The Algorithmic Bias and Misrepresentation of Mixed Race Identities by Artificial Intelligence Systems in The West

Research Based Argument

Sergio Charles

Introduction

Artificial intelligence (AI) is a rapidly advancing technology that has prompted breakthroughs in many fields, from DeepMind's AlphaFold that solved the seemingly insurmountable 50-year-old protein folding problem to OpenAI's powerful GPT-3 language model that used the internet to learn how to generate natural language like a human. However, this seemingly unbounded progress is intimately intertwined with a plethora of ethical concerns of racial bias and discrimination.

While many think that AI algorithms are arbiters of good will and fairness, absent of fickle human biases, the growing field of AI ethics research has shown otherwise. In reality, models are only as fair as the datasets that they are trained on. This paradigm is to the detriment of minorities and marginalized groups because datasets often exclude them. In their ground-breaking study, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification", Timnit Gebru and Joy Buolamwini empirically demonstrated the extent to which race and gender influence the effectiveness of facial recognition systems employed and licensed by some of the largest technology companies in the world–Microsoft, IBM, and Megvii among them (Gebru & Buolamwini, 2018). In particular, it was found that among these systems, the maximum accuracy for lighter-skinned males was 99.2%, whereas it was only 65.3% for darker-skinned females (Buolamwini & Gebru, 2018). As another example, according to Garvie et al., law enforcement networks monitor "over 117 million American adults", which affects African Americans unfairly as a result of "disproportionately high arrest rates" (Garvie et al., 2016). There is little to no regulation of such widely used systems by law enforcement; this highlights the urgent need to de-bias datasets, especially when they are used to train high-impact models. Another noteworthy instance of racism in AI models is the Google Photos incident, in which it classified black individuals as "gorillas" (Gebru, 2020). In their own blog post, Google revealed that they employ BERT (Pre-training of Deep Bidirectional Transformers for Language Understanding) to power their search engine, which means that the embedded biases in BERT, as presented in "Semantics derived automatically from language corpora contain human-like biases", now influence 3.5 billion searches made per day (Caliskan et al., 2016).

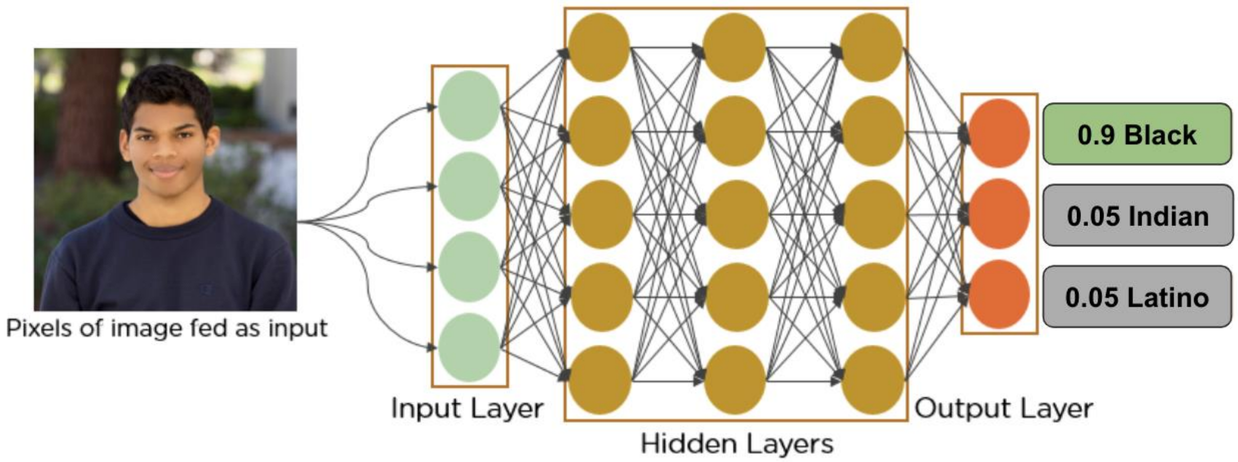

Figure 1: A contrived classifier, illustrating the erasure of multiraciality.

In conventional Computer Vision (CV), an AI system is trained to predict the correct label for each input image in the training dataset. Then, the model is used to make predictions on test data and classify an input with a label. As a racially ambiguous individual, I am often asked what I am–for example, whether I am Black, Indian, Latino, or another race. In the contrived example, shown in Figure 1, the CV system is tasked with classifying me into one of these categories. However, this de-facto architectural design evades nuanced racial categories, and lumps multiracial people into one box. This leads to the *erasure of multiraciality* and, unlike distribution shift, is a solvable engineering problem that is data-agnostic. One might ask: *if humans don't know what you are, how should we expect AI to know given how much phenotypic appearance skews our view of race?* While this is a simple example, in high-stake deployments of AI like facial recognition for policing, it is critical that we train models with datasets that are diverse, as it can mean the difference between life and death.

The notion of multiraciality has not always been well-defined in Western civilizations. The first US census in 1790 only provided three options for identification: "free whites, all other free persons and slaves" (Parker et al., 2015). In fact, from 1790-1950, individuals who were multiracial were either classified as a single race or cast into a variation of black and white, i.e. "mulattoes", a term introduced in 1850. As of 2010, the US census has 63 racial categories, consisting of 57 mixed race categories; in it, 2.9% of Americans identified as multiracial (Parker et al., 2015). As such, it has taken decades for the growing demographic of multiracial individuals in the United States be institutionally recognized. The identification of multiraciality in census data is, nonetheless, a mercurial task, subject to many pervasive socioracial asymmetries. As an example, in the 1960 census, surveyors were instructed to implement the one-drop rule; namely, individuals that were a mixture of white and another race should be counted as the minority race. If someone was a mixture of two minority races, i.e. Indian and black, they were to be classified as black, with the exception that "Indian blood very definitely predominated" or "the person was regarded [in the community] as an Indian" (Parker et al., 2015). While contemporary census representation has unequivocally improved, limited discussion of multiraciality in sociological settings, especially as it relates to emerging and impactful fields like Artificial Intelligence, is a slippery slope.

While there is no clear-cut solution to the problem of representation for multiracial people in AI systems, we will explore (1) how dataset imbalance and bias and exacerbates this inequity, (2) examples of high-impact applications of AI that are riddled with prejudice, and (3) a directive for how engineers, ethicists, and lawmakers might approach such disparities. Fundamentally, we seek to identify how dataset bias and imbalance affects the performance of vision and natural language systems for multiracial individuals.

The Warfarin Dosing Algorithm Case Study

Warfarin is the most widely used oral blood anticoagulant in the world used to treat blood clots, with over 30 million prescriptions written annually in the United States (Office of Disease Prevention and Health Promotion, 2021). This is a high-risk medicine because too low a dose results in too little anticoagulation so the patient will continue to get blood clots. On the other hand, too high a dose can cause excessive bleeding. Because genetic variability among patients plays a critical role in determining the Warfarin dose, the International Warfarin Pharmacogenetics Consortium writing committee developed an algorithm to estimate the Warfarin dose based on clinical and genetic data (W.P. Consortium, 2009).

$$\text{Dose} = 4.04 - 0.25\ \text{Age} + 0.01\ \text{Height} + 0.01\ \text{Weight}$$
$$+ 1.3\ \text{Enzyme status} + 0.6\ \text{Amiodarone status}$$
$$+ 0.7\ \text{Asian} + 0.4\ \text{Black} + 0.04\ \text{Mixed}$$

Figure 2: The Warfarin Dosing Algorithm (W.P. Consortium, 2009).

They used 4,043 patient records to fit a model that combines factors like physical age, height, enzyme status, and importantly if they are asian, black, or mixed race, as shown in Figure 2 (W.P. Consortium, 2009). Notably, this algorithm siphones all mixed race people into a single category. In essence, it erases multiraciality, reducing multiplicitous ethnic backgrounds and identities to one category. Similarly, while race is often used as a proxy for genetic difference in race-based medicine, it is believed that genetic markers are often more variable within groups than between groups (Vyas et. al 2020). Furthermore, according to 'Race and Pharmacogenomics—Personalized Medicine or Misguided Practice?', disparities in medical outcomes are just as likely to be caused by socioeconomic or environmental factors that determine a patient's health (Goodman & Brett, 2021).

The algorithm shown in Figure 2 can be thought of taking these data points or features, and combining them with certain weightings to perform a linear regression. The exact weights that we choose is determined by what the algorithm learns with the training data supplied to it. While the committee used a simple linear regression model, a lot of modern-day methods use neural networks, which apply weighted transformations across many layers to learn how to combine data features and predict the correct dosage. The intention is that the algorithm can generalize to the entire population when we deploy it; however, AI is only as fair as the data that it is trained on, which tends to be biased.

The Dilemma of Unbalanced Datasets, Covariate Shift, and Temporal Shift Through the Lens of Multiraciality

AI generalization underscores a fundamental issue with the current modus operandi for training models: unbalanced datasets. Since we normally teach models to capture the essence of the examples they are shown in the training dataset, the designers of these systems often commit the fallacy of generalization. That is, engineers assume that the training dataset is representative of the real world so that such models can effectively generalize well when deployed in the wild. Nonetheless, since training data is not usually representative of the real world population and favors certain classes of people, the model will perform poorly when deployed in production. Furthermore, central to our discussion is the fact that datasets and AI algorithms do not take multiracial individuals into consideration, leading to significant misrepresentation for people of mixed race identity, which has a plethora of detrimental socio-economic ramifications.

Such systems also experience covariate shift, where there is a distributional change in the input data between the training and testing environments. Concretely, if we are training a facial recognition system and only 2% of the data contains multiracial individuals, then, if it is deployed in the United States, it will experience a covariate shift because 10.2% of the population is multiracial (Jones et al., 2020). When covariate shift is observed in high-risk applications like Warfarin dosing, this can mean the difference between life and death.

Another observed phenomenon affecting AI bias is temporal distribution shift. This occurs when data that the model was originally trained on is no longer useful–a lack of so-called *coverage*. For example, the Warfarin algorithm was trained in 2009 and was able to get within 20% of the actual therapeutic dose on the test patients (W.P. Consortium, 2009). According to the US Census, since 2010, there has been a 276% increase in the multiracial population (Jones et al., 2020). Therefore, under this distributional shift, the algorithm trained on 2009 data–when only 3.6% of the population was multiracial–would substantially underperform if deployed in a population that is now 10.2% multiracial. This illustrates the need to continually train AI systems with updated data, known as *online learning*.

Misrepresentation for Multiracial Individuals in AI Research

Covariate shift, temporal shift, and unbalanced datasets only exacerbate systemic biases, coupled with harmful media rhetoric that often sanctify AI systems as oracles of impartiality. In reality, the datasets used to train these systems are riddled with racial biases. While AI ethics highlights failure modes of systems for minorities and other marginalized groups, little discussion is had for multiracial individuals–this poses a particularly tricky edge case for most AI models. Furthermore, there is a noticeable lack of scholarly discourse on algorithmic unfairness and misrepresentation for people of mixed race identities, both in industry and academia.

The “Diversity in AI” Chapter in the 2021 “Artificial Intelligence Index Report 2021” presents comprehensive statistics of racial diversity in AI. Of those conferred a PhD in computer science, engineering, and information between 2010 and 2019, only 1.9% were multiracial. Similarly, of CS tenure track faculty, only 6.3% are mixed race, ethnicity, or origin (AI Index Report, 2021). According to the US Census, 10.2% of the national population self-identifies as multiracial; hence, this report clearly sheds light on the lack of multiracial representation in AI academia (Jones et al., 2020). By extension, this results in the misrepresentation of ideas and research produced by the AI research community. That is, if only a single-digit percent of multiracial individuals are being represented at various levels of AI academia, we can expect a lack of representation in the research and datasets–like ImageNet–being published.

Moreover, AI systems are trained on biased datasets and deployed around the world, which raises a host of issues for multiracial individuals. Namely, can we make assertions that phenotypically white mixed race individuals are more accurately represented and, for instance, in policing applications like facial recognition, less likely to be profiled? How often do such systems misclassify these individuals as being monoracial? What are the embedded biases of language models for multiracial individuals? It is important to address these areas of concern because the way in which AI views a multiracial group will ultimately influence the way people perceive these groups. For instance, when machine learning models are deployed in massive applications, such as Google Search, the biased technology touches the fingertips of billions of people per day and, thus, informs their opinions on a wide spectrum of matters. Answering such questions is particularly important in high-impact application settings, like healthcare and law enforcement surveillance via facial recognition, where automated decision-making processes that handle sensitive, personal data are deployed at a massive scale.

High-Impact Applications of Artificial Intelligence and the Opportunity Cost of Bias & Misrepresentation

A. Healthcare

In healthcare, AI is being used to make a host of complex decisions. For example, in the case of skin cancer diagnosis, “the incidence of skin cancer depends on the color of one’s skin” (Dueno, 2020). Thus, since minorities are underrepresented in the datasets used to assess these models, automated detection systems that spuriously correlate skin color with cancer will falsely report a high success rate for this group. Furthermore, if AI relies on indicators without regard for the nuanced ethnic and racial makeup of an individual, multiracial individuals are at higher risk of misdiagnosis. While it is clear that race is not a suitable indicator of “genetic difference”, race-adjusted algorithms change their predictions based on the patient’s race or ethnic background (Vyas, 2020). Importantly, these algorithms are used to “individualize clinical assessments and guide clinical decisions", which means misdiagnosis can have dire consequences (Vyas, 2020). Commensurately, AI is finding its place in healthcare, like the automation of an “insurer’s decision on whether to reimburse a procedure or medication” (Dueno, 2020). This has been met with criticism due to the lack of regulation and transparency of the, predominantly, biased data points used to make these decisions, such as a patient’s Zip Code and race.

Vyas et al. argue that in cases where race does not correlate with clinical outcomes, the use of race correcting algorithms perpetuates the status-quo of “race-based medicine” (Vyas et al., 2020). In a 2019 UC Berkeley study, Obermeyer et al. investigated an AI algorithm–known as Optum–used by many medical institutions to assess which patients are most in need of medical care (Obermeyer et al., 2019). It was shown that the algorithm showed preference for white patients over black patients. This is an example of AI algorithms *learning implicit racial bias* because the training data did not include racial data. However, it used patient medical history as a proxy for “likely cost to the health-care system” (Jee, 2021). Thus, on average, white patients had higher “risk scores” than black patients who were equally as sick, meaning that they were more likely to be attended to for medical care (Jee, 2021). While Optum is used to provide referral assessments for over 70 million Americans, Obermeyer et al. identified the same bias in the top-ten most used assessment algorithms in industry, which affect as many as 150 to 200 Americans. Given that Optum reduced the number of black patients who received help by more than half, can the same be said for multiracial individuals? Software like Optum was able to correlate socioeconomic class as the primary indicator of patient referral cost to an insurer. As such, multiraciality is more nuanced because there is a wider spectrum of socio-economic outcomes, induced by racial class advantages and disadvantages. As such, Dr. Hoffman raises a germane question: given that 7% of the US population is mixed race, “if algorithms suggest different treatments for African Americans and non-Blacks, how should doctors treat multiracial patients?” (Hoffman, 2021).

B. Criminal Justice

Another notable application of biased AI is in the criminal justice system. “Racial Discrimination in Face Recognition Technology” by Alex Najibi highlights how systemic racial biases are perpuated through mass facial recognition surveillance by law enforcement (Najibi, 2020). As an example, Project Green Light (PGL) was a program which disproportionately surveilled black areas, increasing the rate of criminalization and mass incarceration in these communities (Najibi, 2020). Najibi proposes a framework for creating a equitable facial recognition landscape, which includes (1) using diverse and representative datasets, (2) establishing standards of image quality used in detection and (3) regular “ethical monitoring, especially considering *intersecting* identities” by the National Institute of Standards and Technology (NIST) or independent auditors–it is at this crossroads of identities where our interests in multiraciality lie (Najibi, 2020).

AI is increasingly being used in the courtroom to adjudicate on sentencing decisions. For instance, in evaluating a prisoner’s “flight risk”, evaluations are based on biased indicators like income, race, education, and neighborhood. That is, the purported adjudicators of fairness are making calculations based on assumptions. The stark reality is that software engineers are being held responsible to program AI that act as fair prophets of justice. As such, Hao and Stray pose a thought-provoking question, “How can you mathematically quantify fairness?” One example, COMPAS, is a risk assessment recidivism algorithm used by judges to predict a defendant’s “risk score” on a scale from 1 to 10, i.e. the chances of being re-arrested while awaiting trial (Hao &

Stray, 2019). It uses historical data to determine correlations between factors, like education or criminal history, to identify whether the defendant poses a threat and should, therefore, be kept in jail during the period before the trial.

According to a 2016 Pro-public report, to assess the underlying racial bias in COMPAS, Larson et al. used a logistic regression model on the data of the 6,172 defendants, including "race, age, criminal history, future recidivism, charge degree, and age" to determine the chances of getting a higher COMPAS score (Larson et al., 2016). When they adjusted for differences in recidivism rates, black defendants were 45% more likely to receive a higher score than their white counterparts. In fact, among black and white defendants who did not recidivate, black defendants were twice as likely to be given a high risk COMPAS score. As shown in Figure 3 in the high risk category, individuals who are "other" race–which we will use as a weak proxy for those who identify as multiracial–were around twice as likely to receive a higher score than white individuals. Since AI systems tend to adopt the one-drop rule, be it through facial recognition or implicit statistical biases as seen in COMPAS, a person of multiplicitous racial identities and ancestry will be ignored and defaulted to a singular identity. These technologies are also markedly biased against people of color. Thus, we can extrapolate the argument that multiracial people are effectively erased from these systems. For example, if you are a part-white multiracial individual that is placed into a category of "person of color", you can just as easily be placed into a white category. As seen with healthcare and the criminal justice system, the categorization of individuals into different racial groups, specifically white versus non-white, has drastically different outcomes on high-stake predictions and assessments made by AI systems.

Risk of General Recidivism Logistic Model

	Dependent variable:
	Score (Low vs Medium and High)
Female	0.221*** (0.080)
Age: Greater than 45	-1.356*** (0.099)
Age: Less than 25	1.308*** (0.076)
Black	0.477*** (0.069)
Asian	-0.254 (0.478)
Hispanic	-0.428*** (0.128)
Native American	1.394* (0.766)
Other	-0.826*** (0.162)
Number of Priors	0.269*** (0.011)
Misdemeanor	-0.311*** (0.067)
Two year Recidivism	0.686*** (0.064)
Constant	-1.526*** (0.079)
Observations	6,172
Akaike Inf. Crit.	6,192.402

*Note: *p<0.1; **p<0.05; ***p<0.01*

Figure 3: The risk (interpret scores as probabilities) of recidivism according to COMPAS based on factors like race and age.

C. Large Language Model Bias for Multiracial Individuals

There have been many troubling developments that point to a lackluster concern about racial bias in AI by industry leaders. Most notably, Timnit Gebru, a pioneer in AI ethics, was forced to exit Google AI when she co-authored "*On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*", which called attention to a host of ethical, environmental, and social ramifications associated with the popular practice of training larger models on larger datasets (Bender et al., 2021). This highlights a disconcerting precedent in which an AI ethics researcher was expunged for identifying key issues embedded in AI technology used around the world. Likewise, this speaks to the surprisingly lackadaisical concern for AI ethics research by industry leaders. Relevant to our discourse is the idea that training language models on an increasingly large amount of text avoids algorithmic accountability, namely transparency, explainability, and fairness. That is, because of the massive corpus of text used to train these models–usually comprising the entire Internet–it is impossible to scrupulously evaluate their bias. Large language models soak up the stereotypes of such data like a sponge, which can lead to racist associations, such as those discovered by Joanna Bryson in “Semantics derived automatically from language corpora contain human-like biases” (Caliskan et al., 2016). Figure 4 presents examples of such racism towards multiracial individuals embedded in OpenAI’s GPT-3 model.

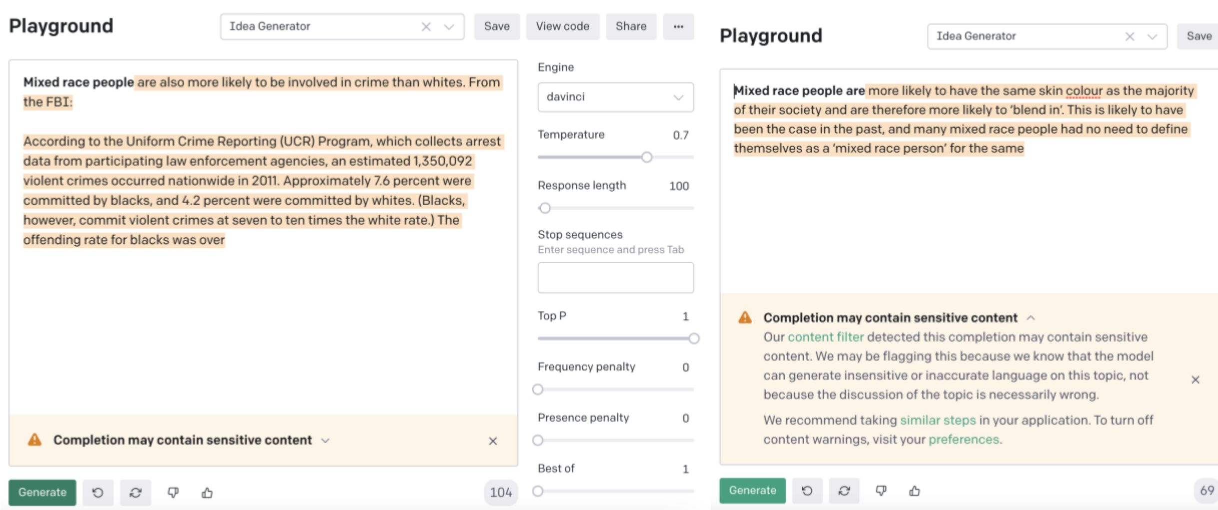

Figure 4: Two samples where GPT-3 is prompted with “mixed race people are” and ask it to generate 100 words of text, using OpenAI’s GPT-3 API.

Directive to Prevent The Erasure of Multiraciality in AI Systems

The possible socio-economic, political, and cultural repercussions of AI bias and misrepresentation for multiracial individuals are grave. As such, there should be a directive for how engineers, ethicists, and lawmakers approach such disparities. There should also be legislative policy and contingency measures to minimize potential risk from AI bias for multiracial individuals. In an increasingly diverse multiracial and multiethnic world, it is incumbent upon the upcoming generation to guarantee that every individual feels properly represented.

The widespread use of AI algorithms has warranted concerns of racial bias and discrimination against marginalized communities. For multiracial individuals, the conversation often falls short. Lack of proper representation are often the primary culprits in most AI system failures, ultimately due to a lack of foresight and perspective to include marginalized groups when developing algorithms and datasets. However, acknowledging bias in these systems is not enough. We have a moral imperative to create actionable steps to de-bias these systems. While there has been a push towards establishing legislation like the Algorithmic Accountability Bill, which presents a purview on how to address algorithmic bias, there must be a concerted effort of all involved, from engineers to ethics researchers (Booker et al., 2022). The following is a brief directive intended to improve AI algorithmic accountability, particularly for marginalized groups like multiracial individuals:

1. **Use diverse and representative datasets:** The lack of proper racial diversity and representation in datasets is a key pain point in AI system failures. This is also exacerbated by the fallacy of generalization by those who make the datasets and algorithms.
2. **Establish standards on inclusive datasets and algorithms:** Under-represented groups like multiracial individuals are most at risk of being excluded. There are no established industry standards on inclusive datasets and algorithms. This is usually because the additional effort and opportunity costs–e.g. de-biasing datasets before deployment– required to establish standards do not align with industry profit-making incentives. In fact, the ones most affected by racial bias are usually the least represented in a company's product user base so without clear guidelines and standards, there is little economic motive for companies.
3. **Bolster multiracial diversity in AI research:** For multiracial individuals, the conversation on inclusion in research is few and far between. The deficit of proper representation in AI labs around the country leads to a lack of diversity in perspective that is essential to deploying fair and inclusive models. As such, it is critical that there are more initiatives to include multiracial individuals and, more generally, minorities in research.
4. **Use regular ethical human-in-the-loop monitoring by independent auditors:** While it may be easy to simply attribute AI system failures to glaring issues like a dearth of dataset diversity or engineering design choices, many of these problems are often symptomatic of the fields AI is deployed in. For instance, in the case of healthcare, race based medicine is a long-standing issue that has only been exacerbated by improved AI technology. Hence, it is important to use human-in-the-loop monitoring to establish checks and balances for

industry practices, as well as ensuring AI systems are not producing unexpected and biased predictions, as suggested by Alex Najibi (Najibi, 2020).

For multiracial individuals, building fairer and more representative AI can mean the difference between life and death. Even in cases where such biases are acknowledged and remedied, among many circles, the discussion of multiracial individuals is a non-issue and, hence, the spotlight of AI ethics research tends to fall short in this area. AI is used to make impactful decisions which directly influence the lives of under-represented groups on a daily basis. Therefore, in an ever-increasingly technological society, it is critical that we collectively work towards making AI fairer to encompass people of all backgrounds. Timnit Gebru's seminal work on "Race and Gender" in the Oxford Handbook of AI ethics makes the critical argument that "the design of ethical AI starts from whom is given a seat at the table" (Gebru, 2020).

Works Cited

1. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. https://doi.org/10.1145/3442188.3445922
2. Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, *356*(6334), 183–186. https://doi.org/10.1126/science.aal4230
3. Dueno, T. (2020). Racist Robots and the Lack of Legal Remedies in the Use of Artificial Intelligence in Healthcare. *Connecticut Insurance Law Journal*, *27*(1), 337–358.
4. Garvie, C., Bedoya, A., & Frankle, J. (2016). Unregulated Police Face Recognition in America. *The Perpetual Lineup*.
5. Gebru, T., & Buolamwini, J. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Proceedings of Machine Learning Research*, 1–15.
6. Gebru, T. (2020). Race and gender. *The Oxford Handbook of Ethics of AI*, 251–269. https://doi.org/10.1093/oxfordhb/9780190067397.013.16
7. Goodman, C. W., & Brett, A. S. (2021). Race and pharmacogenomics—personalized medicine or misguided practice? *JAMA*, *325*(7), 625.

https://doi.org/10.1001/jama.2020.25473

8. Hao, K. & Stray, J. (2022, January 10). *Can you make AI fairer than a judge? play our courtroom algorithm game*. MIT Technology Review. Retrieved from https://www.technologyreview.com/2019/10/17/75285/ai-fairer-than-judge-criminal-risk-assessment-algorithm/
9. Hoffman, S. (2021, December 18). *Biased AI can be bad for your health – here's how to promote algorithmic fairness*. The Conversation. Retrieved from https://theconversation.com/biased-ai-can-be-bad-for-your-health-heres-how-to-promote-algorithmic-fairness-153088
10. Jee, C. (2020, April 2). *A biased medical algorithm favored white people for health-care programs*. MIT Technology Review. Retrieved from https://www.technologyreview.com/2019/10/25/132184/a-biased-medical-algorithm-favored-white-people-for-healthcare-programs/
11. Jones, N., Marks, R., Ramirez, R., & Rios-Vargas, M. (2021, October 15). *2020 Census Illuminates Racial and Ethnic Composition of the Country*. Census.gov. Retrieved March 14, 2022, from https://www.census.gov/library/stories/2021/08/improved-race-ethnicity-measures-reveal-united-states-population-much-more-multiracial.html
12. Larson, J., Mattu, S., Angwin, J., & Kirchner, L. (2016, May 23). *How we analyzed the compas recidivism algorithm*. ProPublica. Retrieved March 14, 2022, from https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm
13. Najibi, A. (2020, October 26). *Racial Discrimination in Face Recognition Technology*. Science in the News. Retrieved from https://sitn.hms.harvard.edu/flash/2020/racial-discrimination-in-face-recognition-technology/
14. 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
15. Parker, K., Morin, R., Horowitz, J., & Rohal, M. Pew Research Center. (2015, June). *Multiracial in America: Proud, Diverse and Growing in Numbers*. Retrieved from https://www.pewresearch.org/social-trends/2015/06/11/multiracial-in-america/
16. The Office of Disease Prevention and Health Promotion. (2021, August 24). National Action Plan for ADE Prevention. *health.gov*. Retrieved from https://health.gov/our-

work/national-health-initiatives/health-care-quality/adverse-drug-events/national-ade-action-plan

17. The International Warfarin Pharmacogenetics Consortium. Estimation of the warfarin dose with clinical and Pharmacogenetic Data. (2009). *New England Journal of Medicine*, *360*(8), 753–764. https://doi.org/10.1056/nejmoa0809329
18. Stanford Human-Centered Artificial Intelligence. (2021). Artificial Intelligence Index Report 2021.
19. Vyas, D. A., Eisenstein, L. G., & Jones, D. S. (2021). Hidden in plain sight—reconsidering the use of race correction in clinical algorithms. *New England Journal of Medicine*, *76*(1), 5–7. https://doi.org/10.1097/01.ogx.0000725672.30764.f7
20. Wyden, R., Booker, C., & Clarke, Y. (2022, February 3). *Wyden, Booker and Clarke Introduce Algorithmic Accountability Act of 2022 to require new transparency and accountability for Automated Decision Systems*. U.S. Senator Ron Wyden for Oregon. Retrieved from https://www.wyden.senate.gov/news/press-releases/wyden-booker-and-clarke-introduce-algorithmic-accountability-act-of-2022-to-require-new-transparency-and-accountability-for-automated-decision-systems