Implicit Bias Effects of Artificial Intelligence Implementation in Cardiovascular Care

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

Jashanjeet Matharoo
Spring, 2021

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

Signature Jashanjeet Matharoo Date 05/14/2021
Jashanjeet Matharoo

Approved H Rogers Date _5/4/2021_____
Dr. Hannah Rogers, Department of Engineering and Society

Abstract

Artificial intelligence (AI) is a groundbreaking technology that has become ubiquitous in everyday life in the form of smartphone personal assistants and email spam filters, to name a few and more. However, recent years have revealed biases presented by these AI systems in applications ranging from advertisements during web browsing to hiring practices of international corporations. Preliminary applications in healthcare show a similar trend, raising alarming concerns about the downstream effects on patient outcomes of widespread clinical AI implementation. To probe this possibility, the analysis presented here viewed the issue through the lens of technological momentum and utilized the historical case study method to explore COVID-19 vaccine distribution in the state of Virginia to address the following: How will the implementation of artificial intelligence in cardiovascular medicine directly and indirectly impact the prevalence of medical bias as measured by outcome inequity, and what can be done about it?

The results showed underrepresentation in vaccinated populations of African American (AA) residents of the ten localities with the largest AA populations at rates significantly higher than those of the ten localities with the smallest AA populations and with the highest average income. Further exploration of AI-assisted healthcare resource allocation practices common in industry indicated parallel contexts in both the aforementioned case study and the current landscape of cardiovascular care, highlighting the need for active intervention to avoid the propagation of healthcare disparities by AI in clinical settings. Specifically, higher rates of cardiovascular disease comorbidities in minority communities lead to greater risk for these patient populations, but lower rates of access to quality healthcare for these same communities deceptively indicates lower usage and subsequently lower allocation of healthcare resources by AI systems reliant on historical data.

To stop the vicious cycle of technological momentum in healthcare AI, transparency must be promoted in machine learning algorithms. Additionally, steps taken by physicians to acknowledge social determinants of health, and by insurers and legislators to dissect complex data fueling policy decisions, can promote the ethical implementation of AI in clinical settings. Together, these efforts can give way to a brighter and more innovative future.

Implicit Bias Effects of Artificial Intelligence Implementation in Cardiovascular Care

Introduction

Can a machine be racist, sexist, or homophobic? This question is now at the forefront of the technology industry's collective consciousness and specifically refers to artificial intelligence (AI), a relatively new innovation that is projected to contribute over $15 trillion to the global economy in the next decade (Neiger, 2018). Although AI has become ubiquitous in society in the form of smartphone personal assistants, spam filters on emails, and countless other applications, the technology has displayed a propensity for propagating societal biases. From seemingly benign digressions such as displaying images of predominantly white males as search results for "CEO" to more destructive displays such as racial discrimination in granting parole, the implicit bias held by AI systems is detrimental to various industries and sectors (Mattu, n.d.; Smith, 2020). Though there are multiple potential contributors to AI bias, the likely cause is the training of these software on historic data sets that reflect the aggregate set of prejudices held by members of society. However, evaluating the counter-effect that AI technology has on society is essential to deciding whether novel applications, such as those in healthcare, will promote or deteriorate societal welfare. Moreover, such evaluation is vital to determining more ethical approaches to innovation in AI, laying the framework for how engineers, physicians, and legislators should move forward in this regard. Viewing this issue through the lens of STS scholar Thomas Hughes' Technological Momentum, aided by the use of historical case studies and statistical analyses, offers a framework with which the complete societal impact of AI is analyzed (Hughes, 1994). Ultimately, this analysis provides a response to the following: How will the implementation of artificial intelligence in cardiovascular medicine directly and indirectly impact the prevalence of medical bias as measured by outcome inequity, and what can be done about it?

Background

Artificial Intelligence is Already Biased

The early advances in artificial intelligence (AI) that are used widely today have highlighted potential issues that indicate implicit bias in machines as a possibility. Criticisms of AI often include implicit biases that are "baked in" to systems by the developers that make them, and the data sets they are trained on. For example, image databases will return images of primarily white males when searching "CEO" and resume screening software will discriminate against names aligning with certain ethnicities despite identical applicant qualifications (Smith, 2020).

Also, word recognition software has shown heavy bias to associating words with "male" or "female" following common gender stereotypes (*Biased Bots*, n.d.). These outcomes may be due to the data sets that AI systems are trained on, as these could reflect human decision-making and societal inequities (Lum & Isaac, 2016).

The multi-causal phenomenon of implicit bias in AI has impacts that are tangible across many sectors of society. Online advertisement delivery programs target ads for high-interest credit cards and criminal background checks to Black users (Sweeney, 2013). Facial recognition software ubiquitous in modern devices misidentifies Black women between a quarter and a third of the time (*Study Finds Gender and Skin-Type Bias in Commercial Artificial-Intelligence Systems*, n.d.). Software used to determine whether or not defendants can be granted parole consistently report higher risk rates for Black individuals than White individuals, even if the Black individual has no prior record and the White individual does (Mattu, n.d.). Even household names like Amazon use hiring software that discriminates against racial minorities and women, as reports found that the AI used by the online retail giant perpetuated the 60-75% male composition of the workforce by citing a lack of "fit" as a reason to hire less women (*No Surprise Amazon's AI Was Biased against Women, Says Sandra Wachter - Business Insider*, n.d.).

Implicit bias in AI systems will potentially have a particularly devastating impact in the field of medicine. Specifically, training healthcare AI technologies on past diagnoses that are influenced by provider biases will scale up the effect of harmful stereotypes prevalent in medicine today (Hall et al., 2015). For example, this technology will further exacerbate the unacceptable statistic that Black women are 2-3 times more likely to die from pregnancy-related complications (*Racial and Ethnic Disparities Continue in Pregnancy-Related Deaths | CDC Online Newsroom | CDC*, 2019). Additionally, bias detected by patients reduces physician trust and leads to worse outcomes in the long-term (*How Does Implicit Bias by Physicians Affect Patients' Health Care?*, n.d.). This is potentially due to patients being increasingly perceptive to indications of bias – well-founded or not – and subsequently doubting the physician's consideration of their best interests and being reluctant to follow treatment plans and recommendations. Thus, it is vital to understand the potential implications of implementing AI in various healthcare settings.

Artificial Intelligence and Society Affect One Another

Although it is the prevailing attitudes of a subset of society that influence the actions of AI, the technology itself has an influence on the attitudes of the entire society. With AI specifically, the developers that create these programs are largely White or Asian males and therefore have a collective set of overlapping biases that are translated onto the technology they are developing (Manyika et al., 2019). Additionally, the technology they develop is trained on huge amounts of historical data that often reflect historical status quos in society. However, AI's subsequent effect on society is often overlooked. It is perceived by those without a technical background as completely objective and incapable of prejudice; thus, when inequity as described above is a result of the implementation of AI systems, it is prescribed to fundamental differences or deficiencies in groups of people rather than systemic effects and intersectionality, diverting social attention and resources away from combating these pressing issues.

Given this positive feedback loop solidifying implicit bias and prejudice prevalent in today's society resultant of AI, the technology is best viewed through the lens of technological momentum. This framework is receptive to the social influences on technological artifacts but is also able to prospectively view the technology's counter influence on society (Hughes, 1994). An example of technological momentum lies in the rise of the internet; humans' needs to communicate and drive to innovate led to the internet's creation, which in-turn has transformed behaviors in our daily lives such as shopping and entertainment as well as how society conducts commerce. Critiques of this approach come from proponents of social constructivism and technological determinism, with the former stating that technological artifacts are given meaning and agency by the society that adopts them, and the latter arguing that technology shapes society independent from the actions of that its members (Klein & Kleinman, 2002; Smith, 1994). Additionally, critics cite technological momentum's lack of originality and label the framework derivative of its aforementioned predecessors (Colarossi, n.d.). However, technological momentum lies between social constructivism and technological determinism on the same continuum, and therefore affords the ability to determine how society has shaped artificial intelligence and predict how society will be influenced by artificial intelligence.

Research Question and Methods

In relation to the technical component of this report, the technological momentum framework was used to explore how societal bias has shaped artificial intelligence (AI) technology and how, in turn, that technology has influenced societal bias. Ultimately, this analysis has sought to provide a response to the following: How will the implementation of AI in cardiovascular medicine directly and indirectly impact the prevalence of medical bias as measured by outcome inequity, and what can be done about it?

The method used to approach the above question utilizes the historical method of examining case studies, and the United States' response to the COVID-19 pandemic serves as an optimally representative historical case study to explore the link between society and the AI technology used to allocate healthcare resources. To better elucidate and understand this relationship, background research on healthcare disparities across minority populations and on the healthcare resource allocation algorithm was conducted. Additionally, statistical analysis was conducted on data representing COVID-19 resource allocation across the state of Virginia by county to explore continued disparities and trends. Then, the information gathered through the aforementioned historical case study and statistical analysis was used to predict the effect of AI implementation in cardiovascular medicine on medical bias as measured by patient outcome inequity.

Healthcare Access

The United States healthcare system is a complicated blend between a private, multi-payer and public, single-payer system, and due to this unique set-up brings particular challenges in terms of accessibility to care. To further elucidate this, an overview of the U.S. healthcare system is presented below, followed by a brief review of the disparities prevalent in this system.

The U.S. Healthcare System is Complicated

As described earlier, the U.S. healthcare system is a mix of a multi-payer system and single-payer system, with the latter taking the form of government programs such as Medicare and Medicaid providing healthcare to the elderly and socioeconomically disadvantages, respectively, as well as the Veterans Administration system providing healthcare to military personnel and their families. Different insurers pay different rates to health systems and hospitals for procedures and care, and these rates are negotiated between the payer and the provider. However, rates can vary based on insurance, with private insurance generally being more expensive than federal and state programs like those listed above (De Lew et al., 1992).

Minorities Face Disparities in Access to Care

The healthcare system detailed above oftentimes leaves large groups of individuals uninsured or underinsured, with these groups frequently being ethnic minorities. More specifically, groups such as Latinos and African Americans are uninsured or underinsured at rates multiple times higher than white non-Latino populations. A report published by the UCLA Center for Health Policy Research conducted a detailed exploration of healthcare disparities across these minority groups and found data supporting the above. In particular, the group found that 37% of Latinos and ~25% of African Americans were uninsured in comparison to the 14% of white non-Latinos who are uninsured. Additionally, only 43% of Latinos and 53% of African Americans received job-based health insurance in comparison to 73% of white non-Latinos who did (Brown et al., 2000).

In 2014, the Obama administration passed the Affordable Care Act (ACA), which took steps towards attempting to bridge the healthcare gap and improve disparities for minority populations. Analysis of national healthcare statistics reported that states adopting the expansion of Medicaid, a central provision of the ACA, saw increases in insured adults, those with usual care providers, and those receiving necessary care among non-Hispanic Whites, Latinos, and African Americans. Additionally, there was a decrease in the gap between these metrics among the aforementioned ethnic populations (Hayes et al., 2017).

However, there still exists difficulties in access to healthcare for the aforementioned minority populations. A report published by the U.S. Department of Health and Human Services Agency for Healthcare Research and Quality (AHRQ) found that, of the standardized quality metrics used, African Americans and Latinos receive worse care than non-Latino Whites according to 40% and ~33% of the metrics, respectively. For African Americans, 52% of quality measures had either remained stagnant or worsened in the last two decades. For Latinos, 42% of quality measures had either remained stagnant or worsened in this same time period. Both groups faced the largest disparities in comparison to non-Hispanic Whites in new HIV cases in persons over 13 and HIV infection deaths per 100,000 population, with African American children aged 2-17 also facing the largest disparity in asthma-related hospitalizations when compared to non-Hispanic Whites (*2019 National Healthcare Quality and Disparities Report*, n.d.).

It is clear that despite government intervention, the disparities prevalent across the U.S. healthcare system have continued to persist. As discussed below, the COVID-19 pandemic has highlighted these disparities and led to even further healthcare inequity for minority populations. The analysis presented here aims to uncover any potential correlations between the implementation of artificial intelligence in the COVID-19 response and this outcome.

AI and the COVID-19 Response

The onset of the coronavirus pandemic early last year brought many challenges and required swift action by federal, state, and local governments to respond to the rising case numbers. One of the largest questions faced by decision makers was how to allocate resources, as there were nationwide shortages in personal protective equipment (PPE) for healthcare workers, coronavirus testing supplies, and even common household products such as toilet paper. Oftentimes, officials turned to advisors, consultants, and technology to better inform their decisions regarding how coronavirus response funding should be allocated.

Resource Allocation During COVID-19 Depends on Historical Data

A highly cited article published by a group in the *New England Journal of Medicine* presented recommendations on medical resource allocation during the pandemic based on prevailing ethical principles converged upon by other research in the space. The group found that following six recommendations would allow for resource allocation in an ethically responsible manner; the recommendations were as follows. First, benefits of actions must be maximized in the form of saving the most lives and preserving the most quality of life post-infection. Next, healthcare workers must be prioritized, as they are the ones caring for the infirm. Also, an allocation by first-come-first-serve should be avoided to prevent the influence of systemic inequities and advantages in resource distribution. Additionally, officials should remain responsive to evidence in a rapidly changing landscape surrounding the coronavirus. Officials are also recommended to recognize the risks assumed by research participants and take these into account. Finally, the same standards should be applied to COVID-19 and non-COVID-19 patients (Emanuel et al., 2020).

The above recommendations encompass the anatomy of an ethical response to absolute scarcity such as that faced in the early days of the coronavirus pandemic, but they all largely rely on access to information. The first point about maximizing benefits is particularly reliant on historical data, as actions taken in that vein are largely based on predictions and hypotheses. Thus, technology to sift through this data and information was utilized to assist in making accurate predictions and allocating COVID-19 resources.

The AI Solution to COVID-19 Resource Allocation was Biased

Artificial intelligence (AI) is part of the aforementioned technology that is apt at sifting through large amounts of data and was utilized during the pandemic for this very reason.

Specifically, AI was used to determine proper allocation of ventilators and ICU beds in various areas of the country by implementing an algorithm that based predicted need on historic healthcare spending. This seemingly objective measure, however, actually resulted in bias against African American populations (Röösli et al., 2021).

The commercialized AI program introduced above resulting in the aforementioned bias utilized healthcare spending as a metric for future need. A recent study, however, found that this metric led to African Americans being attributed the same need index as their less infirm non-Hispanic White counterparts. This result is due to the reduced access to quality healthcare available to minority populations, as described in an earlier section of this paper, which results in less clinic and hospital visits, fewer and less frequent procedures, and subsequently lower overall expenditure. The study also found that addressing this bias would result in an increase from 17.7% to 46.5% of African Americans receiving additional resources (Obermeyer et al., 2019). It should be noted that this study was conducted prior to the coronavirus pandemic but analyzed a widely used algorithm in the healthcare industry that was later used in COVID-19 resource allocation.

Unfortunately, the biased resource allocation was followed by disparities in COVID-19 infection and death rates. The U.S. CDC found that COVID-19 hospitalizations and deaths are ~3 and ~2 times more common in minority populations when compared to non-Latino Whites, respectively (CDC, 2020). Additionally, separate work found in an analysis of COVID-19 data from counties across the U.S. that those in the most vulnerable counties had a 1.73 times higher mortality rate than those in the least vulnerable counties, and that when only minority vulnerable communities were compared the factor rose to 4.74 (Khazanchi et al., 2020). However, it should be noted that these are correlative rather than causal findings in relation to the aforementioned AI technology.

COVID-19 Resource Disparities in Virginia

The resource allocation disparity faced nationwide is not lost at the state level. In the state of Virginia in particular, 14.2% and 8.0% of administered vaccine doses have gone to African Americans and Latinos, respectively, as reported by the Virginia Department of Health (*COVID-19 Vaccine Demographics – Coronavirus*, n.d.). These percentages are in context of the 2019 U.S. Census findings of African Americans and Latinos being 19.9% and 9.8% of the state population, respectively (*U.S. Census Bureau QuickFacts*, n.d.). To further explore the resource allocation and

COVID-19 outcome disparity across the state of Virginia, the following analysis was conducted on the counties with the largest and smallest minority demographics.

Examining Disparities Quantitatively

To conduct the aforementioned analysis, population and demographic data from the ten counties with the highest African American population, the ten counties with the lowest African American population, and the ten counties with the highest average income was gathered through the U.S. Census, the Virginia Department of Health, and the UVA Weldon Cooper Center for Public Service (*Population Estimates for Age & Sex, Race & Hispanic, and Towns | Weldon Cooper Center for Public Service*, n.d.). Then, data regarding vaccination distribution was aggregated for these counties through the public data made available by the Virginia Department of Health. Next, vaccines per capita and vaccination representation (measured as percentage of the total population who are African American minus percentage of the vaccinated population who are African American) were calculated for each of the 30 counties. Finally, these values were compared across groups using Unpaired Student's T-tests.

After conducting the analysis detailed above, it was revealed that the 10 counties with the highest African American populations had an average value for vaccines per capita of 0.768 doses per person, compared to the 0.375 and 0.271 doses per person for the 10 counties with the lowest African American populations and the 10 highest-income counties, respectively (though only the latter difference was statistically significant with $p = 0.035$). When comparing vaccine representation (as defined above), the 10 counties with the largest African American populations had vaccine underrepresentation for African Americans of 10.04%, compared to -0.21% and 3.25% for the 10 counties with the lowest African American populations and the 10 highest-income counties, respectively (statistically significant difference with $p < 0.0001$ and $p = 0.0043$, respectively). The p-value for statistical significance was considered to 0.05 for these analyses.

Discussion of Results

The above results deviated from the expectation that more vulnerable counties with larger minority populations would receive fewer vaccine doses per capita than both the counties with smaller minority populations and wealthier counties (classified as less vulnerable), as the opposite was found through data analysis. The methods above refuted the original hypothesis that vaccine distribution would follow the same trends as the resource allocation earlier in the pandemic detailed in prior sections.

Further exploration of Virginia's vaccine distribution plans and policies was conducted, revealing that Virginia officials had pivoted the vaccine distribution strategy in late January in an effort to distribute doses in a more equitable manner. Originally, the distribution was through both health districts and individual health systems/hospitals and the allocation was dependent on demand from districts. It was reported, however, that the VDH turned to allocating based on share of state population and not sending doses through both the aforementioned pathways, leading to a wave of appointment cancellations in Fairfax County, among others (Masters et al., 2021).

Though the above results refuted the original hypothesis, review of Virginia's allocation strategy shows that the initial allocation faced a similar bias as the AI algorithm used for ventilator distribution – utilization of historic healthcare usage and demand to predict future needs. However, despite the manual adjustment for this bias, there still appeared to be a disparity in distribution at the locality-level, as African Americans were underrepresented in the vaccinated population to a greater extent in the counties with the largest minority populations. This case study of vaccine distribution in Virginia emphasizes the difficulty of untangling bias stemming from systemic issues. This is the major challenge healthcare AI algorithms face.

Viewing these results through Hughes' lens of technological momentum, the historic and systemic underserved status of minority communities in the healthcare system constitutes society's impact on AI technology. Subsequently, the results of AI systems trained on biased and unrepresentative data sets impacts societal welfare through inequitable distribution of resources and outcomes as shown above. Following this line of inquiry, one can consider the profound impact these inequitable outcomes would have on future machine learning algorithms trained on pandemic healthcare data. This analysis reveals a vicious cycle of bias that is often very difficult to break and should be watched for in other applications of AI – specifically those in healthcare as discussed below.

Medical Bias is Prevalent in the U.S. Healthcare System

The historical case study presented above highlights how artificial intelligence (AI) incorporation in healthcare resource allocation can lead to a propagation of the existing disparities faced by minority populations. However, the inequities in the healthcare system extend beyond access to care and can also have significant impacts in patient outcomes. Due to systemic disadvantages such as low socioeconomic status, higher uninsured rates, and lower levels of quality education, minority populations often suffer multiple comorbid conditions. These may contribute to statistics such as the one of Black women having a 2-3 times higher rate of death due to pregnancy-related complications presented in an earlier section (*Racial and Ethnic Disparities Continue in Pregnancy-Related Deaths | CDC Online Newsroom | CDC*, 2019).

Minorities Face Health Disparities in Cardiovascular Care

Cardiovascular conditions such as cardiovascular disease (CVD) coincide with the largest racial disparities in access to care and patient outcomes, as these conditions have comorbidities that minorities are more vulnerable to and also require extensive longitudinal care. Studies have found that African Americans are more vulnerable to these comorbidities than their non-Hispanic White counterparts; specifically, African Americans have twice the likelihood of having Diabetes, a 40% higher chance of having hypertension (high blood pressure), and face a higher rate of obesity (Graham, 2015). These conditions increase their likelihood of developing cardiovascular conditions that can seriously and sometimes fatally compromise their health.

The disparities continue beyond these comorbid conditions to the family of diseases and conditions that fall under the CVD umbrella. In particular, African Americans are 30% more likely to die from CVD, have higher rates of myocardial infarctions (MIs), and have higher rates of heart failure than their non-Hispanic White counterparts (Graham, 2015). Unfortunately, providers are often unaware of the disparities in access and care faced by these minority populations. An analysis of cardiologists' awareness of these disparities found startling results showing that their awareness was overall low and decreased in reference to their own clinical practice (Lurie et al., 2005). An economic analysis revealed that over $1 trillion is the combined cost of health inequalities and premature death, with over 30% of excess medical care expenditures resulting from health inequalities (LaVeist et al., 2009).

<u>Early AI Implementation in Healthcare</u>

Artificial intelligence (AI) applications in healthcare and clinical settings have been explored and studied extensively, with multiple potential applications in dermatology and radiology, to name a few, presented in literature throughout the last decade (Kulkarni et al., 2020). Though the incorporation of AI into clinical practice faces many challenges and is unlikely to come to fruition within the near future, most researchers believe that its role in healthcare will, more realistically, slowly increase to assist physicians in clinical decision-making.

AI Has Already Been Implemented in Clinical Settings

A recent review on AI applications in healthcare outlined specific breakthroughs in the fields of radiology, pathology, dermatology, and ophthalmology. More specifically, radiology has seen advances in X-ray interpretation, MRI study reads with and without gadolinium, and administrative task workflow improvement. AI applications in pathology have shown improvement in identifying various types of cancers through pathohistology in multiple studies. Dermatology has seen preliminary results indicating a similar improvement in the identification of a multitude of skin conditions. Finally, AI implementation in ophthalmological settings is being explored to differentiate between mild and severe diabetic retinopathy (Kulkarni et al., 2020). Though these applications are in their infancy and still far from mainstream clinical use, they represent the large body of work committed to eventually implementing AI into clinical practice.

In cardiology, AI has been implemented in the practices of echocardiography and nuclear cardiology. In echocardiology, the technology has been used to better analyze and quantify ventricular activity, predict mortality rates, and differentiate between nuanced diagnoses. In nuclear medicine, AI technology has already been implemented in single-photon emission computed tomography (SPECT) myocardial perfusion imaging (MPI) to enhance images. In fact, there is currently an FDA-approved program that collects these MPI results and compares them to larger datasets in order to assist the physicians interpreting them (Lopez-Jimenez et al., 2020).

Preliminary Findings Show Bias in Clinical Implementations of AI

These early implementations of AI were fraught with bias and led to inequitable outcomes across gender and race. In the dermatologic applications mentioned above, the algorithms did a worse job identifying cancerous lesions on dark skin than on light skin due to the training data containing mainly images from light-skinned individuals (Adamson & Smith, 2018). In the aforementioned radiologic applications, the technology underperformed in reading chest X-rays from the gender underrepresented in the training data set (Larrazabal et al., 2020). These results are in line with the previous examination of the algorithm used in the healthcare industry to allocate resources, with both incorporating the historic bias ingrained in the datasets used for learning.

Future AI Implementations in Clinical Medicine Will Likely Be Biased

In comparison with the historical case study presented earlier in this paper regarding AI bias in COVID-19 resource allocation, the above analysis of healthcare disparities and AI's propagation of those disparities is eerily similar. Thus, the outcome of widespread clinical implementation of AI can be reasonably predicted to increase societal bias and inequities around healthcare. Specifically, the mechanism of this continued disparity in care would be through the incorporation of current healthcare inequities, such as a lack of access to primary care, affordable medication, or nutritional food leading to comorbidities, into training data sets (Hague, 2019). The field of cardiology is especially vulnerable to biased data sets, as many conditions comorbid with cardiovascular disease (CVD) are heavily influenced by lifestyle and require longitudinal care that disadvantaged groups often do not have access to (Muncan, 2018). Moreover, the metrics used by these algorithms would project objectivity to policy decision makers, healthcare providers, and the public, solidifying the aforementioned disparities by discouraging meaningful change in the future.

Discussion

This analysis found disparities in healthcare access that influenced the AI algorithms tasked with the allocation of ventilators and ICU beds early in the pandemic. Extending this analysis to the allocation of vaccine doses in the state of Virginia, it was found that despite conscious efforts by officials to counteract resource allocation disparities, minority populations were still underrepresented in vaccinated populations on a county basis. This analysis revealed that removing systemic bias from systems that require historical data is a complex challenge that is not easily overcome.

Drawing parallels between resource allocation during COVID-19 and the current state of the U.S. healthcare system, it is evident that a similar environment is being created in the latter. Specifically, bias leading to disparities in care and patient outcomes was found to be prevalent in clinical settings, and preliminary analyses of AI systems being researched for potential clinical applications revealed that this bias was permeating into the technology's performance. Thus, this analysis made the reasonable prediction that the implementation of AI in routine clinical practice would lead to greater disparities of outcome for minority populations.

The common theme of AI systems simply accepting the bias from historical data and the apparent inability of basic human intervention to rectify this is a product of algorithms following "black box" models. The concept of a "black box" refers to a system that takes inputs and transforms them into outputs without offering clarity on the steps taken in between. Thus, developers and users cannot easily trace (nor validate) the decision matrix used by the system, leading to the results described throughout this analysis (Kulkarni et al., 2020). However, this black box is what gives AI systems an advantage over humans - it promotes objective analysis of connections based on the data it is given. Despite these advantages, though, the issues outlined in prior sections must still be addressed in ways detailed below.

Conclusion

Historical case studies were used to view artificial intelligence (AI) implementation in healthcare settings through the lens of technological momentum. By analyzing resource allocation during COVID-19, it was found that bias from society influenced the AI technology used, which in turn promoted continued disparities in that same society. Extending this analysis to clinical settings, a similar bias influence from society was found to affect preliminary AI systems being researched for clinical implementation. Thus, it is predicted that routine clinical use of these systems will lead to continued disparities in clinical outcomes for minority populations unless proactive steps are taken to mitigate the influence of biased historical data and training sets. Accepted ethical systems do not simply suggest mitigative measures be taken, but rather demand their implementation.

The goal of this paper, however, is not to dissuade the scientific community from pursuing AI applications in healthcare. In fact, the “black box” model employed by these algorithms can be advantageous in clinical settings by making connections between various symptoms and outcomes that physicians could not (Miotto et al., 2016). Instead, this paper vouches for responsible AI development and implementation through increased transparency in the inner workings of the algorithms used. Examples of methods that promote transparency are model induction, the slight alteration of inputs and observation of changes in outputs, or development of a second neural network to decode the decision-making process of the first (Bleicher, n.d.). If complete transparency can be achieved in AI algorithms and data used to train them, concerted and focused efforts can be made to meaningfully combat the bias attributed to historical data and reduce the likelihood of AI propagating disparities of the past. Specifically, various players involved in healthcare can all make an impact; engineers can develop more transparent AI systems, physicians can work to better acknowledge social determinants of health, and insurers and legislators can thoroughly dissect the complex data they are using for policy decisions. By taking these necessary steps in a concerted effort, society can break the vicious cycle of technological momentum and work towards a brighter and more innovative future.

References

2019 National Healthcare Quality and Disparities Report. (n.d.). Retrieved April 6, 2021, from http://www.ahrq.gov/research/findings/nhqrdr/nhqdr19/index.html

Adamson, A. S., & Smith, A. (2018). Machine Learning and Health Care Disparities in Dermatology. *JAMA Dermatology*, *154*(11), 1247–1248. https://doi.org/10.1001/jamadermatol.2018.2348

Biased bots: Artificial-intelligence systems echo human prejudices. (n.d.). Princeton University. Retrieved October 16, 2020, from https://www.princeton.edu/news/2017/04/18/biased-bots-artificial-intelligence-systems-echo-human-prejudices

Bleicher, A. (n.d.). *Demystifying the Black Box That Is AI*. Scientific American. Retrieved February 26, 2021, from https://www.scientificamerican.com/article/demystifying-the-black-box-that-is-ai/

Brown, E. R., Ojeda, V. D., Wyn, R., & Levan, R. (2000). *Racial and Ethnic Disparities in Access to Health Insurance and Health Care*. https://escholarship.org/uc/item/4sf0p1st

CDC. (2020, February 11). *Cases, Data, and Surveillance*. Centers for Disease Control and Prevention. https://www.cdc.gov/coronavirus/2019-ncov/covid-data/investigations-discovery/hospitalization-death-by-race-ethnicity.html

Colarossi, A. (n.d.). *Summary and critique of Technological Momentum by Thomas P. Hughes*. Retrieved November 2, 2020, from https://www.academia.edu/1881558/Summary_and_critique_of_Technological_Momentum_by_Thomas_P_Hughes

COVID-19 Vaccine Demographics – Coronavirus. (n.d.). Retrieved April 6, 2021, from https://www.vdh.virginia.gov/coronavirus/covid-19-vaccine-demographics/, https://www.vdh.virginia.gov/coronavirus/covid-19-vaccine-demographics/

De Lew, N., Greenberg, G., & Kinchen, K. (1992). A layman's guide to the U.S. health care system. *Health Care Financing Review*, *14*(1), 151–169.

Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., & Phillips, J. P. (2020). Fair Allocation of Scarce Medical Resources in the Time of Covid-19. *New England Journal of Medicine*, *382*(21), 2049–2055. https://doi.org/10.1056/NEJMsb2005114

Graham, G. (2015). Disparities in Cardiovascular Disease Risk in the United States. *Current Cardiology Reviews*, *11*(3), 238–245. https://doi.org/10.2174/1573403X11666141122220003

Hague, D. C. (2019). Benefits, Pitfalls, and Potential Bias in Health Care AI. *North Carolina Medical Journal*, *80*(4), 219–223. https://doi.org/10.18043/ncm.80.4.219

Hall, W. J., Chapman, M. V., Lee, K. M., Merino, Y. M., Thomas, T. W., Payne, B. K., Eng, E., Day, S. H., & Coyne-Beasley, T. (2015). Implicit Racial/Ethnic Bias Among Health Care Professionals and Its Influence on Health Care Outcomes: A Systematic Review. *American Journal of Public Health*, *105*(12), e60–e76. https://doi.org/10.2105/AJPH.2015.302903

Hayes, S. L., Riley, P., Radley, D. C., & McCarthy, D. (2017). Reducing Racial and Ethnic Disparities in Access to Care: Has the Affordable Care Act Made a Difference? *Issue Brief (Commonwealth Fund)*, *2017*, 1–14.

How does implicit bias by physicians affect patients' health care? (n.d.). Https://Www.Apa.Org. Retrieved October 16, 2020, from https://www.apa.org/monitor/2019/03/ce-corner

Hughes, T. P. *Technological Momentum*. (1994). Cambridge, Massachusetts. London, England. The MIT Press.

Khazanchi, R., Beiter, E. R., Gondi, S., Beckman, A. L., Bilinski, A., & Ganguli, I. (2020). County-Level Association of Social Vulnerability with COVID-19 Cases and Deaths in the USA. *Journal of General Internal Medicine*, *35*(9), 2784–2787. https://doi.org/10.1007/s11606-020-05882-3

Klein, H. K., & Kleinman, D. L. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values,* 27(1), 28-52. https://doi.org/10.1177/016224390202700102

Kulkarni, S., Seneviratne, N., Baig, M. S., & Khan, A. H. A. (2020). Artificial Intelligence in Medicine: Where Are We Now? *Academic Radiology*, *27*(1), 62–70. https://doi.org/10.1016/j.acra.2019.10.001

Larrazabal, A. J., Nieto, N., Peterson, V., Milone, D. H., & Ferrante, E. (2020). Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences*, *117*(23), 12592–12594. https://doi.org/10.1073/pnas.1919012117

LaVeist, T., Gaskin, D., & Richard, P. (2009). The Economic Burden of Health Inequalities in the United States. *Health Policy and Management Faculty Publications*. https://hsrc.himmelfarb.gwu.edu/sphhs_policy_facpubs/225

Lopez-Jimenez, F., Attia, Z., Arruda-Olson, A. M., Carter, R., Chareonthaitawee, P., Jouni, H., Kapa, S., Lerman, A., Luong, C., Medina-Inojosa, J. R., Noseworthy, P. A., Pellikka, P. A., Redfield, M. M., Roger, V. L., Sandhu, G. S., Senecal, C., & Friedman, P. A. (2020). Artificial Intelligence in Cardiology: Present and Future. *Mayo Clinic Proceedings*, *95*(5), 1015–1039. https://doi.org/10.1016/j.mayocp.2020.01.038

Lum, K., & Isaac, W. (2016). To predict and serve? *Significance*, *13*(5), 14–19. https://doi.org/10.1111/j.1740-9713.2016.00960.x

Lurie, N., Fremont, A., Jain, A. K., Taylor, S. L., McLaughlin, R., Peterson, E., Kong, B. W., & Ferguson, T. B. (2005). Racial and ethnic disparities in care: The perspectives of cardiologists. *Circulation*, *111*(10), 1264–1269. https://doi.org/10.1161/01.CIR.0000157738.12783.71

Manyika, J., Silberg, J., & Presten, B. (2019, October 25). What Do We Do About the Biases in AI? *Harvard Business Review*. https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai

Masters, K., January 27, V. M., & 2021. (2021, January 27). Virginia shifts vaccine distribution strategy, triggering thousands of canceled appointments. *Virginia Mercury*. https://www.virginiamercury.com/2021/01/27/virginia-shifts-vaccine-distribution-strategy-triggering-thousands-of-canceled-appointments/

Mattu, J. A., Jeff Larson,Lauren Kirchner,Surya. (n.d.). *Machine Bias*. ProPublica. Retrieved October 16, 2020, from https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing?token=nponEfVIhbu_Yb4_j1EeltF4h5FW_Uc4

Miotto, R., Li, L., Kidd, B. A., & Dudley, J. T. (2016). Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records. *Scientific Reports*, *6*. https://doi.org/10.1038/srep26094

Muncan, B. (2018). Cardiovascular disease in racial/ethnic minority populations: Illness burden and overview of community-based interventions. *Public Health Reviews*, *39*(1), 32. https://doi.org/10.1186/s40985-018-0109-4

Neiger, C. (2018, July 28). *5 Reasons You Should Own at Least 1 Artificial Intelligence Stock*. The Motley Fool. https://www.fool.com/investing/2018/07/28/5-reasons-you-should-own-at-least-1-artificial-int.aspx

No surprise Amazon's AI was biased against women, says Sandra Wachter—Business Insider. (n.d.). Retrieved October 16, 2020, from https://www.businessinsider.com/amazon-ai-biased-against-women-no-surprise-sandra-wachter-2018-10

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

Population Estimates for Age & Sex, Race & Hispanic, and Towns | Weldon Cooper Center for Public Service. (n.d.). Retrieved April 6, 2021, from https://demographics.coopercenter.org/population-estimates-age-sex-race-hispanic-towns

Racial and Ethnic Disparities Continue in Pregnancy-Related Deaths | CDC Online Newsroom | CDC. (2019, September 6). https://www.cdc.gov/media/releases/2019/p0905-racial-ethnic-disparities-pregnancy-deaths.html

Röösli, E., Rice, B., & Hernandez-Boussard, T. (2021). Bias at warp speed: How AI may contribute to the disparities gap in the time of COVID-19. *Journal of the American Medical Informatics Association*, *28*(1), 190–192. https://doi.org/10.1093/jamia/ocaa210

Smith, C. S. (2020, January 2). Dealing With Bias in Artificial Intelligence. *The New York Times*. https://www.nytimes.com/2019/11/19/technology/artificial-intelligence-bias.html

Smith, M.R. (1994). Technological Determinism in American Culture. *Does Technology Drive History?: The Dilemma of Technological Determinism*. (pp. 1-17). Cambridge, Massachusetts. London, England. The MIT Press.

Study finds gender and skin-type bias in commercial artificial-intelligence systems. (n.d.). MIT News | Massachusetts Institute of Technology. Retrieved October 16, 2020, from https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212

Sweeney, L. (2013). *Discrimination in Online Ad Delivery* (SSRN Scholarly Paper ID 2208240). Social Science Research Network. https://doi.org/10.2139/ssrn.2208240

U.S. Census Bureau QuickFacts: Virginia. (n.d.). Retrieved April 6, 2021, from https://www.census.gov/quickfacts/VA