-Assisted Plagiarism and Digital Provenance

AI introduces an additional challenge related to plagiarism and content provenance. LLMs may generate text or figures that resemble or incorporate information from existing publicly available sources, and it may be difficult for authors, reviewers, or editors to determine the origin of every generated passage. Conventional plagiarism-detection systems are primarily designed to identify textual similarity and may not reliably establish whether AI-generated content has incorporated material from multiple sources in ways that are difficult to detect.

This creates an important distinction between content similarity detection and digital provenance. Plagiarism detection attempts to identify whether content resembles previously published material, whereas digital provenance aims to establish when a particular version of a work was created, submitted, or registered and by whom.

A potential solution is the use of timestamped, on-chain digital publication. Blockchain-based

systems may provide an immutable record of the existence of a particular digital artifact at a specific point in time. Such records could potentially complement traditional publication and plagiarism-detection systems by establishing a verifiable history of content creation and submission. VitaHash.org is an emerging platform exploring this approach through on-chain documentation of publication and digital ownership.3

Importantly, a blockchain timestamp should not be interpreted as definitive proof of authorship, originality, or the absence of plagiarism. Rather, it may provide an additional layer of digital provenance that can help establish the existence and timing of a particular version of a work.

In an AI-enabled publishing environment, AI disclosure and digital provenance may therefore represent complementary components of publication integrity: disclosure describes how AI was involved, whereas provenance can help document when and how a digital work entered the scholarly record.

Preliminary Interpretation of the Clinical AI Score

The proposed scoring system is intended as a practical framework rather than a validated regulatory instrument.

Clinical AI Score <10:
Limited-to-moderate AI involvement with appropriate human oversight.

Clinical AI Score ≥10:
Increased AI contribution that may warrant enhanced human verification, revision, and disclosure.

These thresholds should be considered preliminary and require future validation through expert consensus, reliability testing, and evaluation in real-world publications.

TABLE 1. Proposed Clinical AI Score for Quantifying AI Contribution in Medical Publication

Task	AI Involvement	Score	Comment
Conceptualization	AI-assisted idea development	1	Human authors developed the final research question
	AI-generated concept or hypothesis	5	AI substantially generated the research concept
Literature search	AI-assisted search of selected literature	1	Human-directed literature review
	AI performed broad literature identification	3	Requires human verification
Writing – Original Draft	<30% AI contribution	1	Limited writing assistance
	30–50% AI contribution	2	Moderate writing assistance
	>50% AI contribution	3	Extensive AI drafting
	>70% AI contribution	5	Predominantly AI-generated draft
Writing – Review and Editing	AI performed critical review, restructuring, or major revision	5	Significant intellectual editing assistance
Figures	AI-assisted figure creation	1	Human-designed with AI support
	AI-created figure with human supervision	3	Requires expert verification
Tables	AI-assisted table generation	1	Human-designed content
	AI-created table	3	Requires validation
References	AI-assisted reference identification	1	Human verification required

Future Validation of the Clinical AI Score

The proposed scoring system can be further refined using established methods for developing clinical and quality-assessment scores.9,10 Future work could evaluate whether the individual domains, scoring levels, and proposed threshold appropriately reflect the degree of AI involvement in medical publication.

Future validation could include expert Delphi consensus to refine the domains and scoring definitions, analytic hierarchy process (AHP)-based weighting to establish the relative importance of different publication tasks, interobserver reliability testing to determine whether independent reviewers consistently assign similar scores, and correlation with independent editorial assessments of AI involvement.

The preliminary point assignments and threshold of 10 should therefore be considered hypotheses rather than empirically established standards. The initial Clinical AI Score is intended as a practical disclosure instrument that can be systematically refined and validated through future research.

Role of Human Oversight

AI should be considered an assistive technology rather than a replacement for scientific judgment. If AI improves efficiency in manuscript preparation, visualization, organization, or analysis under appropriate supervision, the time saved may allow researchers to focus on higher-value activities, including clinical reasoning, study design, interpretation of findings, and advancement of medical knowledge.

Human authors must remain responsible for the accuracy, originality, interpretation, and integrity of the final publication. AI systems cannot assume authorship responsibility or scientific accountability, nor can they assume responsibility for errors.

The goal of the Clinical AI Score is not to determine whether a publication is "human" or "AI-generated." Instead, it provides a transparent framework to communicate the degree of AI involvement while maintaining the fundamental principles of scientific integrity and author accountability.

Preparing for the Future of AI in Publishing

AI technologies will continue to evolve and become increasingly capable. The goal of medical publishing should not be to prevent AI use but to establish appropriate AI disclosure systems and responsible standards that protect scientific integrity.

The Clinical AI Score provides a potential framework for transparency by allowing readers, reviewers, and editors to understand the proportion, purpose, and type of AI involvement in a publication.

Medical publication integrity must continue to evolve in the AI era. The future standard of scientific publishing should not be determined by whether AI was used, but by whether AI involvement was transparent, appropriately supervised, and consistent with human accountability.

Ultimately, the question should not be whether a manuscript was written by a human or assisted by AI, but whether the scientific work remains transparent, verifiable, and accountable to human authors.

Conflict of Interest: Dr. Krittanawong is the founder of VitaHash.org.

AI Disclosure: CK wrote and critically edited the article, including its ideas, analysis, and conclusions. An AI model was used for grammar and spelling edits (Clinical AI Score: 1 point). The central illustration was generated using AI based on CK's concepts and subsequently refined through CK's inputs and review (Clinical AI Score: 3 points). The total Clinical AI Score for this manuscript is 4 points.

References

1. Yoo JH. Defining the Boundaries of AI Use in Scientific Writing: A Comparative Review of Editorial Policies. J Korean Med Sci. 2025 Jun 16;40(23):e187. doi: 10.3346/jkms.2025.40.e187. PMID: 40524628; PMCID: PMC12170296.
2. Holzwarth L, González-Márquez R, Kobak

D. Most Biomedical Publications Show Signs of LLM-Assisted Writing. arXiv. 2026. doi: 10.48550/arXiv.2608.10715.

3. Krittanawong C. Why I built VitaHash. VitaHash. 2026. STAMP-2026-0907-IWR2DGI4.

4. Flanagin A, Perlis RH, Bibbins-Domingo K. Updated Guidance for Author Use of AI in Medical Publication. JAMA. 2026 Aug 10. doi: 10.1001/jama.2026.16613. Epub ahead of print. PMID: 42574208.

5. Eldesoqui M, Albadawi EA, AlQumaizi KI, Radwan MNM, Ebrahim HA, Elsaid MAE. Assessing the Anatomical Accuracy of AI-Generated Medical Illustrations: A Comparative Study of Text-to-Image Generator Tools in Anatomy Education. Clin Anat. 2025 Sep;38(6):712-717. doi: 10.1002/ca.70002. Epub 2025 Jul 9. PMID: 40635255.

6. Wong M. America Has a Pangram Problem. The Atlantic. May 30, 2026.

7. Vaccaro M, Almaatouq A, Malone T. When combinations of humans and AI are useful: A systematic review and meta-analysis. Nat Hum Behav. 2024 Dec;8(12):2293-2303. doi: 10.1038/s41562-024-02024-1. Epub 2024 Oct 28. PMID: 39468277; PMCID: PMC11659167.

8. Sears S, Weisberg DS. Bot or not: Can people tell the difference between stories written by a human or by an AI system? Judgment and Decision Making. 2026;21:e21. doi: 10.1017/jdm.2026.10042.

9. Baltussen R, Niessen L. Priority setting of health interventions: the need for multi-criteria decision analysis. Cost Eff Resour Alloc. 2006 Aug 21;4:14. doi: 10.1186/1478-7547-4-14. PMID: 16923181; PMCID: PMC1560167.

10. Chakraborty S, Raut RD, Rofin TM, Chakraborty S. A Comprehensive and Systematic Review of Multi-Criteria Decision-Making Methods and Applications in Healthcare. Healthcare Analytics. 2023;4:100232. doi: 10.1016/j.health.2023.100232.