Promoting Health Equity through Improved Regulation of Artificial Intelligence Medical Devices

Kristina Dortche[1,2], Grace McCarthy[3,4], Sara Banbury[1], Isabel Yannatos[5]

[1]Perelman School of Medicine, University of Pennsylvania, Division of Urology, Philadelphia, PA, USA
[2]Mailman School of Public Health, Columbia University, Department of Health Policy and Management, New York, NY, USA
[3]Oregon Health & Science University, Department of Surgery, Portland, OR, USA
[4]Oregon Health & Science University, Brenden-Colson Center for Pancreatic Care, Knight Cancer Institute, Portland, OR, USA
[5]University of Pennsylvania, Department of Neurology, Philadelphia, PA, USA
https://doi.org/10.38126/JSPG210302
Corresponding author: yannatos@pennmedicine.upenn.edu
Keywords: artificial intelligence; healthcare; equity; FDA; SaMD regulation; medical device

Executive Summary: Existing health disparities in the United States are partially driven by the way healthcare is delivered. There is interest in using Artificial Intelligence (AI)-driven software as medical devices (SaMD) to aid in healthcare delivery and reduce health disparities. However, AI-driven tools have the potential to codify bias in healthcare settings. Some AI-driven SaMDs have displayed substandard performance among racial and ethnic minorities. Auditing these tools for biased output can help produce more equitable outcomes across populations. However, there are currently no explicit Food and Drug Administration (FDA) regulations that examine bias in AI software in healthcare. Therefore, we propose the FDA develop a distinct regulatory process for AI-driven SaMDs that includes assessing equitable output across populations and avoiding potential health disparity exacerbation. This change could help prevent AI-driven health disparities nationwide.

I. Health disparity crisis in the United States

Racial and ethnic health disparities persist in the United States despite governmental attempts to improve health outcomes. In the United States, racial and ethnic minorities are 1.5 to 2 times more likely to experience most chronic illnesses (Price et al. 2013). For example, childhood asthma rates are higher among Puerto Rican (13%) and African American (9.5%) children compared to White children (7.2%; Price et al. 2013). Black children experience higher rates of cerebral palsy, HIV/AIDs, and death from Type I Diabetes than White children. Hispanic children have higher rates of spina bifida, depression, and death from acute leukemia than White children (Berry et al. 2010). African Americans are 77%, Hispanics are 66%, and Native Americans are almost 300% more likely to be diagnosed with diabetes than non-Hispanic White Americans. Additionally, African Americans are 30% more likely to die from heart disease and 50% more likely to have a stroke compared to White Americans (Minority 2021). After the onset of the Covid-19 pandemic, age-adjusted death rates climbed by 16.8% from 2019 to 2020 for the total population. This increase in death rate, however, was higher in Black (28.0% for males, 24.9% for females) and Hispanic (42.7% for males, 32.4% for females) Americans in comparison to White (13.4% for males, and 12.1% for females) Americans, highlighting underlying health disparities that are exacerbated during health crises (Murphy et al. 2020). Health disparities are also associated with profound

national economic losses, with an estimated annual cost of $35 billion in excess healthcare expenditures, $10 billion in illness-related productivity loss, and nearly $200 billion in premature deaths (LaVeist et al. 2011).

II. Artificial intelligence in healthcare

i. Benefits of AI technology in healthcare

Healthcare providers are exploring how to use Artificial Intelligence (AI)-driven software as a medical device (SaMD) (IMDRF 2013; IMDRF 2013) to strengthen and revolutionize healthcare. AI technology, sometimes referred to as machine learning, is distinct from traditional software because it typically "learns" from a large amount of data to predict a specific outcome (Chen & Decary 2019). This technology can assist with medical diagnosis and clinical decision-making among other benefits (Amisha et al. 2019; Briganti & LeMoine 2020; Chen & Decary 2020). For example in the pediatric field, AI-based algorithms streamlined test ordering and got faster results for 22.3% of pediatric patients at a children's emergency department (Singh et al. 2022), identified rare subtypes of pediatric pulmonary hypertension difficult to identify from clinical observations (Ong et al. 2017), and created a non-invasive neonatal jaundice detection system (Aydin et al. 2016). Furthermore, the use of AI-based algorithms is currently being explored across numerous pediatric conditions such as prematurity, neonatal seizures and mortality, autism, ADHD, and epilepsy, among others (Hoodbhoy et al. 2021).

If designed properly, AI even has the potential to reduce disparities in healthcare. For example, while a traditional model for breast cancer risk was more accurate for White than African American women, an AI model showed more accurate overall results and equal accuracy across racial groups (Yala et al. 2019).

ii. Bias in AI

However, AI-driven software also has the potential to amplify healthcare bias and further exacerbate racial and ethnic health disparities (Gianfrancesco et al. 2018). Because AI technology learns from data, biases in the dataset—such as underrepresentation of minority groups, unequal amounts of missing or inaccurate data, or existing disparities that are reflected in data—can be encoded into the technology's predictions. These biases can lead to less accurate or less beneficial predictions for medically-underserved patients. For example, an AI-driven diagnostic tool using chest X-rays results in under-diagnosis of lung disease in minority populations (Seyyed-Kalantari et al. 2021). This highlights the tool's biased inputs that were generated by human-based under-diagnosis among marginalized populations. There are also concerns about AI in fields such as dermatology and cardiology, where minority patients are under-represented in the databases used to train healthcare algorithms (Adamson & Smith 2018; Tat et al. 2020). Design of the AI can also lead to bias, for example if the chosen output of the model is a poor proxy for the clinical outcome of interest (Leslie et al. 2021). An algorithm that uses healthcare spending as a proxy for sickness overestimates the health of Black patients, resulting in resources being disproportionately allocated to white patients (Obermeyer et al. 2019).

iii. Growth in implementation

The FDA is increasingly approving AI-driven medical devices with over 750% growth in approvals from 2015 to 2019 (Figure 1A). As AI-driven tools are increasingly incorporated into healthcare, they have the potential to encode and exacerbate existing biases in disease diagnosis and classification, and in the allocation of healthcare resources. It is important to ensure that these tools display equitable performance across racial and ethnic subgroups, but evidence of bias in AI SaMD function indicates that existing regulations are not sufficient.

III. Current regulation of AI medical devices

Currently, there is no regulation in the US specific to AI in healthcare. Recent proposals by the FDA for regulatory frameworks aim to limit regulation, prioritize innovation, and quicken time to market (Health 2022d). In contrast, the European Union (EU) proposed an AI regulatory framework that enhances governance and enforcement of fundamental rights and safety laws by developing a single market for AI applications (European Commission 2021).

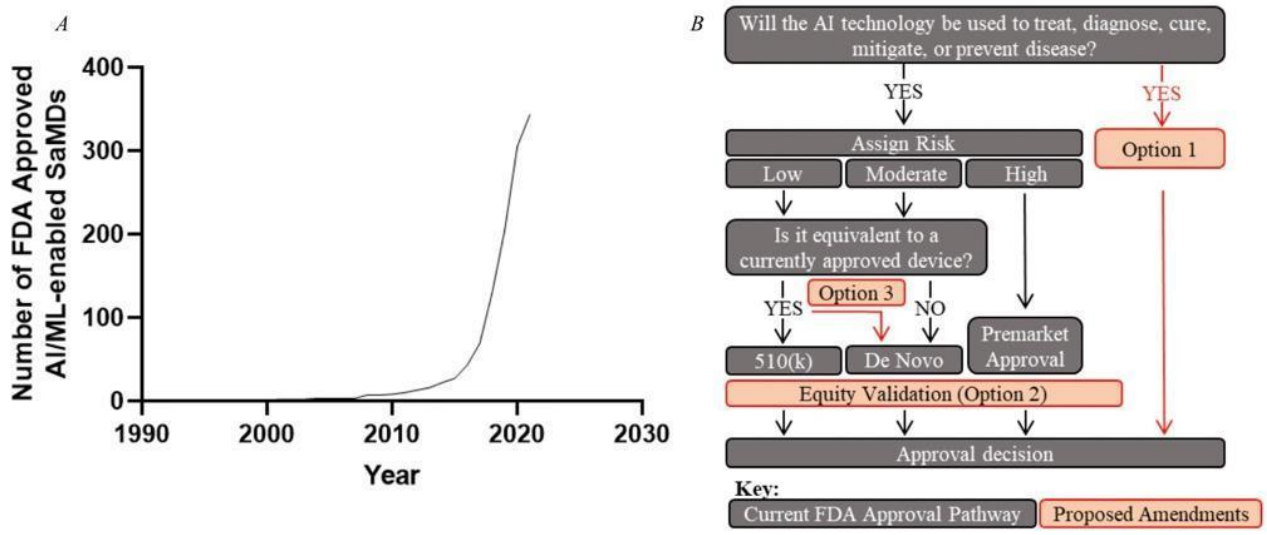

Figure 1: FDA Approval of AI-driven SaMDs. **A)** The number of FDA approved AI-enabled medical devices marketed in the United states over time as of September 22, 2021. **B)** Decision chart for determining the FDA review process for AI-driven SaMDs. Information sourced from Muehlematter et al. 2015.

Currently in the US, AI-driven SaMDs are subject to various degrees of regulatory scrutiny depending on their classification and risk level. If used to treat, diagnose, cure, mitigate, or prevent disease or other conditions, AI technology is regulated by the FDA under Software as a Medical Device (SaMD) as if it were traditional technology (Muehlematter et al. 2015; Pew 2022).

Most Class I (low-risk) and Class II (moderate-risk) AI-driven SaMDs (i.e., devices of a simpler design that present minimal to intermediate potential harm to users) undergo a 510(k) review. This is a premarket submission made to FDA which requires manufacturers to demonstrate that the AI-driven SaMD is safe, effective, and substantially equivalent to an existing device (Health 2022). New Class I-II SaMDs that are not "substantially equivalent" to a predicate device but whose underlying technology is well understood may go through the De Novo Classification Process. The De Novo process requires clinical and/or non-clinical data that reasonably ensures safety and effectiveness (Health 2022b). Class III (high-risk) AI-driven SaMDs (i.e., life-supporting/sustaining or important in preventing impairment of health) must undergo the full premarket approval process of scientific and regulatory review to prove that the benefits of the device outweigh the risks. This is the most stringent FDA marketing application and requires both clinical and non-clinical scientific evidence for safety and effectiveness (Health 2022c).

Once on the market, if changes are made to the device, the FDA makes a case-by-case decision on whether a review process is required (Muehlematter et al. 2015; Pew 2022). The process for re-evaluation is largely dependent on the device's risk classification and the nature of the change (Pew 2022). This, in addition to the multiple pre-market approval processes, leads to great risk of inconsistency in AI-driven SaMD FDA review.

As the presence of AI in healthcare continues to skyrocket (Figure 1A), the lack of regulation to prevent bias is a key issue in FDA regulation of AI-driven SaMDs. Current recommendations are vague and unclear. The FDA, Department of Health and Human Services, and World Health Organization all recommend that AI in healthcare be regulated to ensure equity yet have not implemented explicit standards to accomplish this goal (OCIO 2021; FDA 2021; WHO 2022). Only the Federal Trade Commission has provided specific recommendations such as testing algorithms before and after approval to ensure

they don't introduce new or worsen existing health inequities (FTC 2021).

Current FDA regulations do not address biased AI SaMD outputs. We can no longer rely on broad recommendations to ensure that these tools do not amplify existing health disparities between different racial and ethnic groups.

In this policy brief, we discuss the following policy options for the commissioner of the FDA to address bias in AI-driven SaMDs: 1) Creating a new and separate regulatory process for AI-driven SaMDs, 2) amending the proposed Total Product Life Cycle regulatory pathway to include equity validation, and 3) increasing the number of applications that go through premarket approval (Figure 1B). We also describe the policy option with the best strategy to mitigate healthcare bias in AI technology.

IV. Policy options

i. Option 1: Create a distinct FDA regulatory process for AI healthcare devices.

Currently, AI-driven SaMDs are reviewed in the same manner as non-AI SaMDs. This fails to address the unique challenges presented by AI compared to traditional software. We propose the FDA commissioner create a new regulatory process specifically dedicated to AI SaMDs. In this new framework, a panel that includes experts in algorithmic justice and healthcare equity will develop benchmarks and requirements for the investigation of bias at every development stage (i.e., from the datasets used to develop the device to post-market performance). Additionally, the panel will review all AI-driven SaMD functions to ensure that they demonstrate a low chance of exacerbating existing health disparities. Requirements could include evaluation of the training and testing datasets, reporting on existing disparities in the target outcome, and metrics to evaluate device accuracy across groups. Specific requirements may vary based on context, such as the risk level of the device and the target population of the product.

Advantages

This proposed policy would create a standardized and streamlined process that all AI-driven SaMDs would undergo, instead of being distributed to three possible processes (510(k), De Novo, or premarket approval), which simplifies the review process while maintaining rigor. We estimate that eliminating the decision of where to submit AI-driven SaMDs and the evaluation to determine if it was submitted to the correct process could reduce the time to approval by a month. This is significant as the approximate time to approval for 510(k), De Novo, or premarket approval is 90 days, 150 days, and 180 days, respectively. We believe this option is achievable as the previous commissioner of the FDA proposed and piloted a digital health precertification program intended to replace premarket submission for AI-driven SaMDs (Health 2022d). In this new proposed model, however, the FDA evaluates the developer instead of the individual SaMDs. Pending demonstrating the ability to develop safe and effective SaMDs, products would undergo a streamlined review process and could be placed on the market without premarket review.

Disadvantages

This policy option would require a significant amount of funding and effort from both FDA and SaMD developers. Determining appropriate evaluation standards will require nuanced ethical, legal, and technical considerations (Hoffman & Podgursky 2019). Specific methods of ensuring unbiased output may be difficult to achieve—for example, diversifying datasets through more open sharing of data is inhibited by privacy and proprietary concerns. This may hinder approval of devices and dissuade small businesses with fewer resources to overcome regulatory burdens from developing AI-driven SaMDs. Such companies are already less likely to achieve FDA approval of these devices compared to larger ones (Muehlematter et al. 2015).

ii. Option 2: Amend the FDA total product life cycle regulatory approach to include equity validation standards.

The FDA has adopted the International Medical Device Regulators Forum's (IMDRF) clinical evaluation framework for SaMDs within its Total Product Life Cycle regulatory approach (FDA 2019). This framework provides guidance on evaluating SaMDs across the realms of clinical association validation, analytical validation, and clinical validation. However, the IMDRF framework

lacks explicit instruction on equity validation to ensure SaMDs avoid mirroring existing healthcare bias. While the FDA currently supports regulatory science efforts to develop methods for identifying and eliminating bias in AI-based SaMDs (FDA 2020), it does not yet incorporate explicit regulations for detecting and evaluating bias in SaMD outcomes.

We propose the FDA add a category for equity validation in SaMD regulation. Equity validation should include 1) assessments that ensure SaMDs effectively and accurately achieve their intended outcomes and 2) evidence that SaMDs produce the intended outcomes equitably across populations (e.g., accurate chest x-ray diagnoses in all populations). Further, high-risk AI SaMDs should not rely only on retrospective data but undergo prospective randomized trials that demonstrate the medical device does not replicate existing biases.

Advantages
The proposed regulations would ensure developers evaluate for and address bias in AI SaMD output. This maintains discretion for proprietary algorithmic designs while assuring the devices avoid exacerbating existing disparities. Additionally, it allows SaMD developers to use novel methods of achieving equitable accuracy.

Disadvantages
This strategy focuses regulation on the output of AI-driven SaMDs and is independent of the methods used to achieve equitably accurate predictions. This may allow other sources of bias to persist and may be less robust at preventing AI-exacerbated healthcare disparities.

iii. Option 3: Expand the scope of review within existing regulatory pathways.
Currently, many AI-driven SaMDs are approved through the 510(k) premarket submission based on having predecessors on the market that are "substantially equivalent." The definition of equivalency is vague and broad, allowing novel AI-driven SaMDs to bypass premarket review and clinical trials. A device approved under 510(k) may remain on the market even if its predecessor is demonstrated to be unsafe or biased. In addition, clinical decision software is currently exempt from review if it "supports or provides medical recommendations to a health care professional" but is not the primary driver behind medical decisions (PEW 2021). This definition is also vague and allows AI-driven devices that influence treatment and healthcare outcomes to be introduced on the market without rigorous review.

We propose that the FDA commissioner lower the thresholds for FDA review so that more devices undergo thorough De Novo or premarket approval. We recommend narrowing the equivalency standard for 510(k) review and the criteria for exemption.

Advantages
This approach would require the least change to current regulations, enabling FDA efforts to maintain focus on other important issues. These changes would also apply to non-AI SaMDs, increasing the rigor of review for all medical software devices.

Disadvantages
This approach does not explicitly include consideration of bias or disparities in outcomes during the review process. Additionally, time to market would be increased due to some devices requiring longer and more thorough review processes.

V. Policy recommendation

We recommend Option 1: for the commissioner of the FDA to develop a distinct regulatory process specifically designed for AI-driven SaMDs. This policy option would help prevent approval of AI medical devices that mirror or exacerbate existing health disparities. This could also increase trust from the public for whom AI bias and its potentially negative impact on health disparities is a significant concern (Richardson et al. 2021). Moreover, these devices have the potential to reverse existing healthcare biases if regulated properly. We believe this is of utmost importance to the commissioner of the FDA whose mission is to protect and promote public health (FDA 2021). As healthcare is on the cusp of grand technological transformation, the FDA has an opportunity to craft a safe and equitable healthcare environment for all Americans.

References

Adamson, Adewole S., and Avery Smith. 2018. "Machine Learning and Health Care Disparities in Dermatology." JAMA Dermatology 154 (11): 1247–48. https://doi.org/10.1001/jamadermatol.2018.2348.

"Aiming for Truth, Fairness, and Equity in Your Company's Use of AI." 2021. Federal Trade Commission. April 19, 2021. http://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai.

Amisha, Paras Malik, Monika Pathania, and Vyas Kumar Rathaur. 2019. "Overview of Artificial Intelligence in Medicine." Journal of Family Medicine and Primary Care 8 (7): 2328–31. https://doi.org/10.4103/jfmpc.jfmpc_440_19.

"Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan". 2021. Accessed May 3, 2022. PDF file. https://www.fda.gov/media/145022/download

Aydın, Mustafa, Fırat Hardalaç, Berkan Ural, and Serhat Karap. "Neonatal Jaundice Detection System." Journal of Medical Systems 40, no. 7 (July 2016): 166. https://doi.org/10.1007/s10916-016-0523-4.

Berry, Jay G., Bloom Sheila, Foley Susan, Palfrey Judith S. 2010. "Health Inequity in Children and Youth With Chronic Health Conditions". Pediatrics 126, no. 3 (December 2010): S111-S119. https://publications.aap.org/pediatrics/article-abstract/126/Supplement_3/S111/29999/Health-Inequity-in-Children-and-Youth-With-Chronic?redirectedFrom=fulltext

Briganti, Giovanni, and Olivier Le Moine. 2020. "Artificial Intelligence in Medicine: Today and Tomorrow." Frontiers in Medicine 7. https://www.frontiersin.org/article/10.3389/fmed.2020.00027.

Chen, Mei, and Michel Decary. "Artificial Intelligence in Healthcare: An Essential Guide for Health Leaders." Healthcare Management Forum 33, no. 1 (January 2020): 10–18. https://doi.org/10.1177/0840470419873123.

———. "Embedding Health Literacy Tools in Patient EHR Portals to Facilitate Productive Patient Engagement." Studies in Health Technology and Informatics 257 (2019): 59–63.

Commissioner, Office of the. 2022. "Robert Califf." U.S. Food and Drug Administration. March 10, 2022. https://www.fda.gov/about-fda/fda-organization/robert-califf.

European Commission. "Proposal For a Regulation of The European Parliament and of The Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts". (2021). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206

FDA. "Artificial Intelligence/Machine Learning (Ai/Ml)-Based Software as a Medical Device (Samd) Action Plan." (2021). https://www.fda.gov/media/145022/download

FDA. "Developing the Software Precertification Program: Summary of Learnings and Ongoing Activities." (2020). https://www.fda.gov/media/142107/download

Flores, Laura E., Walter R. Frontera, Michele P. Andrasik, Carlos Del Rio, Antonio Mondríguez-González, Stephanie A. Price, Elizabeth M. Krantz, Steven A. Pergam, and Julie K. Silver. 2021. "Assessment of the Inclusion of Racial/Ethnic Minority, Female, and Older Individuals in Vaccine Clinical Trials." JAMA Network Open 4 (2): e2037640. https://doi.org/10.1001/jamanetworkopen.2020.37640.

Gianfrancesco, Milena A., Suzanne Tamang, Jinoos Yazdany, and Gabriela Schmajuk. "Potential Biases in Machine Learning Algorithms Using Electronic Health Record Data." JAMA Internal Medicine 178, no. 11 (November 1, 2018): 1544–47. https://doi.org/10.1001/jamainternmed.2018.3763.

Health, Center for Devices and Radiological. "Premarket Notification 510(k)." FDA. FDA, October 3, 2022. https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/premarket-notification-510k..

———. "De Novo Classification Request." FDA, October 3, 2022. https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/de-novo-classification-request.

———. "Premarket Approval (PMA)." FDA. FDA, October 3, 2022. https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/premarket-approval-pma.

———. "Digital Health Software Precertification (Pre-Cert) Pilot Program." FDA. FDA, March 8, 2022. https://www.fda.gov/medical-devices/digital-health-center-excellence/digital-health-software-precertification-pre-cert-pilot-program.

Hoffman, Sharona, and Andy Podgurski. "Artificial Intelligence and Discrimination in Health Care." *Yale J. Health Pol'y L. & Ethics* 19 (2019): 1. https://openyls.law.yale.edu/bitstream/handle/20.500.13051/5964/Hoffman_v19n3_1_49.pdf?sequence=2&isAllowed=y#page=13

Hoodbhoy, Zahra, Sarah Masroor Jeelani, Abeer Aziz, Muhammad Ibrahim Habib, Bilal Iqbal, Waqaas Akmal, Khan Siddiqui, Babar Hasan, Mariska Leeflang, and Jai K. Das. "Machine Learning for Child and Adolescent Health: A Systematic Review." Pediatrics 147, no. 1 (January 1, 2021): e2020011833. https://doi.org/10.1542/peds.2020-011833.

"How FDA Regulates Artificial Intelligence in Medical Products." n.d. Accessed May 14, 2022. https://pew.org/3yglbCS.

LaVeist, Thomas A., Darrell Gaskin, and Patrick Richard. 2011. "Estimating the Economic Burden of Racial Health Inequalities in the United States." International Journal of Health Services: Planning, Administration, Evaluation 41 (2): 231–38. https://doi.org/10.2190/HS.41.2.c.

Leslie, David, Anjali Mazumder, Aidan Peppin, Maria K. Wolters, and Alexa Hagerty. "Does 'AI' Stand for Augmenting Inequality in the Era of Covid-19 Healthcare?" BMJ 372 (March 16, 2021): n304. https://doi.org/10.1136/bmj.n304.

"Minority Population Profiles - The Office of Minority Health." 2021. https://minorityhealth.hhs.gov/omh/browse.aspx?lvl=2&lvlID=26.

Muehlematter, Urs J., Paola Daniore, and Kerstin N. Vokinger. 2021. "Approval of Artificial Intelligence and Machine Learning-Based Medical Devices in the USA and Europe (2015–20): A Comparative Analysis." The Lancet Digital Health 3 (3): e195–203. doi.org/10.1016/S2589-7500(20)30292-2

Murphy, Sherry L, Kochanek, Kenneth, Xu, Jiaquan, Arias, Elizabeth. 2021. "Mortality in the United States, 2020". NCHS Data Brief 427: 2.. https://www.cdc.gov/nchs/data/databriefs/db427.pdf

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." Science (New York, N.Y.) 366 (6464): 447–53. https://doi.org/10.1126/science.aax2342.

Officer (OCIO), Office of the Chief Information. 2021. "HHS Artificial Intelligence (AI) Strategy." Text. HHS.Gov. December 22, 2021. https://www.hhs.gov/about/agencies/asa/ocio/ai/strategy/index.html.

Ong, Mei-Sing, Mary P. Mullen, Eric D. Austin, Peter Szolovits, Marc D. Natter, Alon Geva, Tianxi Cai, Sek Won Kong, and Kenneth D. Mandl. "Learning a Comorbidity-Driven Taxonomy of Pediatric Pulmonary Hypertension." Circulation Research 121, no. 4 (August 4, 2017): 341–53. https://doi.org/10.1161/CIRCRESAHA.117.310804.

Price, James H., Jagdish Khubchandani, Molly McKinney, and Robert Braun. 2013. "Racial/Ethnic Disparities in Chronic Diseases of Youths and Access to Health Care in the United States." BioMed Research International 2013 (September): e787616. https://doi.org/10.1155/2013/787616.

"Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD): Discussion Paper and Request for Feedback". n.d. Accessed May 3, 2022. PDF file. https://www.fda.gov/media/122535/download

Richardson, Jordan P., Cambray Smith, Susan Curtis, Sara Watson, Xuan Zhu, Barbara Barry, and Richard R. Sharp. 2021. "Patient Apprehensions about the Use of Artificial Intelligence in Healthcare." Npj Digital Medicine 4 (1): 1–6. https://doi.org/10.1038/s41746-021-00509-1.

Services, U.S. Department of Health and Human. "Artificial Intelligence (AI) Strategy." (2021). https://www.hhs.gov/sites/default/files/hhs-ai-strategy.pdf

Seyyed-Kalantari, Laleh, Haoran Zhang, Matthew B. A. McDermott, Irene Y. Chen, and Marzyeh Ghassemi. 2021. "Underdiagnosis Bias of Artificial Intelligence Algorithms Applied to Chest Radiographs in Under-Served Patient Populations." Nature Medicine 27 (12): 2176–82. https://doi.org/10.1038/s41591-021-01595-0.

Singh, Devin, Sujay Nagaraj, Pouria Mashouri, Erik Drysdale, Jason Fischer, Anna Goldenberg, and Michael Brudno. "Assessment of Machine Learning–Based Medical Directives to Expedite Care in Pediatric Emergency Medicine." JAMA Network Open 5, no. 3 (March 16, 2022): e222599. https://doi.org/10.1001/jamanetworkopen.2022.2599.

"Software as a Medical Device (SAMD): Clinical Evaluation : Guidance for Industry and Food and Drug Administration Staff - Digital Collections - National Library of Medicine." n.d. Accessed May 3, 2022. https://collections.nlm.nih.gov/catalog/nlm:nlmuid-101720008-pdf.

"Software as a Medical Device (SaMD): Key Definitions." 2013, 9. https://www.imdrf.org/sites/default/files/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf

Tat, Emily, Deepak L. Bhatt, and Mark G. Rabbat. 2020. "Addressing Bias: Artificial Intelligence in Cardiovascular Medicine." The Lancet Digital Health 2 (12): e635–36. https://doi.org/10.1016/S2589-7500(20)30249-1.

"Using Artificial Intelligence and Algorithms." 2020. Federal Trade Commission. April 8, 2020. http://www.ftc.gov/business-guidance/blog/2020/04/using-artificial-intelligence-algorithms.

Vokinger, Kerstin N., Stefan Feuerriegel, and Aaron S. Kesselheim. 2021. "Mitigating Bias in Machine Learning for Medicine." Communications Medicine 1 (1): 1–3. https://doi.org/10.1038/s43856-021-00028-w.

Vyas, Darshali A., Leo G. Eisenstein, and David S. Jones. 2020. "Hidden in Plain Sight — Reconsidering the Use of Race Correction in Clinical Algorithms." New England Journal of Medicine 383 (9): 874–82. https://doi.org/10.1056/NEJMms2004740.

"WHO Issues First Global Report on Artificial Intelligence (AI) in Health and Six Guiding Principles for Its Design and Use." n.d. Accessed May 14, 2022. https://www.who.int/news/item/28-06-2021-who-issues-first-global-report-on-ai-in-health-and-six-guiding-principles-for-its-design-and-use.

Yala, Adam, Constance Lehman, Tal Schuster, Tally Portnoi, and Regina Barzilay. "A Deep Learning Mammography-Based Model for Improved Breast Cancer Risk Prediction." Radiology 292, no. 1 (July 2019): 60–66. https://doi.org/10.1148/radiol.2019182716.

Acknowledgements
We thank Victor Acero, University of Pennsylvania, for their assistance drafting and editing this manuscript. We also thank the Penn Science Policy and Diplomacy Group for providing the opportunity and instruction to work together.

Kristina Dortche Kristina Dortche is completing residency in Urology at the Cleveland Clinic Foundation in Cleveland, OH. She obtained her MD at the Perelman School of Medicine where she was involved in developing coursework and national conferences highlighting disparities in health and healthcare. She also earned her MPH at Columbia Mailman School of Public Health in Health Policy and Management. She hopes to use her training to identify and address health disparities that plague pediatric patients in Urology.

Grace McCarthy Dr. Grace McCarthy is a postdoctoral scholar at Oregon Health & Science University, where she obtained her Ph.D. in Cancer Biology. She works with groups such as Alliance for Visible Diversity in Science, the National Postdoctoral Association, and her department's DEI Committee to advocate for diverse, equitable, and safe research environments. She also serves as the Mission Chair for the Portland Affiliate of the Pancreatic Cancer Action Network, overseeing pancreatic cancer education, outreach, and advocacy. Dr. McCarthy is dedicated to increasing access and equity in healthcare.

Sara Banbury Sara Banbury is a third year medical student at the Perelman School of Medicine at the University of Pennsylvania. She has previous experience in community organizing and education, and has worked with Students Opposing Racism in Medicine during her time at medical school. She is interested in the intersection of medicine and advocacy, specifically increasing physician advocacy regarding issues that affect their patient populations. She is currently pursuing her MD, and hopes to integrate advocacy and health equity work into her future career.

Isabel Yannatos Isabel Yannatos is a Ph.D. candidate in biomedical studies at the University of Pennsylvania. She works in Dr. Corey McMillan's lab on racial disparities in aging. Her work focuses on how environmental exposures, such as neighborhood resources and air pollution, contribute to disparities in epigenetic aging and cognitive outcomes. Isabel is passionate about the intersection of research and policy in addressing health disparities. She is pursuing a certificate in public health and cognitive aging and other opportunities in health policy and advocacy. In her career she hopes to translate research into action to promote health equity.