The Future of Algorithmic Nondiscrimination Compliance in the Affordable Care Act

William G. La Cava, PhD[1]*; I. Glenn Cohen, JD[2]; Jaya Aysola, MD, MPH[3]

[1] Computational Health Informatics Program, Boston Children's Hospital, Boston, Massachusetts, USA
Department of Pediatrics, Harvard Medical School, Boston, Massachusetts
*Correspondence: william.lacava@childrens.harvard.edu

[2] Harvard Law School, Cambridge, Massachusetts
Petrie-Flom Center for Health Law Policy, Biotechnology, and Bioethics, Harvard University, Cambridge, Massachusetts, USA

[3] Leonard Davis Institute of Health Economics, University of Pennsylvania
Department of Medicine, Perelman School of Medicine, University of Pennsylvania
Penn Medicine Center for Health Equity Advancement, Philadelphia, Pennsylvania, USA

Preprint version. In press here:
La Cava, W.G., Cohen, I.G. & Aysola, J. The future of algorithmic nondiscrimination compliance in the affordable care act. *npj Digit. Med.* (2025). https://doi.org/10.1038/s41746-025-02224-7

Abstract

A new Section 1557 rule bans discrimination by AI-based clinical decision tools, with compliance required by May 2025. This paper explores challenges in identifying and mitigating algorithmic bias, especially where outcome disparities exist. We emphasize the need to audit high-risk tools, address proxy discrimination, and provide standardized guidance. Political uncertainty around enforcement complicates long-term planning, making expanded regulatory support essential for health systems and developers alike.

Introduction

On May 6, 2024, the Office of Civil Rights (OCR) and the Department of Health and Human Services (HHS) issued a final rule regarding section 1557 of the Affordable Care Act that went into effect on July 5, 2024. In Section § 92.210 ***(Figure 1***), the Rule states that any patient care decision support tools, including artificial intelligence (AI)-based decision support interventions (DSIs), that exhibit discrimination based on several protected attributes, are prohibited and subject to enforcement under existing federal law[1]. This final Rule signals a recognition of the potential of AI to exacerbate disparities in health care and health outcomes.

Here, we discuss the challenges that algorithm designers and health care organizations face to comply with § 92.210, required by May 1, 2025. These challenges include conflicting views on what is considered discrimination; handling combinations of protected attributes; and factors to consider when prioritizing which tools to audit at scale. While the Rule remains in place for now, we discuss how political landscapes and shifting priorities may influence future policy and compliance requirements.

By requiring covered entities to regularly audit DSIs that directly use protected attributes, rulemakers assume that such tools have the highest potential to discriminate. Yet as prior works have pointed out, AI models are most likely to display a more pernicious form of *proxy discrimination* coined "rational" discrimination, in which a model discriminates through the use variables correlated with protected attributes, rather than the attributes themselves. To identify these cases, we advocate for prioritizing the auditing of tools that intervene on outcomes with well-documented disparities, since outcome disparities are a primary driver of DSI performance differentials, which in turn drive disparate health outcomes. We discuss several real-world examples of this and walk through a detailed modeling simulation to illustrate the various challenges that arise in interpreting and complying with the Rule. Finally, we propose the development of a standardized implementation guide to support a diverse set of health systems in coming into compliance.

§ 92.210 Nondiscrimination in the use of patient care decision support tools.

(a) *General prohibition.* A covered entity must not discriminate on the basis of race, color, national origin, sex, age, or disability in its health programs or activities through the use of patient care decision support tools.

(b) *Identification of risk*. A covered entity has an ongoing duty to make reasonable efforts to identify uses of patient care decision support tools in its health programs or activities that employ input variables or factors that measure race, color, national origin, sex, age, or disability.

(c) *Mitigation of risk*. For each patient care decision support tool identified in paragraph (b) of this section, a covered entity must make reasonable efforts to mitigate the risk of discrimination resulting from the tool's use in its health programs or activities.

Figure 1: *Section 92.210 of "Nondiscrimination in Health Programs and Activities",[1]the final rule regarding Section 1557 of the Affordable Care Act.*

Implications of the Rule in Clinical Care

DSIs are commonly used to stratify patient risk. For example, one DSI might estimate the likelihood of cancer from clinical history and lab results to guide biopsy decisions, while another may predict whether an ED patient will require inpatient care, helping to manage patient flow.[2,3]

Section 92.101(a) of the Rule prohibits excluding individuals from participation, denying benefits, or otherwise discriminating based on race, color, national origin, sex, age, or disability. In practice, DSIs can benefit patients by: (1) identifying high-risk patients who need care, and (2) identifying low-risk patients who can avoid unnecessary interventions. Errors in either direction may constitute a denial of benefit, as patients should ideally receive only medically necessary treatment, and may otherwise suffer harms.

However, when outcome rates (e.g., cancer, hospitalization) differ across groups, accuracy trade-offs between subgroups are often unavoidable. These trade-offs complicate the question of what counts as discrimination under the Rule, and more guidance is needed to navigate them.

Disparities in health care and outcomes are well documented.[4-6] For example, Black and Hispanic ED patients often face longer wait times than White patients with similar urgency,[5] and cancer incidence increases 40-fold from age 20 to 60.[6] A miscalibrated DSI might discharge a patient who needs care—leading to harm—or admit one unnecessarily, exposing them to medical risks or costs. In prostate cancer screening, DSIs may over- or under-estimate risk, leading to overtreatment or missed diagnoses.[2]

Training DSIs on historical data requires deciding whether to calibrate to group-specific prevalence or treat observed differences as artifacts of biased data. This may involve resampling or applying fairness frameworks beyond the scope of this discussion.[7] Here, we focus on scenarios where group-level calibration is necessary to ensure accuracy on all groups.

In short, the Rule raises challenges when: (1) the predicted outcome occurs at different rates across subgroups, and (2) the harms of DSI errors are two-sided—e.g. either undertreating or overtreating patients.

Nondiscriminatory or least discriminatory? Inherent Tradeoffs in Tool Metrics

A tool's false positive rate (FPR) reflects how often it recommends unnecessary treatment; its false negative rate (FNR) captures how often necessary care is missed. A fairness standard sometimes called *equalized odds* or *error rate balance* requires that FPR and FNR be equal across protected groups.[8]

However, care providers often focus on a different metric: *positive predictive value* (PPV)—the probability that a flagged patient will experience the outcome. PPV, closely related to calibration, helps determine how many interventions are needed to find a true positive, which is useful in managing clinical workflows. Requiring equal PPV across groups supports equal treatment for individuals with similar model-estimated risk.

Unfortunately, equalizing PPV and error rates is mathematically impossible when outcome prevalence differs across groups, as is common due to structural inequities. This limitation cannot be resolved by training group-specific models, collecting more data, or using statistical adjustments. In a calibrated model, outcome prevalence constrains the tradeoff between FPR and FNR.[9,10] Therefore, fixing PPV within a group sets a linear relationship between false positives and false negatives based on that group's prevalence.

We illustrate this principle through a simple example in ***Figure 2***, detailed in the Supplement. Consider a DSI predicting a clinical outcome *y* using a clinical risk factor *x* and group membership (*a* or *b*). Suppose groups are equal in size (N = 50,000) and have the same distribution of *x*, but group *a* has 20% outcome prevalence, and group *b* has 50%. In practice we simulate this scenario by sampling *x* from a normal distribution with zero mean and unit variance (for both groups) and sampling the outcome as $y \sim x > t_g$, where t_g is a group-specific threshold that achieves the simulated group prevalences (***Figure 2.A***).

To equalize PPV, we can set different probability thresholds for each group. However, doing so results in a higher FNR for group *a* (lower prevalence) and a higher FPR for group *b*. Alternatively, if we equalize FPR and FNR, PPV will be lower for group *a* due to its lower outcome rate (***Figure 2.B***).

Each metric difference could plausibly be viewed as unfair, yet not all can be equalized simultaneously. This makes it essential for developers and providers to define the tool's intended benefit and adopt a fairness definition aligned with clinical priorities. For instance, if false negatives carry greater harm, minimizing FNR among subgroups should take precedence—even if it results in higher FPR or lower PPV in some groups.

Still, given these unavoidable tradeoffs, developers may fear any design choice could violate the Rule. As philosophers note, "ought implies can": the non-discrimination provisions imply a non-discriminatory design available to the developer that often is non-existent. Clearer guidance on which types of errors constitute actionable discrimination would support developers in making principled, compliant design decisions.

Trade-off Transparency

Given the impact of DSI design choices, algorithm developers should clearly communicate trade-offs and collaborate with care providers to assess potential harms and benefits. In its commentary on the final Rule, HHS declined to mandate patient disclosure of training data or

model inputs, citing cost and current practices. However, it acknowledged that a new ONC rule will require developers of certified health IT to publicly share summary information on risk management practices for predictive DSIs.[11] HHS characterized patient disclosure as a "best practice" but stopped short of requiring it. Clearer guidance on transparency expectations would help developers and providers meet the Rule's non-discrimination requirements.

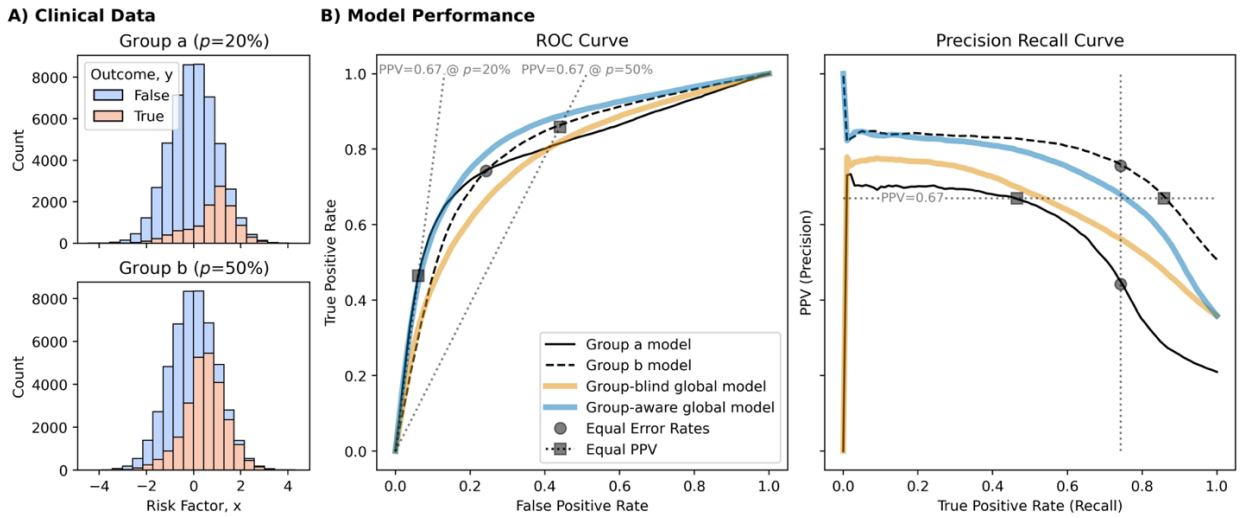

Figure 2*: An illustration of the inherent trade-offs in error rate balance and positive predictive value (PPV) for risk models when the prevalence,* p*, of the outcome,* y*, differs among patient subgroups (*a *and* b*). **A)** The groups have identical distributions of the risk factor, x, and identical sizes (*N*=10, 000), but vary in prevalence of the outcome,* y*. **B)** The receiver operating characteristic curve (left) and precision recall curve (right) for simple logistic regression models under different modeling assumptions. The marked dots illustrate how balancing error rates results in different PPV among the groups, and the marked squares illustrate how balanced PPV leads to different error rates. The lightly dotted lines on the ROC curve indicate how equal levels of PPV relate to True and False positive rates, and how those line levels depend on the prevalence of the outcome in the groups.*

Challenges to Scalability

Handling Combinations of Protected Attributes

Section 92.101a) of the Rule further clarifies that discrimination based on "any combination of" protected attributes should be accounted for, addressing the fact that discrimination along singular attributes does not sufficiently capture discrimination on intersecting subgroups.[12] One must ask, then, how feasible it is to both a) robustly measure discrimination on any combination of protected attributes, and b) address that unfairness through model changes. Algorithm designers should approach both tasks with caution when the number of intersecting groups is

large. Even if we consider only two options for each of the six protected attributes covered under the Rule, this gives 720 potential subpopulations that could be audited for discrimination.

As a concrete example, we look at 198,823 patient admissions from the ED at Beth Israel Deaconess Medical Center from 2011-2019.[13] We group the admissions by four protected attributes: sex, race, age, and language preference (a proxy for national origin) in ***Figure 3***, resulting in 270 groupings. Using a statistical power calculation, we estimate how many samples we would need to robustly detect a disparity between groups (detailed in Supplement). To detect the difference in outcome prevalence from our illustrative example (20% versus 50%), we need at least 175 samples per group, which is possible in all single-attribute groups, but only about 65% of four-way intersecting groups. To detect a smaller difference - say, 40% in one group versus 50% in another – requires more samples; in this case only 20% of the full, four-way intersections of the protected attributes have sufficient sample size. We stress that our starting cohort is large, covering 8 years of visits to a large urban medical center, highlighting the challenge facing algorithm designers monitoring discrimination at smaller centers or over shorter time scales.

In addition to algorithmic challenges related to identifying and removing bias among intersectional groups, the ability to successfully litigate intersectional discrimination is hampered by the piecemeal set of anti-discrimination laws that Section 1557 relies on for enforcement, which often pertain to only one protected attribute (e.g. sex, disability) and are supported by their own case law[14].

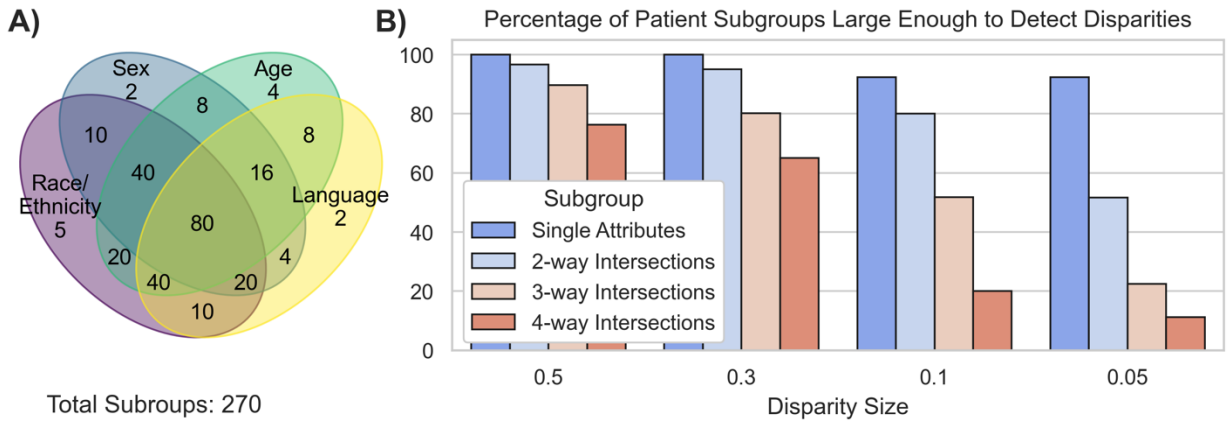

Figure 3: A) *Venn diagram showing the number of distinct patient subgroups when considering multiple, intersecting attributes (Race/Ethnicity, Sex, Age, Language) in the MIMIC-IV ED dataset.[13]* ***B)*** *In order to detect disparities faithfully, a minimum number of patients must be observed in each subgroup. Bars indicate the percentage of 1-,2-,3-, and 4-way intersecting subgroups in MIMIC-IV ED that are large enough to detect a given disparity effect size. As the size of the disparity decreases, so too does the number of patient subgroups that are large enough to power statistical analysis. For further details, see Supplement.*

Choosing the Right Tools to Audit

Commenters noted the challenge of auditing every DSI in large health systems for discrimination. The final Rule clarifies that the regulation (***Figure 1***) requires regular monitoring only of those tools that use protected attributes as input and for each such tool, make reasonable efforts to mitigate discrimination from its use.

This language assumes that discrimination primarily arises from tools that use protected attributes. But if audits focus only on such tools, would they catch the most prominent cases of algorithmic bias? Sometimes, yes—as with kidney function models that explicitly include race.[15] But many widely cited examples of bias involve tools that do *not* use protected attributes.

For instance, a widely used algorithm to identify high-risk patients for care management excluded race but relied on a proxy outcome variable that was racially correlated, producing significant bias.[16] Similarly, chest X-ray AI models have underperformed for underserved populations due to prevalence differences, even without using protected attributes.[17]

In short, limiting audits to tools that include protected variables is likely to miss many biased systems, especially those that discriminate via proxies or confounding. While the Rule acknowledges this possibility, it does not require broader, regular audits that would be more effective at catching such cases.

Although moving away from race-based medicine is a valuable goal, simply removing protected attributes from DSIs does not eliminate discrimination,[19] nor are such attributes the root cause of many recent examples of algorithmic bias.[18]

Addressing Proxy Discrimination

U.S. anti-discrimination law distinguishes between *disparate treatment* (intentional differential treatment based on protected characteristics) and *disparate impact* (facially neutral practices that disproportionately affect protected groups). Most DSI-related litigation is likely to fall under the latter, which does not require proof of intent.[20]

Proxy discrimination is a form of disparate impact in which a model uses variables correlated with protected characteristics, rather than the characteristics themselves. In their seminal article on the subject, Prince and Schwarcz treat redlining, the historical practice wherein financial firms refused to serve predominantly African American geographic regions, as the paradigmatic version of the phenomenon because “the disparate impact it produced was by design: The usefulness to firms of refusing to serve redlined geographic regions was that it allowed them to covertly achieve their discriminatory aims.” [21] While early cases like redlining involved intentional proxy discrimination, intent is not required if the model’s utility depends in part on the disparate impact.[21] This occurs in the AI context when protected attributes provide predictive value

beyond other available risk factors—such as in our earlier example where adding protected attributes improved predictive performance beyond clinical indicators.

This issue is especially challenging when group differences in outcome prevalence cannot be explained by observable social, behavioral, environmental, or biological factors. For instance, disparities in prostate cancer outcomes or ED wait times persist even after accounting for known risk factors.[4,22] In such cases, minimizing disparate impact may require explicitly incorporating protected attributes—a strategy known as *fairness through awareness*.[23] Using group-specific thresholds, as previously described, is one such approach.

Because AI tools naturally pick up on statistical associations in training data, proxy discrimination is often an expected byproduct of their use,[20] yet it remains one of the hardest forms of discrimination to detect and litigate. Although some commenters urged the final Rule to address proxy discrimination explicitly, HHS instead opted to evaluate such claims case by case.

Requiring providers to audit DSIs that use protected attributes could pressure developers to exclude those variables entirely. Given the regulatory burden of ongoing monitoring, providers may favor tools that avoid protected attributes—even if they perform worse or yield less equitable outcomes. Ironically, such oversight may discourage the use of tools designed to reduce algorithmic bias.

Open Issues for Provider Compliance

Whereas prior ONC rulemaking[11] defined requirements for health IT developers, the Rule complements these by making clear that health providers are responsible for any discrimination resulting from the use of any DSIs, even those developed by third parties. To meet these requirements, providers will need to develop assurance processes, for both internal and/or external developers, to ensure that DSIs are compliant with the Rule. These processes ideally ensure that providers arrive at the least discriminatory algorithm choice among potential options. One way of thinking about this is that when it is shown that a facially-neutral practice has disproportionately impacted members of a protected group, the burden then shifts for third-party developers to establish their own processes for providing non-discrimination assurances to the providers making use of their algorithms.[21]

As the Rule stops short of providing comprehensive safe harbor compliance options for all situations in lieu of case-by-case assessments, additional guidance will be crucial to scaling the requirements appropriately to small, rural hospitals and large, multi-center health systems. In its cost estimates, OCR assumes all an organization’s DSIs can be reviewed for discrimination in one hour. Given the complexity of reviewing even a single DSI (consider that it took 10 months to review a two-variable risk model for kidney function[24]), this estimate is likely off by some orders of magnitude.

To support diverse healthcare systems requires the availability of affordable technical assistance to all healthcare organizations.[25] CMS or a similar body could support this need by convening a multidisciplinary task to develop a standardized “Implementation Guide” (IG). This IG could provide both instructive guidance and legal protection, ensuring a clearer compliance path for organizations of all sizes. This work could build upon burgeoning efforts by industry and academic coalitions to standardize DSI development, testing and evaluation frameworks.[26] It might also be possible for OCR to announce a “safe harbor” for compliance in this space, although it is less clear what deference courts would give to OCR’s interpretation given recent changes in administrative law. An implementation guide that adapted compliance processes for resource-constrained providers would be especially helpful, and could, for example, set expectations around the frequency and scope of algorithm audits to match organizational capacities.

Fate of the rule under the new administration

The Rule remains in force. Although the Congressional Review Act allows Congress to nullify late-term agency rules via joint resolution, this Rule falls outside the Act’s “lookback” period and cannot be overturned that way. To modify or rescind it, HHS would need to engage in the same notice-and-comment rulemaking process required to establish it under the Administrative Procedure Act. Since the original Rule took nearly two years from draft to finalization, it is likely to remain in effect for an extended period.

District courts in four states have heard challenges to the Rule, with appeals pending.[27-30] So far, injunctions have applied only to the provision extending sex discrimination protections to transgender individuals. The Rule’s broader anti-discrimination provisions, including those related to DSIs and other protected classes, remain unchallenged.

The impact of the Rule will depend heavily on enforcement. HHS enforcement actions are one way the Rule may curb discriminatory AI. In its supporting materials, OCR outlines how it may interpret compliance[6] — these “rules of the road” can be revised without formal rulemaking, allowing future administrations to shift enforcement priorities without changing the Rule itself. For example, a new administration could redefine compliance expectations and alter enforcement practices, signaling different priorities to covered entities.

Still, as the Biden HHS noted, courts have upheld “private rights of action” under Section 1557, allowing individuals—not just HHS—to sue for violations.[7] That said, whether private suits will be common, how they will proceed, and whether they will offer meaningful enforcement if HHS scales back its efforts remains uncertain.[31]

Final Remarks

In summary, where outcome disparities exist across groups, algorithm developers and providers must assess trade-offs between FPR, FNR, and PPV/NPV across subgroups to evaluate potential discrimination. Developers should clarify how these metrics relate to the tool's expected benefit, enabling providers to assess equity relative to alternative approaches.

Because disparities in outcome prevalence increase the risk of performance differences, focusing solely on tools that explicitly use protected attributes may be less effective than targeting those that affect outcomes with known disparities. Further work would need to clearly specify "known disparities" to operationalize this recommendation but could build upon other federal efforts that regularly quantify health disparities[32] and leverage subgroup metric reporting requirements that arise via ONC rulemaking or as part of future compliance guidelines. As compliance processes develop, both providers and developers would benefit from standardized implementation guidance to navigate the complex trade-offs required to ensure nondiscrimination.

Acknowledgments

This work was supported by the National Institutes of Health under Award R01LM014300. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

References

1. Department of Health and Human Services. *Nondiscrimination in Health Programs and Activities (FR Vol. 89, No. 88)*. Vol 88.; 2024:37522-37703. https://www.federalregister.gov/documents/2024/05/06/2024-08711/nondiscrimination-in-health-programs-and-activities

2. Borre M. Prostate Cancer Screening—The Need for and Clinical Relevance of Decision Analytical Models. *JAMA Netw Open*. 2021;4(3):e212182. doi:10.1001/jamanetworkopen.2021.2182

3. Barak-Corren Y, Chaudhari P, Perniciaro J, Waltzman M, Fine AM, Reis BY. Prediction across healthcare settings: a case study in predicting emergency department disposition. *Npj Digit Med*. 2021;4(1):1-7. doi:10.1038/s41746-021-00537-x

4. Underwood W III. Racial Regional Variations in Prostate Cancer Survival Must Be Viewed in the Context of Overall Racial Disparities in Prostate Cancer. *JAMA Netw Open*. 2020;3(3):e201854. doi:10.1001/jamanetworkopen.2020.1854

5. Aysola J, Clapp JT, Sullivan P, et al. Understanding Contributors to Racial/Ethnic Disparities

in Emergency Department Throughput Times: a Sequential Mixed Methods Analysis. *J Gen Intern Med*. 2022;37(2):341-350. doi:10.1007/s11606-021-07028-5

6. National Cancer Institute. Risk Factors: Age and Cancer Risk. April 29, 2015. Accessed March 5, 2025. https://www.cancer.gov/about-cancer/causes-prevention/risk/age

7. Foulds JR, Pan S. Are Parity-Based Notions of AI Fairness Desirable? *Data Eng*. Published online 2020:51.

8. Hardt M, Price E, Price E, Srebro N. Equality of Opportunity in Supervised Learning. In: Lee DD, Sugiyama M, Luxburg UV, Guyon I, Garnett R, eds. *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc.; 2016:3315-3323. Accessed July 15, 2020. http://papers.nips.cc/paper/6374-equality-of-opportunity-in-supervised-learning.pdf

9. Chouldechova A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *ArXiv170300056 Cs Stat*. Published online February 28, 2017. Accessed January 5, 2022. http://arxiv.org/abs/1703.00056

10. Pleiss G, Raghavan M, Wu F, Kleinberg J, Weinberger KQ. On Fairness and Calibration. In: Guyon I, Luxburg UV, Bengio S, et al., eds. *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc.; 2017:5680-5689. Accessed October 6, 2020. http://papers.nips.cc/paper/7151-on-fairness-and-calibration.pdf

11. Office of the National Coordinator for Health Information Technology (ONC). *Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing*.; 2024. Accessed April 2, 2025. https://www.federalregister.gov/documents/2024/01/09/2023-28857/health-data-technology-and-interoperability-certification-program-updates-algorithm-transparency-and

12. Lett E, La Cava WG. Translating intersectionality to fair machine learning in health sciences. *Nat Mach Intell*. Published online April 28, 2023:1-4. doi:10.1038/s42256-023-00651-3

13. Johnson A, Bulgarelli L, Pollard T, Celi LA, Mark R, Horng S. MIMIC-IV-ED. Published online 2021. doi:10.13026/77Z6-9W59

14. Wójcik-Suffia MA. Algorithmic Discrimination in M-Health: Rethinking the US Nondiscrimination Legal Framework Through the Lens of Intersectionality. *BioLaw J - Riv BioDiritto*. 2024;(1):367-388. doi:10.15168/2284-4503-3002

15. Diao JA, Inker LA, Levey AS, Tighiouart H, Powe NR, Manrai AK. In Search of a Better Equation — Performance and Equity in Estimates of Kidney Function. *N Engl J Med*. 2021;384(5):396-399. doi:10.1056/NEJMp2028243

16. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366(6464):447-453. doi:10.1126/science.aax2342

17. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nat Med*. 2021;27(12):2176-2182. doi:10.1038/s41591-021-01595-0

18. Roberts DE. Abolish race correction. *The Lancet*. 2021;397(10268):17-18. doi:10.1016/S0140-6736(20)32716-1

19. Chohlas-Wood A, Coots M, Goel S, Nyarko J. Designing equitable algorithms. *Nat Comput Sci*. 2023;3(7):601-610. doi:10.1038/s43588-023-00485-4

20. Barocas S, Selbst AD. Big Data's Disparate Impact. Published online 2016. doi:10.2139/ssrn.2477899

21. Prince A, Schwarcz D. Proxy Discrimination in the Age of Artificial Intelligence and Big Data. Published online August 5, 2019. Accessed December 2, 2024. https://papers.ssrn.com/abstract=3347959

22. Lu FQ, Hanchate AD, Paasche-Orlow MK. Racial/ethnic disparities in emergency department wait times in the United States, 2013–2017. *Am J Emerg Med*. 2021;47:138-144. doi:10.1016/j.ajem.2021.03.051

23. Dwork C, Hardt M, Pitassi T, Reingold O, Zemel R. Fairness Through Awareness. In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ITCS '12. ACM; 2012:214-226. doi:10.1145/2090236.2090255

24. Delgado C, Baweja M, Crews DC, et al. A Unifying Approach for GFR Estimation: Recommendations of the NKF-ASN Task Force on Reassessing the Inclusion of Race in Diagnosing Kidney Disease. *Am J Kidney Dis*. 2022;79(2):268-288.e1. doi:10.1053/j.ajkd.2021.08.003

25. Mello MM, Roberts JL. Antidiscrimination Law Meets Artificial Intelligence—New Requirements for Health Care Organizations and Insurers. *JAMA Health Forum*. 2024;5(8):e243397. doi:10.1001/jamahealthforum.2024.3397

26. Home. CHAI - Coalition for Health AI. Accessed April 1, 2025. https://chai.org/

27. *State of Tennessee v. Kennedy*.(S.D. Mississippi 2024). https://www.courtlistener.com/docket/68814195/state-of-tennessee-v-kennedy/

28. *State of Florida v. Department of Health and Human Services*.(M.D. Fla. 2024). Accessed March 7, 2025. https://storage.courtlistener.com/recap/gov.uscourts.flmd.427330/gov.uscourts.flmd.427330.1.0.pdf

29. *State of Texas vs. Becerra*.(E.D. Tex. 2024).

30. *State of Missouri v. Becerra*.(District Court, E.D. Missouri 2025). https://www.courtlistener.com/docket/68930399/state-of-missouri-v-becerra/

31. Dawson L, Sobel L, Pestaina K, Kates J, Artiga S, Published AB. The Biden Administration's Final Rule on Section 1557 Non-Discrimination Regulations Under the ACA. KFF. May 15, 2024. Accessed April 2, 2025. https://www.kff.org/affordable-care-act/issue-brief/the-biden-administrations-final-rule-on-section-1557-non-discrimination-regulations-under-the-aca

32. *National Healthcare Quality & Disparities Reports*. Accessed November 13, 2024.https://www.ahrq.gov/research/findings/nhqrdr/index.html