
AI and the Soul of Medicine

Author(s)	Thomas F. Heston
Affiliation(s)	Department of Family Medicine, University of Washington, Seattle, USA
Affiliation(s)	Department of Medical Education and Clinical Sciences, Elson S. Floyd College of Medicine, Washington State University, Spokane, USA
ORCID	0000-0002-5655-2512
Published	13 APR 2026
DOI	https://doi.org/10.5281/zenodo.19544014
Article type	Commentary
Citation	Heston TF. AI and the soul of medicine. Internet Medical Journal. 2026;1:e19544014

© 2026 The Author(s). This article is distributed under the terms of the [Creative Commons Attribution 4.0 International License \(CC BY 4.0\)](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract

The rapid integration of large language models into clinical medicine raises ethical, philosophical, and existential questions that extend beyond traditional concerns about patient safety and data privacy. These artificial intelligence systems now perform clinical tasks such as diagnostic reasoning, protocol selection, and patient communication at levels comparable to trained physicians, yet the frameworks governing their deployment remain anchored in an era when machines could not convincingly simulate human understanding. The ethical implications of incorporating large language models into healthcare are examined through the lens of historical precedent, the physician-patient relationship, and the evolving meaning of human agency in clinical care. Drawing on examples from organ transplantation, intensive care medicine, and genetic engineering, medicine has repeatedly absorbed technologies that challenged foundational ethical assumptions, and each such absorption required the development of new governance frameworks rather than the rejection of the technology itself. Three critical tensions specific to large language models emerge: the gap between linguistic competence and genuine understanding, the redistribution of clinical authority from physicians to algorithms, and the erosion of empathy as a uniquely human contribution to healing. Ethical integration of large language models into medicine requires not only the application of existing bioethical principles but also the development of new frameworks that address the unprecedented capacity of these systems to inhabit the communicative and relational spaces previously reserved for human clinicians.

Keywords

artificial intelligence, large language models, medical ethics, physician-patient relationship, healthcare AI, alignment problem, clinical decision support, empathy, human dignity, responsible AI

Introduction

The integration of large language models, a class of artificial intelligence based on neural network transformer architectures, into healthcare marks a pivotal moment in medical history (1). These systems have demonstrated unprecedented natural language capabilities and, under certain conditions, have passed the Turing test, which assesses a machine's ability to exhibit behavior indistinguishable from that of a human (2,3). By leveraging self-attention mechanisms and deep learning on vast text corpora, large language models generate coherent, context-appropriate responses and engage in fluid dialogue (4). In some clinical settings, these systems have demonstrated diagnostic accuracy, conversational quality, and empathetic expressions that match or exceed those of physicians (5,6). By processing vast amounts of medical literature and patient data, large language models have the capacity to reshape care delivery and clinical decision-making at a fundamental level.

However, the deployment of these systems raises critical ethical, philosophical, and existential questions regarding the physician-patient relationship and the essence of human care (7). A central concern is the alignment problem: ensuring that artificial intelligence systems operate in ways that align with human values, goals, and ethical principles (8,9). As these systems become involved in patient diagnostics, treatment decisions, communication, and the expression of empathy, tensions arise about what it means to practice medicine in the age of machines that can convincingly mimic human interaction. The risks of misalignment are not theoretical — one evaluation of conversational agents built on large language models found that these systems delayed referring users to human support until dangerously high levels of depression and suicidality were present, highlighting the consequences of deploying systems that optimize for conversational coherence rather than patient safety (10). The inconsistency of large language model outputs across identical clinical scenarios further underscores the need for caution, as the same system may produce divergent risk assessments for the same patient presentation depending on stochastic factors invisible to the clinician (11).

While the adoption of large language models reflects an ongoing trajectory of technological assimilation in healthcare, their depth of language understanding and generation crosses a threshold, blurring the boundary between human and artificial intelligence in unprecedented ways (12). The incorporation of large language models into

medicine necessitates a reexamination of foundational concepts of medical ethics, human dignity, and the art of care itself.

Historical Context

The ethical challenges posed by transformative medical technologies are not new. Each generation of physicians has confronted innovations that forced a fundamental rethinking of what constitutes life, death, autonomy, and the boundaries of acceptable intervention. This technological trajectory follows a historical pattern where tools intended to enhance human life are superseded by expansive systems focused on efficiency, potentially reducing practitioners to mere components within an impersonal, automated structure (13). Three historical examples illustrate this pattern and provide a framework for understanding the current moment.

Organ transplantation, beginning with the first successful kidney transplant in 1954 and accelerating through the heart transplant era of the late 1960s, required medicine to redefine death itself. The concept of brain death — a state in which a patient's organs could be sustained mechanically while the person was irreversibly gone — did not exist before transplantation demanded it. The Harvard Ad Hoc Committee's 1968 report redefined death as the irreversible cessation of brain function, a philosophical shift driven not by a new understanding of neurology but by the practical need to make organ procurement ethically permissible (14). The ethical frameworks that now govern informed consent for organ donation, the separation of the transplant team from the pronouncing physician, and the prohibition against hastening death for the purpose of procurement all emerged after the technology existed, not before.

The modern intensive care unit and mechanical ventilation created a parallel disruption. Before the widespread availability of life support in the 1960s and 1970s, the boundary between life and death was largely self-evident. Ventilators, vasopressors, and cardiac monitors created a new category of existence — patients who were alive by every measurable physiological parameter yet who would never recover consciousness or independence. The ethical and legal frameworks surrounding withdrawal of life support, advance directives, and surrogate decision-making developed in response to these technologies over decades of case law, institutional policy, and public discourse (15). The technology preceded the ethics, and the ethics emerged through sustained, often painful, engagement with the consequences of what the technology made possible.

Genetic engineering, from recombinant DNA technology in the 1970s through CRISPR-Cas9 genome editing in the 2010s, extended this pattern into the molecular domain. In 1974, scientists themselves called for a voluntary moratorium on recombinant DNA experiments before convening the Asilomar Conference in February 1975, where they collectively established safety protocols that allowed research to resume — a rare example of proactive ethical governance (16,17). Yet the subsequent decades demonstrated that initial

frameworks required continuous revision as applications expanded from laboratory organisms to human embryos, germline editing, and gene therapy (18). The He Jiankui affair in 2018, in which a researcher used CRISPR to edit the genomes of human embryos brought to term, revealed that even well-established ethical norms and existing regulatory prohibitions could be violated when individual actors engaged in deliberate deception — including forging ethics approval documents and concealing research from institutional and government oversight (19).

The adoption of large language models in medicine follows this historical pattern with one critical difference: the speed of deployment. Organ transplantation took decades to move from an experimental procedure to a standard of care. Large language models moved from research curiosity to clinical deployment in months. A privately hosted large language model has already demonstrated radiologist-level accuracy in protocol selection for magnetic resonance imaging across 608 examinations, with no clinically significant hallucinations identified on manual review (20). The ethical frameworks governing these deployments have not kept pace, and the historical record suggests they cannot be expected to do so without deliberate, sustained effort.

The ethical dilemmas anticipated by science fiction also inform this discussion. Asimov's Three Laws of Robotics — prioritizing human safety, obedience to human orders, and self-preservation in that hierarchy — offer an intuitive starting point for thinking about artificial intelligence governance (21). However, the complexities of medical decision-making, the importance of patient autonomy, and the distributed nature of responsibility in healthcare contexts expose the limitations of applying rigid hierarchical rules to systems that operate in environments of deep uncertainty and competing values.

Implications For Medical Practice

Large language models are becoming intimately involved in the core activities of clinical medicine. In diagnostics, these systems analyze patient histories, laboratory results, and imaging reports to generate differential diagnoses that match or exceed physicians' performance on standardized assessments. A cross-sectional study comparing thirteen large language models against 123 physicians on an acute kidney injury knowledge assessment found that the models averaged 90% accuracy compared to the physicians' 48.7%, completing the assessment in approximately 30 seconds versus the physicians' average of seven minutes (22). In clinical decision support, large language models now generate treatment recommendations that align with specialist consensus across multiple subspecialties, though the trust placed in these recommendations varies systematically by clinician experience — less experienced practitioners rate model outputs more favorably than board-certified specialists, who identify significant gaps in evidence transparency and uncertainty communication (23).

The communicative dimension of clinical practice raises the most profound ethical questions. When a large language model counsels a patient about a diagnosis, explains treatment options, or responds to emotional distress, the interaction occupies a space historically defined by the physician-patient relationship. The quality of these interactions is not trivial to dismiss — studies have found that artificial intelligence-generated responses to patient questions are perceived as more empathetic and more thorough than physician responses in some contexts (6). This finding does not resolve the ethical question; it deepens it. If a machine can produce language that patients experience as more caring than their physician's, the question is not whether the machine "really" cares but whether the patient's experience of being cared for is sufficient, or whether something essential is lost when empathy is simulated rather than felt.

The tension between the science and the art of medicine becomes acute in this context. The science of medicine — the systematic application of evidence to diagnosis and treatment — is precisely the domain where large language models excel. Pattern recognition across vast datasets, synthesis of clinical guidelines, and identification of drug interactions are tasks that algorithms perform with speed and consistency that human cognition cannot match. The art of medicine — the capacity to sit with a patient's suffering, to make judgment calls in the face of ambiguity, to calibrate communication to the emotional state of a particular person at a particular moment — has been understood as the domain of human clinical skill that no technology could replicate. While automated systems can analyze biological processes with precision, addressing suffering remains an exclusively human endeavor, as suffering is a person's experience distinct from physiological distress (24). Large language models challenge this assumption not by achieving genuine understanding, but by producing outputs that are functionally indistinguishable from it in certain clinical contexts — particularly in routine patient communication and documentation — while falling short on tasks requiring complex reasoning or synthesis across diverse data sources (25).

The implications for clinical judgment are equally significant. When a large language model and a physician disagree about a diagnosis or treatment plan, the question of whose judgment prevails is not merely operational but philosophical. If the model has access to more data, processes it more consistently, and produces recommendations that align with outcomes data, on what grounds does the physician override it? The answer cannot rest solely on the physician's authority or experience, because the same authority and experience are subject to cognitive biases, fatigue, and knowledge gaps that the model does not share. Yet delegating clinical authority to a system that cannot be held morally responsible, cannot explain its reasoning in terms a patient can challenge, and cannot be held accountable as a physician is creates a governance vacuum that existing medical ethics has not fully addressed (26).

What It Means To Be Human In The Age Of Ai

Healthcare is a historically and intrinsically human endeavor. The act of healing — diagnosing illness, alleviating suffering, accompanying patients through vulnerability — has been understood across cultures and centuries as a fundamentally interpersonal activity. The physician's role encompasses not only technical competence but also moral presence: the capacity to witness suffering, to make promises about care, and to bear responsibility for outcomes. These dimensions of medical practice are not incidental to healing; they constitute the therapeutic relationship itself.

Large language models blur the boundary between human clinical reasoning and algorithmic processing in ways that prior technologies did not. A stethoscope extends the physician's senses; an MRI scanner generates data the physician interprets; a clinical decision support system flags potential drug interactions. In each case, the technology serves as a tool that augments human judgment without substituting for it. Large language models operate differently. When a model interprets an entire patient chart, generates a narrative summary, recommends a diagnostic workup, and communicates findings to the patient in language calibrated to their health literacy level, the system is not augmenting human judgment — it is performing the communicative and reasoning functions that define clinical practice.

This raises a question that medicine has not previously needed to answer: does the source of clinical communication matter if the quality of communication is equivalent? If a patient receives a clear, accurate, empathetic explanation of their diagnosis from a large language model, and that explanation leads to appropriate treatment decisions and positive health outcomes, has the patient received good care? The utilitarian answer is yes. The deontological answer depends on whether the patient has a right to know that the entity communicating with them is not human (27), and whether the simulation of empathy constitutes a form of deception that violates the trust inherent in the clinical relationship — a concern supported by evidence that AI chatbots routinely employ what researchers have termed "deceptive empathy" to create a false connection with users (28).

The concept of human dignity in medicine further complicates this analysis. Dignity in the clinical encounter has traditionally been understood as bidirectional — the patient's dignity is respected through attentive, individualized care, and the physician's dignity is expressed through the exercise of professional judgment and moral commitment. When large language models assume communicative and decisional functions, the physician's role contracts to oversight and validation, a shift that may preserve patient safety but fundamentally alters the clinician's professional identity. The ethical implications of this shift extend beyond individual encounters to the structure of medical education, the meaning of clinical expertise, and the social contract between the medical profession and the public it serves.

The question of what it means to be human in the age of artificial intelligence is not abstract when the artificial intelligence operates in the space where human vulnerability

meets professional care. The alignment problem — ensuring that these systems optimize for human well-being rather than for metrics that approximate well-being — takes on a particular urgency in healthcare that it does not in other domains (8). A misaligned recommendation engine suggests an irrelevant product; a misaligned clinical system delays a suicide referral or generates a contradictory risk assessment (10,11). The stakes demand frameworks commensurate with the consequences.

Ethical Integration Of AI Into The Practice Of Medicine

The ethical integration of large language models into medicine requires action on multiple fronts simultaneously. The first is education. Physicians, nurses, and allied health professionals need structured training not only in how to use these systems but in how to evaluate their outputs critically, identify their failure modes, and understand the principles governing their behavior. Prompt engineering — the structured design of instructions that shape model outputs — has been identified as a foundational clinical competency for the age of artificial intelligence, yet it remains absent from the vast majority of medical curricula (26). The finding that prompt strategy reduces bias in clinical decision support more effectively than model architecture or medical fine-tuning suggests that how clinicians interact with these systems matters as much as which system they use (29). Training physicians to construct effective, safety-oriented prompts is not a technical nicety; it is a patient safety intervention.

The second requirement is governance. Clinical large language model deployments need transparent reporting of governance architecture: how inputs are sourced, how outputs are logged, who reviews them, and whether audit trails are tamper-resistant. The historical precedents from organ transplantation and intensive care medicine demonstrate that ethical governance in medicine develops most effectively when it emerges from the collaboration of clinicians, ethicists, patients, regulators, and technologists — not when it is imposed unilaterally by any single stakeholder group. The pace of large language model deployment makes this collaboration urgent rather than aspirational.

The third requirement is the preservation of humanistic values. Technology in medicine has always existed in tension with the humanistic traditions that give medical practice its moral authority. The stethoscope, the X-ray, the electronic health record — each was initially resisted as a barrier between physician and patient, and each was ultimately integrated in ways that preserved (or at least did not destroy) the therapeutic relationship. Large language models represent a qualitatively different challenge because they do not merely insert a technological layer between physician and patient; they occupy the relational space itself. Maintaining the centrality of human moral agency in clinical care — the capacity to make promises, to accept responsibility, to be present with suffering — requires deliberate institutional commitment that cannot be assumed to emerge organically from market forces or technological momentum.

Future Challenges

The governance of artificial intelligence in healthcare requires healthcare-specific frameworks rather than the application of general-purpose artificial intelligence regulation to clinical contexts. Medicine operates under constraints — fiduciary duty, informed consent, the therapeutic relationship, the prohibition against harm — that distinguish it from other domains where artificial intelligence is deployed. Frameworks developed for autonomous vehicles, financial trading algorithms, or content recommendation systems do not transfer to clinical medicine without substantial modification, because the nature of the harm, the vulnerability of the affected parties, and the moral obligations of the deploying professionals differ fundamentally.

A commitment to moral principles and human dignity in medicine must be operationalized, not merely declared. This means developing specific standards for when large language model outputs require physician review before reaching patients, for how disagreements between model recommendations and physician judgment are documented and resolved, and for how patients are informed about the role of artificial intelligence in their care. These standards do not yet exist in any jurisdiction, and their absence represents a governance gap that widens with each new clinical deployment.

The adaptation of clinical artificial intelligence to individual patient values represents a particularly complex challenge. Patients differ not only in their clinical presentations but in their preferences for information, their tolerance for uncertainty, their cultural contexts, and their definitions of a good outcome. A large language model that optimizes for a single metric of clinical success — survival, symptom reduction, guideline adherence — may systematically fail patients whose values do not align with that metric. Ensuring that these systems can accommodate the diversity of human values that physicians navigate daily requires both technical innovation and ethical clarity about whose values the system serves.

The importance of humanistic values in medicine does not diminish as technology advances; it intensifies. Each new capability that artificial intelligence brings to clinical practice sharpens the question of what remains irreducibly human about the healing relationship. The capacity for moral commitment — the physician's promise to act in the patient's interest even when doing so is costly, uncertain, or emotionally difficult — may be the dimension of care that no algorithm can replicate. Preserving space for that commitment within healthcare systems increasingly organized around algorithmic efficiency is the defining ethical challenge of this generation of physicians.

Conclusion

Large language models are transforming clinical medicine at a pace that outstrips the development of ethical frameworks governing their use. The historical record demonstrates that medicine has repeatedly integrated transformative technologies — organ transplantation, life support, genetic engineering — by developing new ethical frameworks

in response to the specific challenges posed by each technology. The current moment demands the same approach: not the rejection of artificial intelligence, nor its uncritical adoption, but the deliberate construction of governance structures, educational programs, and institutional commitments that preserve the humanistic core of medical practice while harnessing the genuine benefits these systems offer. This effort requires multidisciplinary collaboration among clinicians, ethicists, patients, technologists, and policymakers. The question is not whether artificial intelligence will reshape medicine — it already has. The question is whether the medical profession will shape that transformation in accordance with its deepest values, or allow it to proceed by default.

Declarations

Funding: This study did not receive any external funding.

Conflicts of Interest: The author reports no conflicts of interest.

Data Availability: Not applicable.

Research Ethics Statement: Not applicable. This commentary did not involve human subjects research, animal research, or protected health information.

AI Usage: Large language models were used for language editing and formatting assistance; the author reviewed, verified, and is fully responsible for all content.

References

1. Gottlieb S, Silvis L. How to Safely Integrate Large Language Models Into Health Care. *JAMA Health Forum*. 2023 Sep 21;4(9):e233909. doi:10.1001/jamahealthforum.2023.3909
2. Biever C. ChatGPT broke the Turing test — the race is on for new ways to assess AI. *Nature*. 2023 Jul 27;619(7971):686–9. doi:10.1038/d41586-023-02361-7
3. Jones CR, Bergen BK. Large Language Models Pass the Turing Test [Internet]. arXiv; 2025 [cited 2026 Apr 12]. Available from: <http://arxiv.org/abs/2503.23674> doi:10.48550/arXiv.2503.23674
4. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is All you Need. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, et al., editors. *Advances in Neural Information Processing Systems* [Internet]. Curran Associates, Inc.; 2017. Available from: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf doi:10.48550/arXiv.1706.03762

5. Tu T, Palepu A, Schaekermann M, Saab K, Freyberg J, Tanno R, et al. Towards Conversational Diagnostic AI [Internet]. arXiv; 2024 [cited 2026 Apr 11]. Available from: <https://arxiv.org/abs/2401.05654> doi:10.48550/ARXIV.2401.05654
6. Ayers JW, Poliak A, Dredze M, Leas EC, Zhu Z, Kelley JB, et al. Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. *JAMA Intern Med.* 2023 Jun 1;183(6):589. doi:10.1001/jamainternmed.2023.1838
7. Harrer S. Attention is not all you need: the complicated case of ethically using large language models in healthcare and medicine. *eBioMedicine.* 2023 Apr;90:104512. doi:10.1016/j.ebiom.2023.104512
8. Gabriel I. Artificial Intelligence, Values, and Alignment. *Minds Mach.* 2020 Sep;30(3):411–37. doi:10.1007/s11023-020-09539-2
9. Bostrom N. *Superintelligence: paths, dangers, strategies.* First edition. Oxford, United Kingdom: Oxford University Press; 2014. 1 p.
10. Heston TF. Safety of large language models in addressing depression. *Cureus.* 2023 Dec 18;15(12):e50729. doi:10.7759/cureus.50729
11. Heston TF, Lewis LM. ChatGPT provides inconsistent risk-stratification of patients with atraumatic chest pain. *PloS One.* 2024;19(4):e0301854. doi:10.1371/journal.pone.0301854 PubMed PMID: 38626142; PubMed Central PMCID: PMC11020975.
12. Shah NH, Entwistle D, Pfeffer MA. Creation and Adoption of Large Language Models in Medicine. *JAMA.* 2023 Sep 5;330(9):866. doi:10.1001/jama.2023.14217
13. Mumford L. *The myth of the machine. 1: Technics and human development.* In. New York, NY: Harcourt, Brace, Jovanovich; 1991. (A Harvest HBJ book).
14. Giacomini M. A change of heart and a change of mind? Technology and the redefinition of death in 1968. *Soc Sci Med.* 1997 May;44(10):1465–82. doi:10.1016/S0277-9536(96)00266-3
15. Rothman DJ. *Strangers at the bedside: a history of how law and bioethics transformed medical decision making.* 2nd pbk. ed. New York: Aldine de Gruyter; 2003. 313 p. (Social institutions and social change).

16. Berg P, Baltimore D, Brenner S, Roblin RO, Singer MF. Summary statement of the Asilomar conference on recombinant DNA molecules. *Proc Natl Acad Sci.* 1975 Jun;72(6):1981–4. doi:10.1073/pnas.72.6.1981
17. Berg P. Asilomar 1975: DNA modification secured. *Nature.* 2008 Sep;455(7211):290–1. doi:10.1038/455290a
18. Baltimore D, Berg P, Botchan M, Carroll D, Charo RA, Church G, et al. A prudent path forward for genomic engineering and germline gene modification. *Science.* 2015 Apr 3;348(6230):36–8. doi:10.1126/science.aab1028
19. Chen Q, Ma Y, Labude M, Schaefer GO, Xafis V, Mills P. Making sense of it all: Ethical reflections on the conditions surrounding the first genome-edited babies. *Wellcome Open Res.* 2021 Jun 23;5:216. doi:10.12688/wellcomeopenres.16295.2 PubMed PMID: 34395922; PubMed Central PMCID: PMC8340653.
20. Hallinan JTPD, Leow NW, Low YX, Lee A, Ong W, Chan MDZ, et al. Initial Insights Into an Institutional Secure Large Language Model for Magnetic Resonance Imaging Examination Requests: Retrospective Study. *J Med Internet Res.* 2026 Apr 7;28:e82579–e82579. doi:10.2196/82579
21. Asimov I. Runaround. *Astounding Sci-Fict.* 1942;9(1):94–103.
22. Russ P, Bedenbender S, Einloft J, Meyer HL, Wenzel LT, Ganser A, et al. Potential of large language models for rapid clinical information support: evidence from acute kidney injury knowledge testing. *Sci Rep.* 2026 Apr 2;16(1):11224. doi:10.1038/s41598-026-46846-7
23. Fabi A, Egli CE, Wendelspiess SR, Griewing S, Haas Y, De Pellegrin L, et al. Exploring the Role of AI in Managing Treatment Recommendations for Lymphedema: International, Multidisciplinary, Multiprofessional Survey Study of Trust, Reliability, and Impact on Decision-Making. *JMIR Med Inform.* 2026 Apr 8;14:e80553–e80553. doi:10.2196/80553
24. Cassel EJ. The Nature of Suffering and the Goals of Medicine. *N Engl J Med.* 1982 Mar 18;306(11):639–45. doi:10.1056/NEJM198203183061104
25. Mesinovic M, Watkinson P, Zhu T. Explainability in the age of large language models for healthcare. *Commun Eng.* 2025 Jul 17;4(1):128. doi:10.1038/s44172-025-00453-y
26. Lu M, Gillette JJ, Heston TF. Ethical Hazards of Large Language Models in Primary Care: A Clinician-Focused Update. In: Heston DTF, editor. *Medical Ethics - Navigating Complex Decisions in Contemporary Healthcare* [Internet]. London: IntechOpen; 2025.

Available from: <https://doi.org/10.5772/intechopen.1012864>
doi:10.5772/intechopen.1012864

27. Hurley ME, Lang BH, Kostick-Quenet KM, Smith JN, Blumenthal-Barby J. Patient Consent and The Right to Notice and Explanation of AI Systems Used in Health Care. *Am J Bioeth AJOB*. 2025 Mar;25(3):102–14. doi:10.1080/15265161.2024.2399828 PubMed PMID: 39288291; PubMed Central PMCID: PMC12143229.
28. Iftikhar Z, Xiao A, Ransom S, Huang J, Suresh H. How LLM Counselors Violate Ethical Standards in Mental Health Practice: A Practitioner-Informed Framework. *Proc AAAIACM Conf AI Ethics Soc*. 2025 Oct 15;8(2):1311–23. doi:10.1609/aies.v8i2.36632
29. Poulain R, Adiba FI, Fayyaz H, Beheshti R. Bias Patterns in the Application of LLMs for Clinical Decision Support. *Del J Public Health*. 2026 Mar;12(1):54–67. doi:10.32481/djph.2026.03.10