Cureus
Part of SPRINGER NATURE

Open Access Editorial

The Healer's Blind Spot: Hubris in the Age of Artificial Intelligence

Aarti Sarwal [1], Gopal K. Chopra [2], Ann Hui Ching [3], Matthew Crowson [4], Leo Anthony Celi [5]

1. Neurology and Neurocritical Care, Virginia Commonwealth University Medical Center, Richmond, USA 2. Neurosurgery, Validare Capital LLC/TCG LLC, Boston, USA 3. Emergency Medicine, Singapore General Hospital, Singapore, SGP 4. Otolaryngology, Mass General Brigham, Boston, USA 5. Medicine, Beth Israel Deaconess Medical Center, Boston, USA

Corresponding author: Aarti Sarwal, aartisarwal@gmail.com

Review began 09/03/2026
Review ended 11/03/2026
Published 11/07/2026

© Copyright 2026
Sarwal et al. This is an open access article distributed under the terms of the Creative Commons Attribution License CC-BY 4.0., which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

DOI: 10.7759/cureus.119265

Abstract

Medicine stands at its largest technological inflection point in a century, yet the profession remains culturally unprepared to absorb artificial intelligence (AI). We argue that the central obstacle is not technical but dispositional: hubris, the drift of healthy clinical confidence into arrogance and resistance to feedback. Drawing on the psychology of power, the cognitive science of diagnostic error, and a 2,500-year history from the Hippocratic Corpus through Galen, Ibn Rushd, Semmelweis, and the amyloid hypothesis, we trace how individual bias and collective thought have repeatedly delayed beneficial change at the cost of preventable harm. The exposure this time differs in kind: algorithmic judgment is opaque to most of the clinicians who must act on it, and responsibility for it is shared so widely that no single party owns the result. We characterize forms of AI hubris (technological, regulatory, institutional, clinical, replacement, and ethical) and show how each re-enacts an older pattern in which confidence outruns calibration. We propose a corrective agenda: selecting and training for intellectual humility, retiring pedagogies of dominance, and treating AI as a hypothesis requiring prospective, locally validated evidence rather than a product to be trusted on clearance alone. Concretely, this means weighing intellectual humility in medical school admissions, retiring the pimping ritual from bedside teaching, and giving Patient and Family Advisory Councils explicit authority over AI procurement and deployment. Epistemic humility, we conclude, is the precondition for the responsible adoption of intelligent medicine.

Categories: Health Policy, Medical Education, Quality Improvement
Keywords: artificial intelligence, clinical decision-making, epistemic humility, history of medicine, hubris, medical culture, medical education, medical ethics

Editorial

Introduction

Medicine is at its largest technological inflection point in a century. Thanks to artificial intelligence (AI), algorithms already match or exceed clinicians at diagnosing skin cancer, reading chest radiographs, predicting acute decompensation, and stratifying oncologic risk [1]. The capacity to deliver more accurate, equitable, and humane care is genuinely within reach. Yet the profession is poorly prepared to receive it, and the reason is not technical. It is cultural. Medicine has long struggled to admit that it might be wrong, and that disposition has a name: hubris.

Confidence itself is essential to clinical practice. It helps clinicians make decisions under pressure, lead teams, and reassure patients. But when confidence tips into hubris, it becomes dangerous, manifesting as overreliance on personal experience, dismissal of uncertainty, and resistance to feedback. The paradox of expertise is that the more skilled or experienced we become, the more vulnerable we are to this drift [2].

With each wave of medical innovation, from handwashing to anesthesia to germ theory, progress was slowed by this same resistance, often at the cost of decades of preventable harm. AI raises the stakes. It moves faster than prior innovations, distributes accountability more diffusely across vendors, regulators, deployers, and clinicians, and embeds judgment in code that few practitioners can audit [2]. If hubris once blocked the absorption of soap and observation, it risks distorting the absorption of intelligence. The cost will fall on patients at both ends of the curve: those denied the benefit of better tools and those harmed by tools deployed without the humility their use demands. Addressing hubris is no longer a matter of professional culture in the abstract. It is now a precondition for the responsible adoption of AI. In this domain, perfect is not the enemy of good. Hubris is.

The psychology of hubris

We use hubris in a specific sense rather than as a synonym for arrogance. Owen and Davidson defined hubris syndrome as an acquired disorder of power characterized by 14 specified features, including contempt for the advice of others, loss of contact with reality, and a belief that accountability is owed to history rather

than to colleagues [3]. Operationally, we define clinical hubris as confidence that is not calibrated to accuracy and that is insulated from feedback. Both properties are measurable. Miscalibration is the gap between a clinician's stated certainty and the observed correctness of the decision. Feedback resistance is the degree to which disconfirming information fails to revise that certainty. Berner and Graber, reviewing the empirical literature on diagnostic error, found that physicians consistently underestimate the likelihood that their diagnoses are wrong, that this miscalibration contributes to error, and that clinical systems return too little corrective feedback to close the gap [4]. This is the construct we carry through the argument, and it is the construct at issue whenever a clinician accepts or rejects an algorithmic recommendation.

To see why AI may amplify rather than dissolve medical hubris, we must first understand how hubris is reinforced inside the individual clinician. Healthy confidence is grounded in reality and openness to feedback. Hubris involves a loss of humility, in which confidence becomes arrogance and entitlement. This drift is biologically reinforced. When clinicians achieve successful outcomes or earn admiration, dopaminergic reward circuits activate, strengthening a sense of personal mastery. In a profession where confidence is equated with competence, the reinforcement compounds unchecked. Senior leaders, including senior clinicians, develop grandiose self-beliefs upon attaining authority, a phenomenon known as hubris syndrome [2,3].

Three cognitive biases sustain medical hubris: the illusion of control (overestimating one's influence on outcomes), confirmation bias (selectively recognizing evidence that supports prior judgment), and attribution bias (externalizing failure as a system problem rather than a personal error). A fourth mechanism is harder to measure but no less consequential: empathy erosion, that is, the gradual diminishment of regard for others' perspectives as self-importance grows. The hierarchical structure of medicine creates echo chambers in which senior clinicians are rarely challenged [2]. Once an approach becomes mainstream, it is exceedingly difficult to dislodge, even when superior alternatives exist. These same biases will shape how every clinician interprets algorithmic predictions, model-derived risk scores, and chatbot recommendations. Hubris in the AI era begins in the same place it has always begun - in the clinician's mind. But it does not stay there. Hierarchy determines whether a junior clinician will challenge a senior's overconfident endorsement of an algorithm or whether silence will once again rule the room.

A 2,500-year pattern

These individual mechanisms have a long collective history. The Hippocratic Corpus warned against physicians who claimed certainty beyond what their knowledge justified [5], and the author of On the Sacred Disease explicitly criticized practitioners who invoked divine causes rather than seeking rational explanations [6]. Galen's writings, while pioneering, came to dominate medical thought for more than a millennium, stifling innovation until the Renaissance challenged Galenic dogma [7]. Writing in 12th-century Córdoba, Ibn Rushd cautioned that medicine operates in a sphere of possibility, in which general truths must be linked through natural philosophy to the contingent observations of practice, and that arrogance about outcomes is itself a moral failure [8].

The classical pattern is reliably tragic: hubris (arrogance) breeds *atis* (folly), which provokes nemesis (retribution), ending in *tisis* (destruction). Modern medicine has retraced this arc repeatedly. Ignaz Semmelweis identified handwashing as the key to preventing puerperal fever in 1847, but the medical establishment, unwilling to accept that its own hands were vectors of contagion, rejected his findings; women continued to die for decades [9]. Walter Freeman performed thousands of transorbital lobotomies based on his conviction that he understood the brain [10]. The Tuskegee syphilis study [11] and the Trovan trial in Nigeria [12] show how institutional hubris, the confident belief in one's own beneficence, has produced some of medicine's most grievous harms.

The pattern persists in our own time. The amyloid hypothesis of Alzheimer's disease shaped research priorities, National Institutes of Health (NIH) funding, and drug development for many years under the imprimatur of peer review, even as one failed trial accumulated on another and a landmark 2006 paper central to that frame was subsequently retracted [13]. The hubris in this case was not individual but collective: a thought collective so confident in its frame that disconfirming results were treated as noise, and competing hypotheses struggled for funding. The same peer-review apparatus that took decades to question amyloid is now being asked to validate clinical AI.

These are not the failings of a few aberrant individuals. They reflect what Ludwik Fleck called thought collectives: shared cognitive frameworks that become dominant within scientific communities and resist revision even in the face of disconfirming evidence [14]. Thomas Kuhn extended this insight, showing that scientific progress is not cumulative but punctuated by paradigm shifts that often occur only when the old guard retires [15]. Foucault argued that medicine's clinical gaze objectifies patients and exercises a form of biopower that benefits the profession at the expense of the people it serves [16,17]. Each framework points to the same diagnosis: medicine is structurally rigid, and hubris is the mechanism of that rigidity [2]. Each prior technological inflection has tested this rigidity. AI is the next test, and the conditions may favor a more concentrated collision. These parallels are heuristic rather than evidentiary. They do not establish that AI will fail in the way earlier innovations were resisted, and the mechanisms differ: handwashing and germ theory were resisted because the evidence was new, whereas algorithmic judgment is resisted or accepted

because it is opaque. What the parallels establish is narrower. The disposition that delayed earlier innovations has not been removed from the institutions now adopting AI, and no feature of AI corrects for it.

Not every inflection has produced this outcome. In 1975, faced with the new tools of recombinant DNA, a self-organized group of molecular biologists convened the Asilomar Conference and voluntarily imposed a moratorium on certain experiments until safety frameworks could catch up [18]. The exception is instructive precisely because it is so rare. It demonstrates that a profession can pause itself when humility demands it.

The analogy has limits that matter for what can realistically be asked. Asilomar worked under conditions that clinical AI does not share. The work was concentrated in a small number of academic laboratories, the hazardous experiments were few and identifiable, nothing was yet in clinical use, and the participants themselves controlled the means of doing the work. Clinical AI is built by commercial vendors, procured by thousands of independent health systems, regulated across jurisdictions, and already in use at the bedside. No single body could impose a moratorium, and there is no longer a point before deployment at which one could be imposed.

The feasible analogue is a standard rather than a pause. What Asilomar produced was a self-imposed rule about what evidence had to exist before work proceeded, and that is enforceable at the points where the profession already holds authority. Specialty societies can define the prospective validation a tool must have before endorsement. Journals can require external validation before publishing performance claims. Health systems can require local performance data before procurement. None of this requires a moratorium. Medicine has not set such a standard for AI.

How medicine manufactures hubris

If hubris is the mechanism, training is its factory. Medical education has been described as the slow but steady relinquishment of humility, wherein ignorance is replaced by knowledge and, crucially, doubt is replaced by assurance. As Jay Katz observed, trainees are socialized into the role of bearers of certainty, a posture that becomes a psychological defense against the high stakes of clinical work [19].

The practice of pimping, in which attendings interrogate junior trainees in front of peers, exemplifies this socialization. The student humiliated on rounds learns not humility but the choreography of dominance, replicating the behavior once they ascend the ranks. Atul Gawande captured the consequence: medicine is "an imperfect science, an enterprise of constantly changing knowledge, uncertain information, fallible individuals, and at the same time lives on the line" [20]. Humility, not certainty, is what that reality demands. James Li captured a parallel cost: compassion is itself a form of self-effacement incompatible with an overinflated sense of self [21]. AI cannot supply this self-effacement; only the clinician can. And the clinicians trained today are the same clinicians who will decide, over the next decade, whether each algorithmic recommendation is challenged or rubber-stamped. The product of this training is the substrate on which AI will be deployed.

Self-deception and the beginner's mind

Richard Feynman's first principle of doing science was not about data or method. It was about self-deception: "the first principle is that you must not fool yourself, and you are the easiest person to fool" [22]. The Zen concept of *shoshin*, or beginner's mind, arrives at the same place from a different direction. Expertise narrows possibility; openness expands it. Together, they point to something institutions consistently get wrong. The moment accumulated knowledge gets mistaken for complete knowledge, the questions that matter most stop getting asked.

Hubris in medicine and AI is not a character flaw but a structural condition, built into the incentives that reward certainty, punish doubt, and treat technical competence as epistemic completeness [2]. That is the arrogance trap. It swallows disruption quietly, long before anyone declares a failure. Feynman and *shoshin* converge here not as philosophy but as practice, a repeatable willingness to interrogate one's own foundations. Critical thinking does not begin with better tools. It begins with humility - the disciplined commitment to not fool oneself.

A useful contemporary illustration comes from outside medicine. In Andy Weir's Project Hail Mary, the protagonist Ryland Grace survives an impossible problem not through brilliance but through intellectual humility: he admits ignorance, tests variables, revises hypotheses, and stays open to suggestions throughout. Intellectually humble people say things like "Tell me more" or "I wish I had thought of that." They do not feel threatened by admitting vulnerability. The trait is too often undervalued in those who hold authority, which in medicine includes the attendings, chairs, and chief medical officers whose example trainees absorb daily. It also describes the disposition that responsible AI adoption now requires of all of us.

The AI inflection point

Is AI itself inherently hubristic? Medicine's culture of infallibility, sometimes called the God Complex, primes it to become so [2]. The profession valorizes the decisive, the visionary, and the builder, and it

rewards the launch of tools rather than the accumulation of longitudinal safety data. That asymmetry, between what we can build and what we are willing to own, embeds hubris into workflows and shapes clinical behavior long after the press release fades.

Every previous test of medical hubris was eventually resolved by accumulating evidence: handwashing was vindicated by trial data, anesthesia by surgical outcomes, and germ theory by microbiology. AI tests hubris differently. Earlier challenges displaced individual conviction with collective evidence. AI risks re-individualizing it. Judgment is now being embedded in code, accountability dispersed across vendors, regulators, deployers, and clinicians, and consequences concentrated on patients [2]. The same rigidity that delayed soap risks distorting the integration of intelligence, unless we name it explicitly.

Hubris in clinical AI takes six overlapping forms (Table *1*) [2]. Technological and regulatory hubris sit upstream, in the assumption that retrospective accuracy and regulatory clearance are proof of prospective benefit. Institutional and clinical hubris sit at the bedside, where a tool branded evidence-based by leadership is trusted regardless of local performance, and where clinicians either dismiss an algorithm that contradicts expert intuition or defer to one because it sounds authoritative. Replacement and ethical hubris sit at the horizon, in the proposition that AI can substitute for clinicians and in the assumption that an effective tool is automatically a just one, leaving unresolved questions of data provenance, consent, algorithmic bias, and the unequal distribution of benefits. Each form replays an older pattern in which confidence outruns calibration, and the consequences fall on patients. Each is an expression of the same underlying error: mistaking what AI can do for what AI should do.

The issue: a form of AI hubris	The action: what must change
Technological hubris: assuming a model's retrospective accuracy will translate to prospective benefit.	Require prospective, locally validated evidence before deployment (developers; journals).
Regulatory hubris: treating regulatory clearance as validation - permission to sell as proof to heal.	Stop equating clearance with clinical endorsement (regulators).
Institutional hubris: trusting any leadership-branded tool regardless of local performance.	Empower Patient and Family Advisory Councils with authority over procurement and deployment (hospitals).
Clinical hubris: dismissing a model that dissents, or deferring because it sounds authoritative.	Teach calibrated use of AI; write epistemic humility into competency frameworks (specialty boards).
Replacement hubris: believing AI can replace clinicians, discounting relational and ethical care.	Treat AI as augmentation; select and train for humility (training programs).
Ethical hubris: assuming an effective tool is automatically a just one, ignoring bias, consent, and unequal benefit.	Weigh justice, not just efficacy; independently verify vendor claims (institutions).

TABLE 1: From hubris to humility in clinical AI

Epistemic humility is the precondition: treat AI as a hypothesis to be tested, not a product to be trusted on clearance alone. AI: artificial intelligence.

Table created by the authors. The six forms of hubris are based on Patel et al. [2], which is published under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License (http://creativecommons.org/licenses/by-nc-nd/4.0/)

The promise that makes these failure modes so consequential is real. AI tools have demonstrated dermatologist-level accuracy on skin cancer images, radiologist-level performance on chest radiographs, and the ability to predict acute kidney injury hours before clinical recognition [1]. Adjacent advances include image reconstruction and anatomical landmarking in nuclear medicine, radiomics-based phenotyping in radiology, risk stratification and drug discovery in oncology, and systems that tailor treatment recommendations to a patient's genetic, lifestyle, and clinical context [1]. Deployed wisely, these tools could free clinicians from the cognitive load that fragments attention, restoring the patient-physician relationship that hubris and overwork have hollowed out.

But hubris in AI is not an individual moral failing. It is a systems pathology, a distortion of incentives, feedback, and accountability that converts good intentions into structural risk. Consider the Epic Sepsis Model, deployed in hundreds of U.S. hospitals before independent external validation showed substantial alert fatigue [23]. Consider the widely used commercial algorithm shown by Obermeyer and colleagues to systematically underestimate the health needs of Black patients because it used healthcare cost as a proxy for illness [24]. These are not the failures of bad actors. They are the failures of a system that confuses regulatory clearance with proof of benefit, and that distributes accountability so unevenly that no one fully owns the outcome.

Beneath this pathology sits a structural reality the profession is slow to name. Science is not neutral; it is shaped by who funds it and who profits from its applications. The commercial enterprise of healthcare has produced a powerful bias toward successful commercialization of discovery, channeling investment toward device, biotechnology, and pharmaceutical applications and rendering regulatory approval a procedural ritual rather than a substantive judgment about benefit. Clinical AI is being developed and deployed almost entirely by for-profit vendors. The institutional reflex to treat their products as scientific tools, rather than as commercial offerings whose claims require independent verification, is itself a form of hubris, the hubris of mistaking market presence for proof.

The moral geometry is stark: technologists own the patents, executives own the press releases, and clinicians own the consequences. A new species of medical error is already emerging, what may be called errors of deference, in which clinicians override their own judgment to satisfy an algorithm they did not design and do not understand. When both human and machine agree on the wrong answer, retrospective audit reveals nothing; the model is validated within its own echo chamber.

The cost is not only diagnostic. Each time a clinician suppresses judgment to satisfy a black-box recommendation, they lose a fragment of moral agency. Over time, this produces the burnout that institutions misdiagnose as a wellness problem, and the moral injury that drives clinicians out of the profession. At system scale, the same asymmetry generates three predictable secondary harms: defensive medicine driven by liability that never settles where decisions are actually made, suppressed innovation from clinicians unwilling to take on uncompensated risk, and eroded patient trust as the source of clinical judgment becomes increasingly opaque. AI does not create new pathologies so much as it accelerates inherited ones, magnifying the asymmetry between those who design clinical reality and those who carry its consequences.

There is a final caution before turning to remedy. Trust in science is appropriate; belief in science is dangerous, because belief generates hubris and the resistance to revision that hubris reliably brings. The same caution must apply to AI. We are being asked to trust models that few of us can audit, in a profession that has historically rewarded confidence over calibration. A healthy skepticism is not an obstacle to AI adoption. It is its precondition.

What must change

If AI is to fulfill its promise rather than amplify the profession's oldest failures, hubris cannot be legislated away. It must be designed out. Three changes are necessary, each targeting a different point at which hubris enters the system: who is selected into medicine, how they are trained, and what evidence a tool must produce before it reaches a patient. Table *1* pairs each form of hubris with the corrective action and the party accountable for it.

First, medicine must select for humility, and against the predispositions that hierarchy then amplifies. For more than a century, medical schools have favored applicants hardwired for competitive achievement within rigidly hierarchical training structures, then trained them in publish-or-perish academic systems that reward heroic individualism. If we continue choosing hubris and embedding it in hierarchy, no educational reform will succeed, and no amount of AI literacy training will produce clinicians capable of holding an algorithm accountable. Admissions must weigh intellectual humility, comfort with uncertainty, and collaborative instincts alongside grades and test scores, drawing on the multiple mini-interview literature that has shown such traits can be assessed reliably.

Second, medicine must teach anti-hubris competencies. Critical thinking presupposes humility, because it requires accepting that one might be wrong. Trainees should learn to articulate uncertainty, solicit disconfirming evidence, and treat patient input as primary data rather than corroborative anecdote. The pimping ritual should be retired. Bedside teaching can model fallibility explicitly; attendings who say "I don't know, let's look it up together" teach more about medicine than those who never falter. These same competencies are the ones that will distinguish clinicians who interrogate an algorithm's output from those who simply sign beneath it. Patient and Family Advisory Councils should be standard institutional structures with explicit authority over AI procurement and deployment decisions, formalizing what Sackett described as evidence-based medicine's three foundational pillars: patient-reported, clinician-observed, and research-derived evidence [25]. Models trained without the first pillar are not evidence-based; they are partially blind.

Third, AI must be treated as a hypothesis, not a product. Prospective trials should precede wide deployment. Auditability and explainability should be safety features, not marketing claims. Liability should align with influence, so that developers and deployers share clinical risk with practitioners. Psychological safety must allow clinicians to question or reject algorithmic recommendations without professional cost.

The academic medicine that produced this culture is itself being democratized, with declining public research funding, alternative publication channels, and AI-enabled scholarship for junior faculty, clinicians, and patients themselves. This is not a threat to scientific integrity. It is an opportunity to rebuild on

humbler foundations, recognizing, as the Hippocratic Corpus already warned, that medicine operates in the realm of the possible rather than the certain [5].

Operationally, this agenda is concrete. It means investing in strategies that surface and address clinician resistance to AI without dismissing legitimate concerns; building prospective validation frameworks that test models in the populations they will serve; articulating governance for data privacy and algorithmic bias; designing AI-collaboration curricula for clinicians at every career stage; and committing to longitudinal study of how AI integration changes diagnostic accuracy, treatment efficacy, and patient outcomes over time. None of these is a novel idea. What has been missing is the institutional humility to fund and follow them when their findings prove inconvenient.

Conclusion

The Greek tragedians knew that hubris always invites its own destruction. Medicine has rediscovered this truth in every century, and we are about to rediscover it again. AI will not destroy medicine through malice. The risk is corrosion through misplaced certainty. The safeguard is not technological restraint but epistemic humility, embedded as the core operating principle of intelligent care. We swore to heal, not to conquer. The covenant is renewable, but only if we are willing to admit that we are not gods. We are gardeners: tending, listening, and learning. Used well, AI augments rather than replaces clinical judgment. That distinction is not technical; it is ethical, and it is ours to make.

The call is concrete, and the timeline is immediate. Admissions committees can begin selecting for intellectual humility in the next cycle. Department chairs can retire pimping this academic year. Hospitals require prospective, locally validated evidence before procuring any AI tool. Specialty boards can write epistemic humility into competency frameworks. Regulators can stop treating clearance as endorsement. Journals can require external validation before publishing performance claims. None of this requires a new technology, a new study, or a new generation. It requires only the willingness, today, to admit what every generation of medicine has refused: that we are not the source of truth. We are its students. AI will not change that. It will only raise the cost of pretending otherwise.

Medicine had its Asilomar once. The question now is whether it will set a comparable standard for AI, or whether it will be remembered as the profession that had every example to learn from and chose to learn from none of them.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Aarti Sarwal, Gopal K. Chopra, Ann Hui Ching, Matthew Crowson, Leo Anthony Celi

Acquisition, analysis, or interpretation of data: Aarti Sarwal, Gopal K. Chopra, Ann Hui Ching, Matthew Crowson, Leo Anthony Celi

Drafting of the manuscript: Aarti Sarwal, Gopal K. Chopra, Ann Hui Ching, Matthew Crowson, Leo Anthony Celi

Critical review of the manuscript for important intellectual content: Aarti Sarwal, Gopal K. Chopra, Ann Hui Ching, Matthew Crowson, Leo Anthony Celi

Supervision: Leo Anthony Celi

Disclosures

Conflicts of interest: In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** Leo Anthony Celi declare(s) a grant from Leo Anthony Celi is funded by the US National Institutes of Health (DS-I Africa U54 TW012043-01; Bridge2AI OT2OD032701), the US National Science Foundation (ITEST #2148451), the Boston-Korea Innovative Research Project (RS-2024-00403047), and the Korea Health Technology R&D Project (RS-2024-00439677) through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Republic of Korea. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

References

1. Topol EJ: High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019,

25:44-56. 10.1038/s41591-018-0300-7
2. Patel MS, Scales D, Arslan J, Pedrelli P, Crowson MG, Abulibdeh R, Celi LA: Hubris in the age of intelligent medicine-from clinical hierarchies to algorithmic amplification. NPJ Digit Med. 2026, 9:10.1038/s41746-026-02717-z
3. Owen D, Davidson J: Hubris syndrome: an acquired personality disorder? A study of US Presidents and UK Prime Ministers over the last 100 years. Brain. 2009, 132:1396-406. 10.1093/brain/awp008
4. Berner ES, Graber ML: Overconfidence as a cause of diagnostic error in medicine . Am J Med. 2008, 121:2-23. 10.1016/j.amjmed.2008.01.001
5. Jouanna J: Hippocrates. Johns Hopkins University Press, Baltimore and London; 1999.
6. Hippocrates: The sacred disease . Hippocrates, Volume II. Harvard University Press, Cambridge; 1923. 10.4159/DLCL.hippocrates_cos-sacred_disease.1923
7. Nutton V: Ancient Medicine, Second Edition . Routledge, London; 2012. 10.4324/9780203081297
8. Chandelier J: Averroes on medicine . Interpreting Averroes: Critical Essays. Adamson P, Di Giovanni M (ed): Cambridge University Press, 2018. 10.1017/9781316335543.010
9. Kadar N: Rediscovering Ignaz Philipp Semmelweis (1818-1865) . Am J Obstet Gynecol. 2019, 220:26-39. 10.1016/j.ajog.2018.11.1084
10. Lichterman BL, Schulder M, Liu B, Yang X, Taira T: A comparative history of psychosurgery . Prog Brain Res. 2022, 270:1-31. 10.1016/bs.pbr.2021.12.003
11. Tobin MJ: Fiftieth anniversary of uncovering the Tuskegee syphilis study: the story and timeless lessons . Am J Respir Crit Care Med. 2022, 205:1145-58. 10.1164/rccm.202201-0136SO
12. Ezeome ER, Simon C: Ethical problems in conducting research in acute epidemics: the Pfizer meningitis study in Nigeria as an illustration. Dev World Bioeth. 2010, 10:1-10. 10.1111/j.1471-8847.2008.00239.x
13. Lesné S, Koh MT, Kotilinek L, et al.: Retraction note: a specific amyloid-β protein assembly in the brain impairs memory. Nature. 2024, 631:240. 10.1038/s41586-024-07691-8
14. Fleck L: Genesis and Development of a Scientific Fact . Trenn TJ, Merton RK (ed): University of Chicago Press, 1979. 10.7208/chicago/9780226190341.001.0001
15. Kuhn TS: The Structure of Scientific Revolutions . University of Chicago Press, Chicago; 1962. 10.7208/chicago/9780226458144.001.0001
16. Foucault M: The Birth of the Clinic, First Edition . Routledge, London; 1977. 10.4324/9780203406373
17. Foucault M: The History of Sexuality, Volume 1: An Introduction . Pantheon Books, New York; 1978.
18. Berg P, Baltimore D, Brenner S, Roblin RO, Singer MF: Summary statement of the Asilomar conference on recombinant DNA molecules. Proc Natl Acad Sci USA. 1975, 72:1981-4. 10.1073/pnas.72.6.1981
19. Katz J: The Silent World of Doctor and Patient . The Johns Hopkins University Press, Baltimore and London; 1984.
20. Gawande A: Complications: A Surgeon's Notes on an Imperfect Science . Metropolitan Books, 2002.
21. Li JT: Humility and the practice of medicine . Mayo Clin Proc. 1999, 74:529-30. 10.4065/74.5.529
22. Feynman RP: Cargo cult science . Eng Sci. 1974, 37:10-13.
23. Wong A, Otles E, Donnelly JP, et al.: External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med. 2021, 181:1065-70. 10.1001/jamainternmed.2021.2626
24. Obermeyer Z, Powers B, Vogeli C, Mullainathan S: Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019, 366:447-53. 10.1126/science.aax2342
25. Sackett DL, Rosenberg WM, Gray JA, Haynes RB, Richardson WS: Evidence based medicine: what it is and what it isn't. BMJ. 1996, 312:71-2. 10.1136/bmj.312.7023.71