AI Healthcare Politics: An Analysis on How Industry Practices Perpetuate Racism

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Albert Huang

Spring 2025

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Travis Elliott, Department of Engineering and Society

STS Research Paper

The use of algorithms in the healthcare industry is a very delicate process based on identifying current systemic patterns that can have severe ramifications if done incorrectly. The benefits of large scale machine learning models provide a considerable opportunity for the advancement of healthcare systems, but using these tools accurately and indiscriminately proves a significant challenge to the developers of these models. Several studies have been conducted on issues with healthcare algorithms that have been found to perpetrate racial biases due to a lack of careful consideration with training. This can have severe and lasting consequences for patients that are falsely diagnosed from these algorithms. Developers can take steps to diminish the impact of racial biases by taking a more active approach in identifying potential biases and designing algorithms to both mitigate and correct discriminatory practices. Placing a higher emphasis on eliminating both intentional and unintentional racial biases with specific race-conscious guidelines targeting systemic inequalities can create a better healthcare system that accurately diagnoses patients regardless of their demographics. Using a technological politics framework will assist with better understanding the role developers play in the resounding effects of algorithms and how it should shift to handle inequitable decisions in healthcare.

How Can Algorithms Have Politics?

When analyzing the intersection between science, technology, and society (STS), a useful concept is technological politics. Technological politics is an ideology developed by Langdon Winner that argues that technologies are not politically neutral and either reinforce political dynamics or embody political structures themselves. To understand this proposition, Winner defines his respective interpretations of what politics and technology means. Politics refer to the distribution of power, authority, and privilege in a community and interactions between this arrangement. Technology refers to any systematic way of building order in the world. This broader way of interpreting technologies is inclusive of any design that brings order to the chaos of the world and allows us to understand how certain possibilities of ordering human activity help determine societal paths. Winner further identifies two different manners how technological artifacts have politics.

The first is by decision. What Winner means by "decision" is how technological artifacts have the capacity for different logical arrangements that could yield different political results. These political impacts could possibly be unrelated to the intended uses of the technology and furthers the notion that technologies do not have a right setup and are designed and built in ways that further the creator's ideologies. A common example that displays decisive politics in technology is that of Robert Moses's bridges in Long Island, New York. Moses had racial prejudices against people of color which was reflected in the low bridges that he designed for parkways. These bridges were intentionally short to prevent bus travel that was predominantly used by poor, black people. By purposely inhibiting transportation in Long Island to sites like Jones Beach, Moses ingrained racial and class bias into technological artifacts as bridges and displayed the political consequences of technology.

The second is by necessity. What Winner means by "necessity" is how technologies require or are more compatible with certain political arrangements. This understanding of technological politics differs from the previous by investigating technologies that do not reflect flexibility like those that can decisively motivate different political actions, but represent an intrinsic particular form of political life. Adoption of these technical systems gives rise to societal relations that are either centralized vs. decentralized or oppressive vs. liberating. A common framing for this thought is that of a ship, where democracy is infeasible and a captain is needed to commandeer and direct the crew and ship in an authoritarian environment. Winner summarizes the effect of this kind of technology as leading to the "creation and maintenance of a particular set of social conditions as the operating environment of the system" (Winner, 1980).

When evaluating the adoption of technologies, Winner stresses that careful thought is needed on both the form of technology or if the technology is even needed. Technological politics as a whole takes away some of the power of the system the technology is in and gives that power to the technology itself. Winner likens technology to legislative acts or political foundations that "establish a framework for public order that will endure over many generations" (Winner, 1980). In accordance, technology needs to be handled with the same level of attention and care as would be given to rules of society.

Technological politics proves useful for framing our understanding of the issues stemming from algorithms governing healthcare by acting as a bridge between the technical design of these systems and the social, political, and historical structures that shape and are reinforced by them. The design choices that researchers make with technologies like datasets and ML tools impose politics by giving and taking power from groups. These decisions magnify embedded racism in datasets that historically have been skewed. Similar to the low bridges in Long Island, New York, healthcare algorithms are actively perpetuating disparities in medical coverage which widen class divides and worsen quality of life for marginalized groups. Additionally, algorithms further an authoritarian-esque management of healthcare by reducing human interactions in decision-making and narrowing determinations to algorithmic outputs. This output is treated as objectively correct, despite a black-box approach that gives no reasoning on choices which forces consumers to yield to implicit biases placed by developers. Using technological politics to analyze whether algorithms themselves have politics will equip us with adequate tools in order to examine how systems influence decision-making and reinforce existing social inequalities in healthcare.

How Does Racism in Healthcare Algorithms Present?

In the current age of big data, it is now possible for machine learning models to be trained as tools on the vast collection of information generated by the healthcare industry. The possibilities for these artificially intelligent tools are endless: analyzing patient data to detect those at risk for illnesses, determining personalized treatment options, and streamlining healthcare processes. While artificial intelligence has shown promise in a variety of applications in the healthcare field, these advancements are limited by the methodology behind developing these tools along with the data they are aligned on. A model performance's largest limiting factor is the dataset that it is trained on, particularly the trends and patterns that are prevalent in the dataset. These models identify, extract, and learn patterns to make predictions, and with inherently biased data, the predictions will be similarly distorted in manners that are not representative of the model's purpose.

Research and development of medical algorithms have long been built on simplifications and assumptions that then get ingrained in the technology and produce adverse effects. When determining the basis of their algorithms, developers have limited options for publicly available healthcare datasets that are regarded for being well-rounded because of the specificity of tasks, and are forced to face the innate discrimination in niche commercial data (Alberto et al., 2023). Historically, data collection pertaining to healthcare has been filled with holes in information for minority groups. These holes can be attributed to barriers minority groups face when receiving healthcare or the inability of capturing and digitizing relevant health data, whether it be due to the lack of equipment or manpower. Some barriers that minority groups face that prevent access to healthcare are economic inequality, lack of insurance, or geographic limitations which lead to an inability to get the quality care they need and consequently generate crucial medical data. Additionally, documented trials in the past have historically excluded underrepresented groups leading to ungeneralizable data that prioritizes a white population. Another reason that deters certain demographics from receiving and spending on healthcare is discrimination at the hands of providers and a distrust for the healthcare system (Chen et al., 2021). Repeated mistreatment at the hands of doctors has led to certain superstitions developing in disadvantaged groups that develop into full-blown stigmas against healthcare. An example would be the Tuskegee Syphilis Study, which unethically used poor black men to test the natural progression of untreated syphilis and resulted in severe health complications and death, ruining the trust between African Americans and medical practitioners.

Medical algorithms additionally make use of race-based adjustments due to assumptions and stereotypes about biological differences between racial groups. By incorporating race as a variable, developers aim to provide different medical assessments that best fit patients of different racial backgrounds. This design choice often backfires, however, and instead widens disparities in healthcare. An example of the effects of race-based adjustments is how black Americans are systematically undertreated for pain relative to white Americans. In studies conducted by Hoffman et al. (2016), researchers found through studies that white medical students and residents that believed that there were biological differences between races underrated a black patient's pain level when compared to a white patient and consequently would make less accurate treatment recommendations. These beliefs were attributed to stereotypes like "black people have thicker skin than white people" which then led to poor treatment. Another instance of race-based adjustment's effects is how an algorithm that determined whether an NFL player suffered brain injuries was adjusted to assume black men had lower cognitive abilities. This meant that retired Black players had to meet a lower threshold score than their white counterparts to be diagnosed with a neurodegenerative disease, effectively denying them settlement payments ranging from hundreds of thousands to millions of dollars (Associated Press, 2021). Race-based adjustments along with compromised data creates a cycle of self-feeding racially-biased treatment that generates even more flawed data.

When developing and training the models, developers also can inadvertently use wrong metrics or goals to train data. By using certain features or labels to frame a problem the algorithm is trying to optimize, the optimal policy could be built on assumptions that are inherently racist. This possibility is due to algorithms being black-boxes where researchers are unable to determine what features a model uses and how to make predictions. An example would be an algorithm training on patient locations to predict estimated costs, even though it is not pertinent to the task if the desired process is based on symptoms. An algorithm's logic relies on statistical predictions derived from detected patterns, regardless of their origins, as long as they produce accurate results. The black-box nature of algorithms is not a design choice from researchers though, as the technology for interpretability of models does not exist yet and developers are left to generalize decision-making to the best of their knowledge.

The Current State of Racism in Healthcare

This issue has been exacerbated to the point the term, "Health Data Poverty", was coined as the phenomenon where individuals or groups who are underrepresented in health datasets are less able to benefit from data-driven innovations (Ibrahim et al., 2021). In order to address this phenomenon, legislation has been enacted in an attempt to resolve inequities in artificial intelligence for healthcare. In the United States (US), the Food and Drug Administration (FDA) is responsible for approving commercially marketed medical AI devices. Policies and guide-rails lag behind though, with the explosive growth of technology. Some algorithms like those that predict risk of mortality, likelihood of readmission, and in-home care needs are not required to be reviewed and regulated by the FDA or even any regulatory body. Instead, the FDA in their Jan. 2021 "Artificial Intelligence/Machine Learning-Based Software as a Medical Device (SaMD) Action Plan" asked developers to state steps used to mitigate bias and for data transparency. Additionally, they outlined general Good Machine Learning Practice (GMLP) and stated their commitment to evaluate software for bias and discrimination. However, this plan focuses on only monitoring effects without outlining standards for development.

The culmination of decades of bad practice was brought to light in a groundbreaking study done in 2019, which found that a commercial risk-prediction algorithm used to identify potential "high-risk care management" program patients required black people to be substantially less healthy than their white counterparts in order to be accepted into the program. The algorithm was used on 200 million people in the United States yearly, identifying patients with a risk threshold score above 97% as high-risk to allocate resources that ranged from dedicated nurses to extra primary care appointment slots. Researchers found that black patients needed to be sicker in health markers like severity of diabetes, high blood pressure, renal failure, cholesterol, and anemia to be placed in the high-risk care management program which otherwise would have been given to healthier white patients. When the algorithmic bias was corrected in a simulated scenario by replacing healthy white patients above the risk threshold with less-healthy black patients until the level of health was relatively the same, researchers found that the fraction of black patients marked as high-risk increased drastically from 17.7% to 46.5% (Obermeyer et al., 2019).

Other cases where healthcare algorithms perpetuate racial biases have recently been identified. In 2019, Dr. Sendak, who was developing an algorithm to predict childhood sepsis, noticed that after three years of data collection, doctors took longer to order blood tests for Hispanic kids that were eventually diagnosed. This could have introduced bias into the algorithm by potentially teaching AI that Hispanic kids develop sepsis slower which could be fatal. The slow tests were possibly the result of doctors needing interpreters and taking illnesses more seriously in white kids than those of Hispanic kids (Levi & Gorenstein, 2023). In 2021, researchers found that the removal of race-based adjustments in estimating glomerular filtration rate to predict chronic kidney disease would result in 3.3 million more black Americans meeting the criteria for Stage 3 chronic kidney disease with some qualifying for referrals to kidney specialists and kidney transplants (Tsai et al., 2021). These cases all display the vast consequences that racial bias in healthcare algorithms can cause.

Analysis Using Technological Politics

As of now, the current course of action surrounding that of racial bias in algorithms is one of passivity. A recurring theme in cases involving algorithmic discrimination is the late detection after damage has already been wrought. The impact of poor healthcare as indicated before can be life-altering and fatal and reinforces generational health disparities, economic inequality, and mistrust in medical institutions among marginalized communities. Technological politics can be utilized to help understand the unassertive nature behind decisions regarding how to best address racially biased algorithms.

Herd Mentality in Healthcare

Healthcare algorithm developers often exhibit a docile attitude toward the political implications of their products, largely due to their lack of awareness of the political power these algorithms wield. This is an effect of the research ideology adopted by developers that focuses wholly on effectiveness rather than additionally considering equitability. From a survey, out of thirteen academic medical centers, only four considered racial bias when developing machine learning algorithms (Levi & Gorenstein, 2023). It is evident that these medical centers did not consider how and why their algorithms should even be adopted in a political manner, leaving effects for the users to deal with and relinquishing any influence they might have.

Research fields have assumed a herd mentality based on assumptions and stereotypes that are historically incorrect. The unconscious biases that researchers have can cause them to mistakenly believe their assessments are wholly objective and to blindly follow the actions of predecessors without questioning the political effects such decisions caused. These actions take the form of heuristics that can generate systematic cognitive/behavioural biases and can be grouped under sets like availability, representativeness, and anchoring/adjustment. Availability heuristics occur when people make judgements based on the most readily available or recent information and ignore potentially useful information. Representativeness heuristics refer to judging processes based on similarity to other processes despite distinctness and incorrect information. Finally, anchoring/adjustment heuristics refer to how people often anchor judgements based on their own or another's opinions (Baddeley, 2015). Researchers may resort to any of these heuristics to aid with the development of technologies without understanding the political effects they may cause.

Heuristics are evident in the case of the risk-prediction algorithm used to identify potential “high-risk care management” program patients presented above. The developers of the algorithm assumed that predicting healthcare needs was equivalent to predicting a patient's future health costs. This assumption, however, caused the algorithm to disregard poor patients that faced barriers for receiving healthcare and consequently spent less on healthcare. This issue was further construed to be based on race, with black patients generating on average $1801 lower costs per year when compared to white patients with the same number of chronic illnesses (Obermeyer et al., 2019). The reasoning of using costs to predict healthcare needs was cited as “reasonable” according to developers since it was the industry-wide approach. The lack of responsibility the developers showed when discrediting the impact their decision had on the performance on different demographics illustrates the problem with underestimating the political sway technology can have. Instead, developers need to accept that technology will always manipulate the political landscape and seize that opportunity to design systems that actively mitigate bias, promote equity, and ensure accountability in healthcare decision-making.

How are Politics and Healthcare Intertwined?

A short-handed method of attempting to address racism in these algorithms is a push for race-blind training. This path, however, essentially turns a blind eye to the issue where people of different racial backgrounds are treated differently by healthcare by pretending that the effects are not a part of the healthcare system and a political issue when they are instead well-established. By ignoring the immediate problem, researchers are once again assuming a passive approach that does not take any steps to resolve racial bias in algorithms with its inherent political power. While this ideology does not aggravate racial biases, it minimizes the attention the issue deserves and allows biased decisions to continue being made.

Those that did take advantage of the political powers of healthcare algorithms did so in ways that perpetuated ill racial biases for personal gain. The case where race-adjustments in algorithms analyzing the cognitive performance of retired football players effectively reduced settlement payouts is one such instance where logical arrangements in technology were used in ways to enhance power over and take advantage of others. Another instance where algorithms are used to take advantage of others is by predatory insurance companies. Insurance companies purposely design technology to target the less fortunate by raising premiums against those it grades as "at risk", perpetuating racial differences. These insurance companies ingest personal information like race to manipulate healthcare costs for discrimination far beyond the algorithms intended use. It was reported that under the Medicare Advantage program, insurers abused "risk scoring" algorithms and overcharged Medicare nearly $30 billion (Christensen, 2021). Understanding that developers have the ability to imbue politics into technology allows a decisive line to be drawn in how to derive and regulate technologies. In line with technological politics, developers need to ask in what form technology should be adopted and how they should be designed for their intended uses.

Algorithms themselves are also designed with the purpose of optimizing care and ridding humans out of certain processes. By necessity, these algorithms have politics aligning themselves with a more centralized and less personal relation between healthcare and people. At an extreme, algorithms lead us down the path towards authoritarianism where one's healthcare is governed by the creators of algorithms. While algorithms could possibly increase efficiency and correct human error, they also strip nuanced customized care from individuals with unique needs. This approach, which reduces humans down to data points, would learn the systemic discrimination that marginalized groups face and continue to implement them through automated decisions. A dystopian future where algorithms purely determine what care one gets is bleak and begs the question of if automated algorithms are worth it in a field personalized like healthcare where human elements and judgement are of utmost importance. In line with technological politics, developers must ask if algorithms are really needed.

The Search for Accountability

The racial biases perpetuated by algorithms are not altogether the developer's fault since there is currently no structural policies or set teachings that require methodical thinking about racial skews in society and their effects. Recently, independent researchers analyzing how the FDA processes and approves devices found a series of issues that could compromise their supposed good performance. Wu et al. (2021) found that models were evaluated at a single site alone. This entailed limited diversity in the FDA's analysis, which was further compromised by small sample sizes where the median evaluation sample consisted of only 300 people. Almost all devices underwent retrospective studies reinforcing a passive review style that only monitored and addressed issues rather than actively fixing them. None of the 54 high-risk devices were analyzed in prospective studies which are needed for full characterization of the impact of the AI decision tool on clinical practice. There is, however, a growing wave for developers to hold accountability: multiple academics and industry leaders have said they want to see the FDA spell out in public guidelines exactly what developers must do to prove their AI tools are unbiased. The shift from previous views held by algorithm developers indicates a change in roles where it is no longer up to consumers to identify issues and address them.

In the current healthcare landscape, if developers do not step up and take advantage of the political power that technology like algorithms have, algorithms will continue to make poor decisions and inflame current racial discrimination in healthcare. Repeated use and reinforcement by medical professionals will further cement biases and continue the cycle of misdiagnoses and suffering that disadvantaged people will face. Developers need to carefully weigh the benefits and downsides to algorithms in certain situations and determine whether they are truly vital. If these algorithms are necessary and have the power to benefit many, developers need to be more active in regards to how they address and resolve racial biases in healthcare with design choices in their algorithms.

Structural Racism

Race has always been central to American politics as a social and power construct and can not be avoided when dealing with politics concerning technologies. The root cause of racial bias in algorithms are existing inequities in care built into policies and thus society. To better understand how to remedy discriminatory practices and improve equity in healthcare, a closer look at the history of racism in healthcare is needed.

Structural racism refers to the "totality of ways in which societies foster racial discrimination through mutually reinforcing systems of housing, education, employment, earnings, benefits, credit, media, health care, and criminal justice." (Fashaw-Walters, 2023) Racism presents itself in society as an interconnected web with structures like slavery and Jim Crow laws institutionalizing racial hierarchy into these systems which have been passed down generationally. Separation of treatment is evident in the 1946 Hospital Survey and Construction Act, which, despite mandating equal access to health care, allowed states to build racially separate and unequal facilities. The systematic oppression in jobs that people of color face has also presented itself in healthcare, with black people and other minorities having higher rates of unemployment and under-representation in good-paying jobs that included health insurance as benefits when compared to white counterparts (Williams & Rucker, 2014). The federal government has even acknowledged that “inadequate health insurance coverage is one of the largest barriers to health care access, and the unequal distribution of coverage contributes to disparities in health. Low-income [minority people] with bad health had 68% less odds of being insured than high-income [White people] with good health” (Yearby et al., 2022). If structural racism continues to shape healthcare access, racial minority populations will continue to suffer inequitable access to healthcare and in turn generate increasing amounts of skewed, unusable data that will have lasting impacts (Yearby et al., 2022).

Additionally, most of the current legislation regarding discrimination has been ineffective because of their focus on intentional discrimination. From a policy standpoint, it is easier to identify and quantify purposeful discrimination. As shown through the risk-prediction algorithm case above though, the research community has curated a tendency for heuristics that instead make use of unconscious biases that even researchers are unaware about. These unconscious biases are the most hurtful concerning racial biases because they are embedded within supposedly neutral systems, making them harder to detect and challenge. To combat this, policymakers must implement general regulations that hold developers accountable for discriminatory algorithms, formally recognizing their political influence and ensuring that bias is penalized regardless of intent.

How Can the Issue at Hand be Addressed?

As determined through analysis with technological politics, policies also should not be race-neutral and ignore race by disregarding effects of racism and targeting even benefits for everyone. Instead, the political power of technology needs to be used to implement race-conscious policies. Racism-conscious policies allow policymakers to address racism by identifying, understanding, and responding to the structural barriers and inequities that give rise to and maintain the social, political, and economic limitations that minority groups face. Following a framework proposed by Fashaw-Walters et al. (2023), policymakers would first examine current inequities, then identify inequity-related policies, then dissect policy mechanisms and consequences, then elucidate the impact of racism, and finally create new racism-conscious health policies that consider implementation strategies.

To address healthcare inequities, greater awareness must also be given to social inequities. There is a need for educational campaigns about the problem of racial inequities in the US and how they may present in healthcare. The awareness levels of the public and professional community, especially the medical community, must be raised regarding the political power they hold regarding technologies. All of this can be done through the recruitment and retention of minority doctors. Minority doctors provide equitable perspectives on racially biased issues their respective demographics could face and are willingly to actively make design choices regarding them.

Responding to Counterarguments

Proponents of current healthcare algorithms argue that they enhance efficiency, conserve resources, and alleviate the workload on healthcare professionals and are effective for the most part. These supporters gloss over the inequitable treatment of minority groups citing how the pros outweigh the cons, failing to acknowledge the immoral act of prioritizing efficiency over people's livelihoods. The utilitarian approach of current algorithms prioritizes the well-being of the majority while overlooking the disproportionate harm inflicted on historically marginalized groups that are repeatedly oppressed. The principle behind healthcare should be for ensuring everyone's good health through equitable benefits regardless of someone's race. By justifying different treatment based on one's skin color, algorithms as they are implemented now uphold the belief that not everyone's life is valued equally. Instead algorithms need to be designed in a manner that balances efficiency with fairness in order to not sacrifice the rights of a few for the benefit of the many.

Researchers may also believe that existing algorithms are already refined, with no better way to measure or determine certain health factors. After all, relying on established methods has proven effective in previous studies. Researchers, however, need to redirect their priorities from purely building a usable algorithm to one that ensures robustness of their methods with no unintended political side effects. Problem formulation is a common dilemma that arises in areas involving the use of data. The tasks that developers are trying to provide a solution for are often amorphous and do not have set indicators or features to make predictions on. There usually is no easy way of determining how to predict a certain indicator or else it would be already utilized. It is then up to developers to engineer how the data will get manipulated and used by the model. This takes significant effort and experimentation, which some researchers are unwilling to resort to. In the previous example of the risk-prediction algorithm, basing the predictor for healthcare needs on healthcare costs was a flawed relationship which led to biased decisions. Instead, to mitigate bias, researchers need to conduct careful feature and label engineering of the dataset along with pre and post-processing to tweak the algorithm performance towards equitable political outcomes, ensuring that no group is systematically disadvantaged (Cary Jr. et al., 2023).

Looking Forward

The use of algorithms in healthcare is the future of medical advancements, so it is in our best interest to ensure that these algorithms provide an unbiased service to all demographics. Addressing current inequities in healthcare and mitigating future inequities is vital to minimizing potential consequences that come from algorithms in healthcare, as issues in this field can pose severe risks for patients that are improperly diagnosed. Technological politics was used in the analysis of current algorithm practices to indicate decisive and necessary political impacts of technology which need to be resolved. Using the political power of technology and implementing race-conscious policies provides a set of paths for developers to use to remediate discriminative processes. Through understanding the progression of structural racism, it is also evident that simply modifying an algorithm's external decision-making is insufficient and there is a need to address the internal relationships within the data by tackling broader societal inequities. Ensuring equitable healthcare algorithms requires leveraging technology's societal influence while also committing to dismantling systemic biases to create fair and inclusive solutions for all. Moving forward, it is essential to promote collaboration between developers, policymakers, and affected communities in the process of creating universal algorithms that address both historical and present-day biases to eventually rid the healthcare system of racial prejudice.

References

Alberto, I. R. I., Alberto, N. R. I., Ghosh, A. K., Jain, B., Jayakumar, S., Martinez-Martin, N., McCague, N., Moukheiber, D., Moukheiber, L., Moukheiber, M., Moukheiber, S., Yaghy, A., Zhang, A., & Celi, L. A. (2023). The impact of commercial health datasets on medical research and health-care algorithms. The Lancet. Digital Health, 5(5), e288–e294. https://doi.org/10.1016/S2589-7500(23)00025-0

Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan | FDA. (2021, January). Www.fda.gov; FDA. https://www.fda.gov/media/145022

Associated Press (n.d.). *Black NFL players call for end of algorithm that assumes Black men have lower cognitive abilities*. MarketWatch. https://www.marketwatch.com/story/retired-black-nfl-players-and-their-families-call-for-race-norming-practice-to-end-01621018741

Baddeley, M. (2015). Herding, social influences and behavioural bias in scientific research. *EMBO Reports*, *16*(8), 902–905. https://doi.org/10.15252/embr.201540637

Cary, M. P., Zink, A., Wei, S., Olson, A., Yan, M., Senior, R., Bessias, S., Kais Gadhoumi, Jean-Pierre, G., Wang, D., Ledbetter, L., Nicoleta Economou-Zavlanos, Obermeyer, Z., & Pencina, M. J. (2023). Mitigating Racial And Ethnic Bias And Advancing Health Equity In Clinical Algorithms: A Scoping Review. *Health Affairs*, *42*(10), 1359–1368. https://doi.org/10.1377/hlthaff.2023.00553

Chen, I. Y., Pierson, E., Rose, S., Joshi, S., Ferryman, K., & Ghassemi, M. (2021). Ethical Machine Learning in Healthcare. *Annual Review of Biomedical Data Science*, *4*(1), 123–144. https://doi.org/10.1146/annurev-biodatasci-092820-114757

Christensen, D. M., Manley, J., & Resendez, J. (2021). Medical Algorithms Are Failing Communities Of Color. *Forefront Group*. https://doi.org/10.1377/forefront.20210903.976632

Fashaw-Walters, S. A., & McGuire, C. (2023). Proposing A Racism-Conscious Approach To Policy Making And Health Care Practices. *Health Affairs*, *42*(10), 1351–1358. https://doi.org/10.1377/hlthaff.2023.00482

Ibrahim, H., Liu, X., Zariffa, N., Morris, A. D., & Denniston, A. K. (2021). Health data poverty: an assailable barrier to equitable digital health care. *The Lancet Digital Health*, *3*(4), e260–e265. https://doi.org/10.1016/s2589-7500(20)30317-4

Levi, R., & Gorenstein, D. (2023, June 6). *AI in medicine needs to be carefully deployed to counter bias – and not entrench it*. NPR. https://www.npr.org/sections/health-shots/2023/06/06/1180314219/artificial-intelligence-racial-bias-health-care

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science*, *366*(6464), 447–453. https://doi.org/10.1126/science.aax2342

Tsai, J. W., Cerdeña, J. P., Goedel, W. C., Asch, W. S., Grubbs, V., Mendu, M. L., & Kaufman, J. S. (2021). Evaluating the Impact and Rationale of Race-Specific Estimations of Kidney Function: Estimations from U.S. NHANES, 2015-2018. *EClinicalMedicine*, *42*, 101197. https://doi.org/10.1016/j.eclinm.2021.101197

Winner, L. (1980). Do Artifacts Have Politics? *Daedalus*, *109*(1), 121–136. http://www.jstor.org/stable/20024652

Williams, D. R., & Rucker, T. D. (2014). Understanding and Addressing Racial Disparities in Health Care. *Health Care Financing Review*, *21*(4), 75. https://pmc.ncbi.nlm.nih.gov/articles/PMC4194634/#sec4

Wu, E., Wu, K., Daneshjou, R., Ouyang, D., Ho, D. E., & Zou, J. (2021). How medical AI devices are evaluated: limitations and recommendations from an analysis of FDA approvals. *Nature Medicine*, *27*(4), 582–584. https://doi.org/10.1038/s41591-021-01312-x

Yearby, R., Clark, B., & Figueroa, J. F. (2022). Structural Racism In Historical And Modern US Health Care Policy. *Health Affairs*, *41*(2), 187–194. https://doi.org/10.1377/hlthaff.2021.01466