Viewpoint

Four Imperatives for US Health Systems Governing Racial Bias in Health Care AI

Mike Runquist, BS, MLS

JMR Consultancy LLC, Los Angeles, CA, United States

Corresponding Author:
Mike Runquist, BS, MLS
JMR Consultancy LLC
500 S Sepulveda Blvd
4th Floor, Suite 100 A
Los Angeles, CA, 90049
United States
Phone: 1 213 863 9576
Email: miker@jmrconsultancy.net

Abstract

Racial and ethnic disparities in US health care are long-standing, and AI is being adopted across clinical and administrative functions faster than the governance needed to detect inequity. This Viewpoint is written for the leaders of US health systems who decide what AI tools are bought, deployed, and monitored and who therefore hold the most direct leverage over whether those tools narrow or widen disparities. It advances 4 imperatives: govern before you deploy, demand transparency from vendors, build equity into clinical workflows, and make disparities visible. The imperatives are a leadership-level operationalization of existing frameworks (the Coalition for Health AI Blueprint, the National Institute of Standards and Technology AI Risk Management Framework, and the Rajkomar et al health equity framework), not a new framework. The contributions claimed are limited to 4 main points: allocating specific decisions to specific leadership functions; 3 reusable tools provided as Multimedia Appendices (a preprocurement equity audit checklist, illustrative contract clauses, and criteria for retaining race or ancestry as a model variable); a tiered approach for small and safety-net systems; and an assessment of how the imperatives hold up as US federal requirements are narrowed. The evidence base is concentrated on anti-Black bias, where it is deepest, and the paper identifies evidence gaps for other groups. A brief mapping to the European Union AI Act and the Assessment List for Trustworthy AI is offered for readers outside the United States.

(JMIR AI 2026;5:e103163) doi: 10.2196/103163

KEYWORDS

racial disparities; algorithmic bias; health care AI; health equity; AI governance; health system leadership; clinical decision support; United States; artificial intelligence; AI

Introduction

Background

In 2023, the US maternal mortality rate was 50.3 deaths per 100,000 live births for non-Hispanic Black women and 14.5 for non-Hispanic White women, a ratio of about 3.5 to 1 [1]. In a US outpatient registry, Black patients who were treated with a statin were less likely than White patients to receive it at guideline-recommended intensity (33% vs 44%) [2]. In vignette studies, medical trainees who endorsed false beliefs about biological differences between Black and White people rated Black patients' pain as lower and made less accurate treatment recommendations [3]. These gaps are the product of differential access, quality of care, and trust, not of innate biology.

AI is being adopted in these same systems. Predictive models flag patients for care management, imaging algorithms assist with diagnosis, and large language models (LLMs) summarize records, draft prior authorization requests, and increasingly enter decision support. Published evidence shows what can happen when such tools are built on unrepresentative data or on proxies for need. A commercial care management algorithm that used health care cost as a proxy for need assigned Black patients lower risk scores than equally sick White patients [4]. Chest radiograph classifiers underdiagnosed Black, Hispanic, female, and Medicaid-insured patients [5], and dermatology algorithms performed worse on darker skin tones [6]. These studies do not show that every deployed tool is biased. They show that bias can be invisible in aggregate accuracy reporting, which makes equity-focused oversight a precondition for knowing.

XSL•FO
RenderX

Aim

The aim of this Viewpoint is to set out 4 operational imperatives for US health system leaders, to state precisely what is and is not original to them relative to existing frameworks, and to supply tools that leaders can adapt. The argument is not against AI. It is for AI governance worthy of the trust placed in it.

Scope and Terminology

This Viewpoint is deliberately US focused. The regulatory instruments it discusses (Section 1557 of the Affordable Care Act, the Office of the National Coordinator for Health Information Technology [ONC] certification program, and the Food and Drug Administration [FDA] and Centers for Medicare & Medicaid Services [CMS] programs) are US instruments, and the author's practice experience is with US health systems. Nevertheless, the imperatives are written to be portable, even if their regulatory grounding is US specific. A short mapping to the European Union (EU) AI Act and Assessment List for Trustworthy AI (ALTAI) appears in the Beyond the United States section.

The empirical examples are predominantly Black-White comparisons because that is where the evidence is deepest. The evidence for other groups exists but is thinner. For example, in a single-center cohort, estimated occult hypoxemia (arterial saturation below 88% despite a pulse oximetry reading of 92% or higher) was 6.2% in Black, 6.6% in Asian, and 6.6% in American Indian patients, compared with 3.6% in White patients; the difference was statistically significant after adjustment only for Black patients (odds ratio 1.65, 95% CI 1.28-2.14), and the CIs for the smaller groups were wide [7]. This pattern—a plausible signal that is underpowered because the subgroup is small or poorly measured—is itself an argument for imperative 4. Readers should treat the imperatives as group-agnostic in design and anti-Black bias as the best-evidenced application.

Race is used here as a social category. Where the clinical signal in question is genetic ancestry or genotype, the text says so (see the Race, Ancestry, and Clinical Signal section).

A Conceptual Model of Inequity in the AI Life Cycle

Figure 1 shows the model that organizes the argument. Structural inequity in society and the health system produces inequitable data: datasets that underrepresent minoritized groups, misclassify social experience as biological difference, and undermeasure the most marginalized patients. Inequitable data feed inequitable models, which optimize for majority populations, encode flawed proxies as signal, and lack fairness constraints [8]. Inequitable models yield inequitable deployment, in which performance differs across institutions, workflows amplify algorithmic error through automation bias [9], and monitoring is too weak to detect disparate impact. Deployment yields inequitable outcomes, which reenter the data environment and reinforce the cycle. The imperatives are positioned at the stages where a health system leader holds decision authority.

Figure 1. Conceptual model of how structural inequity moves through the health care AI life cycle, with the 4 imperatives positioned at the stages where health system leaders hold decision authority.

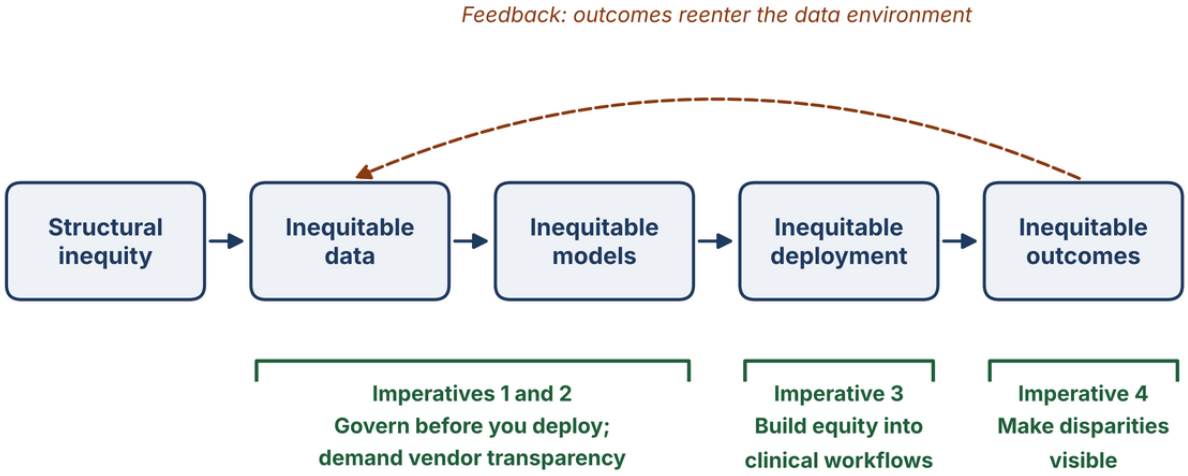

What Is and Is Not New Here

The Coalition for Health AI (CHAI) Blueprint [10], the National Institute of Standards and Technology (NIST) AI Risk Management Framework [11], and Rajkomar et al [8] already call for governance structures, documentation, subgroup evaluation, and monitoring. The imperatives restate none of these as new principles. Table 1 separates what the source frameworks supply from what this Viewpoint adds, so that readers can judge the contribution directly. The additions are modest by design: they concern who decides, with what artifact, and under what resource constraint.

XSL•FO
RenderX

Table 1. Source frameworks versus the contribution of this Viewpoint for each imperative.

Imperative	Already in the source frameworks	Added in this Viewpoint
Govern before you deploy	Life cycle governance, risk management, and fairness assessment are described at the organizational or system level	Assigns predeployment decision rights to named leadership functions; supplies an equity audit checklist and governance charter elements (Multimedia Appendix 1)
Demand transparency from vendors	Model cards, datasheets, and documentation expectations	Converts documentation expectations into enforceable purchase conditions with illustrative clause language (Multimedia Appendix 2); adds vendor-supplied guidance on where a tool is least reliable and notice of model updates
Build equity into clinical workflows	Human oversight and monitoring of deployed systems	Moves part of the training burden upstream to design evaluation and vendor disclosure and keeps clinician training risk-tiered rather than universal
Make disparities visible	Continuous monitoring and postdeployment surveillance	Specifies disaggregated metrics, ties them to executive dashboards, and requires preagreed action thresholds
Cross-cutting	Fairness definitions and their trade-offs	Criteria for retaining race or ancestry as a model variable, reconciled with compliance duties (Multimedia Appendix 3); a tiered approach for resource-limited systems; a test of the imperatives against a narrowing US regulatory floor

Race, Ancestry, and Clinical Signal

A foundational distinction deserves explicit treatment: the difference between race as a proxy for social experience and genetic ancestry or genotype as a clinical signal. AI systems that conflate them will produce both inequity and clinical error.

There are contexts in which ancestry-linked variation carries clinical evidential weight. Pathogenic variant prevalence in BRCA1 and BRCA2 differs across ancestry-defined populations, and some pharmacogenomic variants that affect drug metabolism vary in frequency with ancestry [12]. In these contexts, the clinically meaningful variable is the genotype or ancestry-informative marker, and self-reported race is at best an imperfect stand-in for it. Where a genotype can be tested, the genotype should be used instead of race.

The problem arises when race functions as a proxy for social experience and the model or equation treats it as innate. The now-deprecated race coefficient in estimated glomerular filtration rate equations encoded patterns of access to care and structurally driven comorbidity as physiology, producing higher kidney function estimates for Black patients; race-free equations have replaced it [13,14]. Spirometry reference equations that applied lower normal values to Black patients carried forward 19th-century assumptions [15]. Whether such corrections reached downstream predictive tools, and with what documented effect, is not established by the sources cited here. The more defensible claim is that any model trained on or built with race-corrected inputs or labels inherits the correction, and that leaders should ask vendors to say so.

Fairness criteria also conflict. Calibration, equalized odds, and demographic parity generally cannot all be satisfied at once when base rates differ [16,17]. That is a reason to choose among them deliberately and transparently in each clinical context, with an equity-informed governance body involved, rather than leaving the choice to developers.

The operational principle for leaders is that race should be an input to a clinical algorithm only when there is clear, contemporaneous evidence of an independent clinical signal that cannot be captured by a measured biological or social variable. Multimedia Appendix 3 turns this principle into a 5-question decision aid.

This principle must be reconciled with a compliance duty. Under 45 CFR 92.210, in effect since the compliance date of May 1, 2025, covered entities must make reasonable efforts to identify patient care decision support tools that use input variables measuring race, color, national origin, sex, age, or disability and to mitigate the risk of discrimination from their use [18,19]. The provision does not prohibit race as an input. It requires that its use be identified, reviewed, and mitigated. A documented clinical justification under Multimedia Appendix 3 is therefore both a clinical governance record and evidence of the reasonable efforts that the rule requires. Legal counsel should confirm the application to specific tools.

Four Imperatives for Health System Leaders

Imperative 1: Govern Before You Deploy

The most consequential governance decisions precede deployment: what tools to procure, from whom to procure them, on which populations to validate them, and what contractual accountability to require. It is reasonable to expect that systems that build governance after tools are embedded face greater organizational and contractual barriers to correction; this is a predictive judgment grounded in how contracts and workflows become entrenched, not a measured finding.

An AI governance committee with decision authority, not only advisory capacity, is the institutional mechanism. Membership should include clinical leadership, compliance and privacy, data science, health equity, and community representation from the populations served. Community members bring knowledge of the contexts in which tools will operate and of which outcomes matter to those populations.

Equity audits should be required before any tool that materially influences clinical decisions or resource allocation is approved. Multimedia Appendix 1 provides a checklist covering documentation, training data, subgroup performance, fairness

XSL•FO
RenderX

choices, race-variable justification, interface design, monitoring, and contract terms. Obermeyer et al [4] showed why: at a given risk score, Black patients were substantially sicker than White patients, and correcting the label would have raised the share of Black patients identified for extra care from 17.7% to 46.5%. The bias was invisible in aggregate reporting and appeared only in disaggregated analysis. The investigators also reported that, working with the manufacturer, they reduced the measured bias by 84% using an alternative label [4], which is evidence that vendor engagement can work when the purchaser asks for disaggregated evidence.

Design evaluation should include human factors criteria. Does the interface display uncertainty? Does it prompt for clinical override and record overrides? Does it flag patient groups in which the tool is least reliable? These criteria place part of the automation bias risk [9] where it can be addressed once, at selection, rather than repeatedly through training.

Procurement contracts are governance instruments. Commitments to subgroup performance reporting, notice of model changes, and bias remediation timelines are enforceable in ways that sales-stage assurances are not. Illustrative clause language is in Multimedia Appendix 2.

Imperative 2: Demand Transparency From Vendors

Vendors respond to what purchasers consistently require, and a purchaser who accepts documentation without subgroup data signals that such data are optional. This is a normative argument about market signals, not an empirical finding about vendor behavior.

Model cards [20] and datasheets for datasets [21] are the minimum documentation: intended use, training data composition, performance disaggregated by race and ethnicity, limitations, and failure modes. Health systems should require both and treat refusal to disclose subgroup performance as informative. Two additions go beyond the documentation standards: first, vendor-supplied guidance on populations and settings where the tool is least reliable, which feeds clinician training and interface design, and second, advance notice of model updates with revised subgroup results, so that the evaluation and training burden is not reincurred silently each time a tool changes.

LLMs need the same scrutiny. Commercially available models have been shown to propagate race-based medicine in response to clinical questions [22], so LLM-based decision support should be evaluated for subgroup performance and for output that reflects racial stereotypes.

Imperative 3: Build Equity Into Clinical Workflows

Governance will not produce equitable outcomes if the clinicians who use tools cannot recognize their limits. Automation bias, the tendency to overrely on algorithmic output, is a deployment-phase risk that clinicians' training alone cannot remove [9]. When a model systematically underestimates risk for a group, automation bias turns a statistical disparity into a clinical harm.

Training should be proportionate to risk. Tools that influence care management enrollment, referral, or treatment escalation warrant tool-specific training on known performance limits and on the patient groups in which reliability is lowest. Low-risk tools can rely on short, vendor-supplied guidance. Placing disclosure requirements on vendors (imperative 2) and override prompts in the interface (imperative 1) reduces the retraining load when tools change, which is a real cost that a training mandate alone would ignore.

Equity checkpoints complement training: defined moments in the workflow at which clinicians are prompted to consider whether the recommendation fits the direct clinical assessment and whether patients' characteristics might affect reliability. For high-stakes decisions, these prompts can be built into the electronic health record. A nonpunitive incident reporting pathway for suspected disparate recommendations closes the loop to governance and supplies context that quantitative monitoring misses.

Imperative 4: Make Disparities Visible

Periodic quality reports and annual outcome reviews are retrospective and aggregate. Equity surveillance monitors care gap closure, readmissions, time to diagnosis, and adherence, disaggregated by race, ethnicity, language, and insurance status, and surfaces results to leaders with authority to act. Presenting these metrics alongside financial and quality metrics changes who owns them: when a chief medical officer sees that overall readmission is within benchmark but readmission for Black patients with heart failure is not, the disparity becomes a leadership problem, not just a quality team problem.

Surveillance requires data infrastructure, analytic capacity, and preagreed action thresholds, including willingness to suspend or modify a tool when monitoring indicates disparate impact. Reports that no one acts on document harm and do not prevent it. Disaggregation also depends on the quality of race, ethnicity, and language data; where subgroups are small or collected as coarse categories, signals such as those described for pulse oximetry above [7] remain undetectable, so improving data collection is part of the imperative.

Feasibility for Resource-Limited and Safety-Net Systems

The imperatives, as written, assume capacity that small, rural, and safety-net systems may lack, and these systems often serve the populations most affected. They also have less purchasing leverage than large academic systems. Public data on the cost of governance programs are sparse, and this Viewpoint does not estimate costs; the absence of such data is a limitation and a research need. Table 2 proposes a tiered approach in terms of activities, so that every system can do something meaningful.

XSL•FO
RenderX

Table 2. Tiered approach to implementing the imperatives by organizational capacity.

Tier	Who it fits	Minimum actions
Baseline	Any system using AI that affects care or coverage	Keep an inventory of AI tools; require vendor answers to the Multimedia Appendix 1 core questions before purchase; assign a named accountable executive; review the race-variable justification for any tool that uses race; and provide an incident reporting route
Intermediate	Systems with an analytics team or a compliance function	Add a governance committee with community representation; run annual subgroup performance checks on high-risk tools using local data; and adopt standard contract clauses (Multimedia Appendix 2)
Advanced	Large academic and integrated systems	Add continuous equity surveillance dashboards, electronic health record–embedded checkpoints, independent validation, and published results

Smaller systems can narrow the gap by purchasing through consortia or group purchasing arrangements, by sharing templates and audit results with peers, and by drawing on shared or third-party assurance resources where available. Where a vendor will not provide subgroup data at any scale of purchase, that refusal is itself a decision-relevant finding. These mitigations reduce but do not remove the asymmetry, and leaders of small systems may reasonably conclude that some tools should not be adopted until evidence exists.

The Regulatory Horizon

The United States

US requirements bearing on this topic are shifting in both directions, which is the reason the imperatives are framed as leadership commitments rather than compliance tasks.

Under Section 1557, the AI-relevant provision (45 CFR 92.210) has been in effect since its May 1, 2025, compliance date [18,19]. The litigation to date has addressed other provisions. A federal court's October 2025 final judgment in *Tennessee v Kennedy* vacated provisions of the 2024 rule only to the extent that they extend sex discrimination to gender identity, and the Department of Health and Human Services has stated that the remaining provisions of the rule stay in force [23]. The characterization of Section 1557 as a single contested rule is therefore inaccurate for present purposes.

Health Data, Technology, and Interoperability (HTI-1) established source-attribute transparency requirements for predictive decision support interventions (DSIs) in certified health IT [24]. The Health Data, Technology, and Interoperability: ONC Deregulatory Actions to Unleash Prosperity (HTI-5) final rule, published in October 2026, removed the DSI source-attribute requirements and the intervention risk-management requirements for predictive DSIs, thereby eliminating the certification requirements that ONC had characterized as supporting "model cards" for predictive DSIs [25]. The documentation that imperative 2 asks vendors to supply is therefore no longer supplied by default through certification, and procurement contracts become the main lever.

CMS equity measures in the Hospital Inpatient Quality Reporting Program, including Hospital Commitment to Health Equity and the social drivers of health screening measures, were removed in the fiscal year 2026 final rule [26].

FDA guidance on AI-enabled device software functions points toward life cycle management and postmarket performance monitoring [27], although it governs devices rather than the health system's own procurement choices. The FDA's January 2025 guidance remains in draft form and had not been finalized as of November 2026.

Taken together, the one binding antidiscrimination provision (92.210) applies to the use of tools, the transparency floor that HTI-1 created has been removed, and the financial equity signal has been withdrawn. This supports the Viewpoint's central claim: for the foreseeable future, health system leaders, through governance and procurement, rather than federal mandates, will determine what equity evidence exists about the tools they use.

Beyond the United States

The EU AI Act [28] classifies many AI systems used in health care, including AI that is a medical device subject to third-party conformity assessment, as high risk. Providers of high-risk systems must operate a risk management system (Article 9), apply data governance practices that include examining and mitigating bias (Article 10), maintain technical documentation (Article 11), provide information that lets deployers interpret outputs (Article 13), design for human oversight (Article 14), and monitor performance after market placement (Article 72); deployers have their own obligations (Article 26).

The implementation timetable for the AI Act's high-risk requirements was subsequently amended through the EU Digital Omnibus on AI. The principal obligations for stand-alone high-risk systems (Annex III) now apply from December 2, 2027, while obligations for high-risk AI embedded in regulated products, including many medical devices (Annex I), apply from August 2, 2028 [29].

ALTAI [30] operationalizes 7 requirements, including diversity, nondiscrimination and fairness, transparency, human agency and oversight, and accountability. Table 3 maps the imperatives onto these instruments.

XSL•FO
RenderX

Table 3. Mapping of the 4 imperatives to the European Union (EU) AI Act and the Assessment List for Trustworthy AI (ALTAI).

Imperative	EU AI Act provisions (high-risk systems)	ALTAI requirements
Govern before you deploy	Risk management (Article 9); data governance (Article 10); deployer obligations (Article 26)	Accountability; diversity, nondiscrimination, and fairness
Demand transparency from vendors	Technical documentation (Article 11); transparency to deployers (Article 13)	Transparency; privacy and data governance
Build equity into clinical workflows	Human oversight (Article 14)	Human agency and oversight
Make disparities visible	Postmarket monitoring (Article 72)	Technical robustness and safety; societal well-being

Two caveats limit generalization. First, imperative 4 depends on the collection of race and ethnicity data. Under the General Data Protection Regulation, data revealing racial or ethnic origin are a special category with restricted processing [31], and some EU member states discourage or prohibit the collection of ethnicity data, so EU institutions may need other proxies for equity monitoring, such as language, migration background (where lawful), or geography. Second, the EU framework is binding on system providers and deployers in ways US certification rules are not, which strengthens the case for the vendor documentation that imperative 2 requests. US systems that buy from vendors serving both markets may be able to request EU-conformity documentation.

Limitations

This is a Viewpoint and does not evaluate the imperatives empirically. The evidence base is concentrated on US settings and on Black-White comparisons. Cost and effectiveness data for governance programs are largely unavailable. The regulatory landscape is changing quickly, and some statements may be outdated by publication. Multimedia Appendices 1 to 3 are starting points for adaptation with legal and clinical counsel and are not legal advice.

Conclusions

AI will not narrow racial disparities in health care unless health system leaders decide to build equity into each stage of its life cycle. The 4 imperatives are commitments leaders can make now through governance structures, procurement terms, clinical design, and data infrastructure. As federal requirements narrow, the evidence leaders insist on will matter more. The patients who most need this governance are those whose needs have been underestimated, whose data are least visible, and whose trust has been damaged by unequal care.

Data Availability

Data sharing is not applicable to this article as no datasets were generated or analyzed during this study.

Funding

The work was conducted independently, and the author declared that no financial support was received for this work.

Generative AI Use

The author used Claude (Claude Sonnet 5.5; Anthropic, accessed through claude.ai between May and November 2026) to assist with drafting, structural editing, and document formatting. All intellectual content represents the author's own work. All AI-assisted content was reviewed, verified against primary sources, and revised by the author before submission. The author takes full responsibility for the accuracy, integrity, and originality of the manuscript.

Conflicts of Interest

The author is the principal of JMR Consultancy LLC, which provides health care compliance and AI governance consulting. The imperatives and tools in this paper are drawn from published literature and general practice. They do not describe a proprietary framework, product, or service of the author's firm, and the Multimedia Appendices are published under the journal's open license.

Multimedia Appendix 1

Preprocurement equity audit checklist and governance charter elements.
[DOCX File , 39 KB-Multimedia Appendix 1]

Multimedia Appendix 2

Illustrative AI contract clauses for subgroup performance, change notice, and bias remediation.
[DOCX File , 38 KB-Multimedia Appendix 2]

XSL•FO
RenderX

Multimedia Appendix 3

Decision aid for retaining race or ancestry as a model variable.
[DOCX File , 37 KB-Multimedia Appendix 3]

References

1. Hoyert DL. Maternal mortality rates in the United States, 2023. National Center for Health Statistics Health E Stats. 2025. URL: https://doi.org/10.15620/cdc/174577 [accessed 2026-11-14]
2. Nanna MG, Navar AM, Zakroysky P, Xiang Q, Goldberg AC, Robinson J, et al. Association of patient perceptions of cardiovascular risk and beliefs on statin drugs with racial differences in statin use: insights from the Patient and Provider Assessment of Lipid Management Registry. JAMA Cardiol. Aug 01, 2018;3(8):739-748. [FREE Full text] [doi: 10.1001/jamacardio.2018.1511] [Medline: 29898219]
3. Hoffman KM, Trawalter S, Axt JR, Oliver MN. Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between Blacks and Whites. Proc Natl Acad Sci U S A. Apr 19, 2016;113(16):4296-4301. [FREE Full text] [doi: 10.1073/pnas.1516047113] [Medline: 27044069]
4. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. Oct 25, 2019;366(6464):447-453. [FREE Full text] [doi: 10.1126/science.aax2342] [Medline: 31649194]
5. Seyyed-Kalantari L, Zhang H, McDermott MB, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. Dec 2021;27(12):2176-2182. [FREE Full text] [doi: 10.1038/s41591-021-01595-0] [Medline: 34893776]
6. Daneshjou R, Vodrahalli K, Novoa RA, Jenkins M, Liang W, Rotemberg V, et al. Disparities in dermatology AI performance on a diverse, curated clinical image set. Sci Adv. Aug 12, 2022;8(32):eabq6147. [FREE Full text] [doi: 10.1126/sciadv.abq6147] [Medline: 35960806]
7. Henry NR, Hanson AC, Schulte PJ, Warner NS, Manento MN, Weister TJ, et al. Disparities in hypoxemia detection by pulse oximetry across self-identified racial groups and associations with clinical outcomes. Crit Care Med. Feb 01, 2022;50(2):204-211. [FREE Full text] [doi: 10.1097/CCM.0000000000005394] [Medline: 35100193]
8. Rajkomar A, Hardt M, Howell MD, Corrado G, Chin MH. Ensuring fairness in machine learning to advance health equity. Ann Intern Med. Dec 18, 2018;169(12):866-872. [FREE Full text] [doi: 10.7326/M18-1990] [Medline: 30508424]
9. Khera R, Simon MA, Ross JS. Automation bias and assistive AI: risk of harm from AI-driven clinical decision support. JAMA. Dec 19, 2023;330(23):2255-2257. [doi: 10.1001/jama.2023.22557] [Medline: 38112824]
10. Responsible AI guidance. Coalition for Health AI. URL: https://www.chai.org/workgroup/responsible-ai/blueprint-for-trustworthy-ai [accessed 2026-11-13]
11. Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. 2023. URL: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf [accessed 2026-11-13]
12. Relling MV, Evans WE. Pharmacogenomics in the clinic. Nature. Oct 15, 2015;526(7573):343-350. [FREE Full text] [doi: 10.1038/nature15817] [Medline: 26469045]
13. Inker LA, Eneanya ND, Coresh J, Tighiouart H, Wang D, Sang Y, et al. New creatinine- and cystatin C-based equations to estimate GFR without race. N Engl J Med. Nov 04, 2021;385(19):1737-1749. [FREE Full text] [doi: 10.1056/NEJMoa2102953] [Medline: 34554658]
14. Miller WG, Kaufman HW, Levey AS, Straseski JA, Wilhelms KW, Yu HY, et al. National Kidney Foundation laboratory engagement working group recommendations for implementing the CKD-EPI 2021 race-free equations for estimated glomerular filtration rate: practical guidance for clinical laboratories. Clin Chem. Mar 31, 2022;68(4):511-520. [FREE Full text] [doi: 10.1093/clinchem/hvab278] [Medline: 34918062]
15. Braun L. Race correction and spirometry: why history matters. Chest. Apr 2021;159(4):1670-1675. [doi: 10.1016/j.chest.2020.10.046]
16. Kleinberg J, Mullainathan S, Raghavan M. Inherent trade-offs in the fair determination of risk scores. In: 8th Innovations in Theoretical Computer Science Conference (ITCS 2017). Schloss Dagstuhl, Germany. Leibniz-Zentrum für Informatik; 2017:1-43.
17. Hardt M, Price E, Srebro N. Equality of opportunity in supervised learning. In: NIPS'16: Proceedings of the 30th International Conference on Neural Information Processing Systems. New York, NY. Curran Associates Inc; 2016:3323-3331.
18. Nondiscrimination in health programs and activities. Federal Register. 2024. URL: https://www.federalregister.gov/documents/2024/05/06/2024-08711/nondiscrimination-in-health-programs-and-activities [accessed 2026-11-13]
19. La Cava WG, Cohen IG, Aysola J. The future of algorithmic nondiscrimination compliance in the Affordable Care Act. NPJ Digit Med. Dec 10, 2025;9(1):49. [FREE Full text] [doi: 10.1038/s41746-025-02224-7] [Medline: 41372617]
20. Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al. Model cards for model reporting. In: FAT* '19: Proceedings of the Conference on Fairness, Accountability, and Transparency. New York, NY. Association for Computing Machinery; 2019:220-229.

XSL•FO
RenderX

21. Gebru T, Morgenstern J, Vecchione B, Vaughan JW, Wallach H, Daumé HIII, et al. Datasheets for datasets. Commun ACM. Dec 2021;64(12):86-92. [doi: 10.1145/3458723]
22. Omiye JA, Lester JC, Spichak S, Rotemberg V, Daneshjou R. Large language models propagate race-based medicine. NPJ Digit Med. Oct 20, 2023;6(1):195. [FREE Full text] [doi: 10.1038/s41746-023-00939-z] [Medline: 37864012]
23. Notice of vacatur regarding certain provisions of the 2024 nondiscrimination in health programs and activities final rule. Federal Register. 2026. URL: https://www.federalregister.gov/documents/2026/06/02/2026-11015/notice-of-vacatur-regarding-certain-provisions-of-the-2024-nondiscrimination-in-health-programs-and [accessed 2026-11-13]
24. Health data, technology, and interoperability: certification program updates, algorithm transparency, and information sharing. Federal Register. 2024. URL: https://www.federalregister.gov/documents/2024/01/09/2023-28857/health-data-technology-and-interoperability-certification-program-updates-algorithm-transparency-and [accessed 2026-11-13]
25. Health data, technology, and interoperability: ONC deregulatory actions to unleash prosperity (HTI-5) final rule overview. Office of the National Coordinator for Health Information Technology. 2026. URL: https://healthit.gov/resources/health-data-technology-and-interoperability-onc-deregulatory-actions-to-unleash-prosperity-hti-5-final-rule/ [accessed 2026-11-13]
26. Medicare program; Hospital Inpatient Prospective Payment Systems for acute care hospitals (IPPS) and the long-term care hospital prospective payment system and policy changes and fiscal year (FY) 2026 rates; changes to the FY 2025 IPPS rates due to court decision; requirements for quality programs; and other policy changes; health data, technology, and interoperability: electronic prescribing, real-time prescription benefit and electronic prior authorization. Federal Register. 2025. URL: https://www.federalregister.gov/documents/2025/08/04/2025-14681/medicare-program-hospital-inpatient-prospective-payment-systems-for-acute-care-hospitals-ipps-and [accessed 2026-11-13]
27. Artificial intelligence-enabled device software functions: lifecycle management and marketing submission recommendations. U.S. Food & Drug Administration. 2025. URL: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing [accessed 2026-11-13]
28. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance). European Union. 2024. URL: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng [accessed 2026-11-13]
29. Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI) (Text with EEA relevance). European Union. 2026. URL: https://eur-lex.europa.eu/eli/reg/2026/1744/oj [accessed 2026-11-16]
30. Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment. European Commission. 2020. URL: https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment [accessed 2026-11-13]
31. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance). European Union. 2016. URL: https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng [accessed 2026-11-13]

Abbreviations

ALTAI: Assessment List for Trustworthy AI
CHAI: Coalition for Health AI
CMS: Centers for Medicare & Medicaid Services
DSI: decision support intervention
EU: European Union
FDA: Food and Drug Administration
HTI-1: Health Data, Technology, and Interoperability
HTI-5: Health Data, Technology, and Interoperability: ONC Deregulatory Actions to Unleash Prosperity
LLM: large language model
NIST: National Institute of Standards and Technology
ONC: Office of the National Coordinator for Health Information Technology

XSL•FO
RenderX

Edited by G Nicora; submitted 31.May.2026; peer-reviewed by M Honarmand, S Poddutoori, M Singhal; comments to author 10.Aug.2026; revised version received 07.Nov.2026; accepted 10.Nov.2026; published 26.Nov.2026

Please cite as:
Runquist M
Four Imperatives for US Health Systems Governing Racial Bias in Health Care AI
JMIR AI 2026;5:e103163
URL: https://ai.jmir.org/2026/1/e103163
doi: 10.2196/103163
PMID:

©Mike Runquist. Originally published in JMIR AI (https://ai.jmir.org), 26.Nov.2026. This is an open-access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in JMIR AI, is properly cited. The complete bibliographic information, a link to the original publication on https://www.ai.jmir.org/, as well as this copyright and license information must be included.

XSL•FO
RenderX