

August 28, 2023

The Promise and Peril of Advancing Health Equity through Artificial Intelligence

By Roma Sharma and Tienne Anderson

Share:



Introduction

Today, whether a child survives cancer depends, in large part, on where that child lives. Children with cancer in high-income countries (HIC) have an over 80% chance of survival. That rate is flipped for children in lower- and middle-income countries (LMIC), where children with cancer have an 80% chance of *mortality*. ¹ This statistic is especially devastating considering that worldwide, more than 90% of children with cancer live in LMIC. ² Artificial intelligence (AI) has the potential to drastically affect this survival disparity and impact other health inequities in a multitude of ways, presenting a unique opportunity to alter the trajectory of humankind. The direction of that change rests upon the intentionality by which we invest in and use AI to address healthcare inequality and the manner in which governments ultimately regulate this technology of infinite possibility. Failure to harness this technology for good and manage its risks may undermine the foundational principles of medicine ³ and exacerbate inequities that drive health disparities around the world.

This article will examine the potential impact of AI on health equity and highlight approaches that should be considered as AI regulations are developed. First, this article briefly describes the fundamental concepts of AI and health equity and some of the challenges that lay at their crossroads. Next, this article provides an overview of various representative legal and regulatory approaches to AI around the world with a view towards healthcare. Finally, it discusses key concepts that should guide the development and regulation of AI to promote health equity: representative data, strengthening political will, accountability and transparency, and elevating marginalized voices. The conclusion proffers two paradigms for consideration in designing a framework for the use of individual health data by AI: data as a common good and data as a

personal asset. Such paradigms, or elements thereof, may be considered to balance individual control with reliance on corporate responsibility for the appropriate collection and use of such data by AI systems.

What is AI and What is Health Equity?

Artificial Intelligence

Artificial intelligence sparked imaginations upon its inception in the 1950s. ⁴ AI continues to evolve and capture our collective attention as its capabilities expand, becoming more human-like over time. Some believe that eventually, AI will surpass human intelligence. The recent release of OpenAI's ChatGPT in November 2022 set off a global race in artificial intelligence. ChatGPT is a natural language processing tool driven by AI that interacts with the user in a conversational manner. ⁵ ChatGPT can answer complex questions in nearly any field with reasonable accuracy, draft papers and articles, create poetry and art (arguably), and converse with humans, among other things. ChatGPT has ignited our collective imaginations—prompting us to consider its wide potential for both productive and destructive use. While we are just starting to scratch the surface of the benefits of AI, including ChatGPT, various experts have issued warnings about artificial intelligence and the negative ramifications if deployed without sufficient regulation and oversight. ⁶

Taking a step back from the potential and risks of artificial intelligence—what *is* artificial intelligence? Put simply, AI is a field that combines computer science and robust datasets to enable problem-solving. ⁷ AI has become a catchall term for applications that perform complex tasks that once required human input, such as communicating with customers online or playing chess. ⁸ While there is no single definition of AI, defined broadly, AI is the capability of a machine to imitate intelligent human behavior to perform complex tasks in a way that is similar to how humans solve problems. ⁹

Machine learning (ML) is a subfield of AI. ML was defined by AI pioneer Arthur Samuel in the 1950s as “the field of study that gives computers the ability to learn without explicitly being programmed.” ¹⁰ ML is behind chatbots and predictive text, language translation apps, the shows Netflix suggests, and how social media feeds are presented. ¹¹ ML powers autonomous

vehicles and machines that can diagnose medical conditions based on images. ¹² The term “ML” is commonly used interchangeably with “AI.” In this article, references to AI may include ML and/or other subfields of AI.

The source of AI is data, e.g., numbers, images, and text. Programmers choose a machine learning model to use, supply the data to train the model, and allow the model to find patterns. Data is fed into the program for evaluation purposes as well, that is, to assess the accuracy of the program.

¹³ Machine learning programs can generally be descriptive, where the system uses data to describe what happened; predictive, where the system uses data to predict what will happen; or prescriptive, where the system uses data to make suggestions on actions to take. ¹⁴

There are a great number of current and potential uses of AI in healthcare. Broadly categorized, the four areas where AI is thought to have the greatest potential impact in healthcare are (1) healthcare administration, (2) clinical decision support, (3) patient monitoring, and (4) healthcare interventions. ¹⁵ With respect to healthcare administration, AI can be used to analyze patterns to detect healthcare fraud or errors, for patient data entry, to record clinical notes, for insurance claims processing, to support scheduling and patient triaging, and more. Clinical decision support algorithms may help to diagnose disease, reduce diagnostic and treatment errors, increase efficiency, personalize treatment, and suggest alternative treatment plans that improve outcomes. AI-powered systems are already available to help medical professionals diagnose cancer, diabetic retinopathy, Alzheimer’s disease, heart disease, and COVID-19, among others. ¹⁶ AI can analyze images and be trained to read medical scans or identify such things as tumors, bone fractures, or other markers of illness. AI-powered software could be used to triage patients, provide emergency assistance, and provide counseling and emotional wellness assistance. Indeed, there is solid evidence that certain AI programs can at least match the diagnostic performance accuracy of radiologists and pathologists. ¹⁷ AI-powered patient monitoring may include all manners of wearable/personal health monitoring devices; testing and monitoring devices used in various settings (i.e., electrocardiographs, Doppler ultrasounds, and respiratory monitoring); and outpatient devices and systems used to augment patient compliance. Finally, tailored healthcare interventions generated through AI with biometric and other personalized data can recommend treatment plans and reduce diagnostic wait times. ¹⁸

All industries are likely to be impacted by AI in the long run. We are just beginning to scratch the surface of its uses in healthcare. Responsible use of AI has enormous potential to improve the

health of humankind—improving significantly the rate of treating and preventing disease on a large scale.

Health Equity

Although there is no one agreed-upon definition of “health equity,” conceived in its broadest sense, it is the ability of every person to reach the highest possible standard of health, with adequate support and consideration for the needs of those at greatest risk of poor health, based on social conditions. ⁽¹⁹⁾ Stated differently, “[h]ealth equity is defined as the absence of unfair and avoidable or remediable differences in health among population groups defined socially, economically, demographically or geographically.” ⁽²⁰⁾

The social factors that impact health, referred to as the social determinants of health (SDOH) include, but are not limited to, income and social protection; education; unemployment and job insecurity; working life conditions; food insecurity, housing, basic amenities and the environment; early childhood development; social inclusion and non-discrimination; structural conflict; and access to affordable health services of decent quality. ⁽²¹⁾ The global COVID-19 pandemic highlighted, nearly in real-time, how SDOH disparately impact health and health outcomes for marginalized communities. ⁽²²⁾ The Centers for Disease Control and Prevention (CDC) has estimated that more than 50% of poor health outcomes are powered by SDOH factors—factors we can control and change in a more equitable society. ⁽²³⁾

As identified with the example of childhood cancer, income is a key variable of health equity, as research has shown that “the relationship between income and health is a gradient: they are connected step-wise at every level of the economic ladder.” ⁽²⁴⁾ Beyond income and health, AI’s possibilities seem best understood as a way to support the road to peace and prosperity, perhaps best typified by the United Nations (UN) 2030 Agenda for Sustainable Development and the 17 Sustainable Development Goals (SDGs) that accompany this bold initiative. ⁽²⁵⁾ Taken as a whole, the SDGs represent a shared vision for desired outcomes for life on this planet, including the challenges presented by SDOH that negatively impact health outcomes. While AI can play a role in helping to achieve the SDGs, research supported by KTH Climate Action Centre and Data Futures has also found that AI may inhibit 58 distinct target goals agreed upon internationally. ⁽²⁶⁾ This study brought together a diverse array of experts to score the potential impact of AI through select literature. Now, the lead researcher and others are pursuing the goal of conducting a similar review, but this time supported by AI such that thousands of pieces of literature may be

examined. The ongoing study is intended to process and synthesize the “trade-offs among the SDGs humans may not pick up on, with the objective of developing an algorithm to help avoid unexpected negative interactions among the SDGs in policy decisions.” (27) The potential for AI to generate both positive and negative healthcare outcomes will play out based on the extent to which societies have thoughtfully harnessed AI to address the SDOH and managed downside risk.

As a society, we must commit ourselves to answering the question of why the SDOH continue to negatively impact healthcare—despite our growing understanding of the SDOH, the health inequities of marginalized populations, and how these challenges intertwine to rob society of human resources. In *Unequal Treatment: Confronting Racial and Ethnic Disparities in Health Care*, (28) published by the National Institutes of Health (NIH) 20 years ago, data collection and monitoring were found to be key challenges impacting the health of racial and ethnic populations, issues that now have new import as we consider the rise of AI and its ability to rapidly analyze and process large amounts of data. (29)

In March 2023, stakeholders met to consider the scale of today’s health equity challenges, in light of the *Unequal Treatment* report, and highlighted three key difficulties that continue to drive health disparities:

- 1 “Political will to address healthcare disparities has vacillated among key stakeholders, particularly as other equity issues have moved to the forefront;
- 2 Neither the 2003 *Unequal Treatment* report nor any mechanism since has identified or addressed the lack of accountability among those in a position to make change (including healthcare payors, policy makers, system administrators, providers, and training institutions); and
- 3 The voices of those who have been most affected by health disparities have rarely been heard or valued, despite their critical perspectives on addressing health disparities by informing research questions and assisting in developing critical solutions.” (30)

These challenges must inform the regulation and governance of AI in a way that will allow us to move beyond rhetoric to action and harness the power of AI to advance health equity and tackle the barriers that cause inequality.

As noted, the impact of AI is not limited to health or healthcare. AI is fast becoming not only ubiquitous, but embedded in our everyday lives and transactions in ways that impact the SDOH. AI has been shown to have potentially discriminatory and/or negative effects in the following, among other ways:

- Being less accessible to, and thus less able to benefit, lower socioeconomic status groups; ³¹
- Law enforcement, through facial recognition technologies, recidivism algorithms that flag minority groups more than whites, and other uses; ³²
- Credit access, whereby bias can arise in the model development or the data used to generate it; ³³
- Employment, whereby resumes listing women's colleges or participation in women's sports were scored lower than others with similar content because the program algorithm, trained mostly on men's resumes, taught itself that male candidates were preferable; ³⁴
- Housing, through biased outcomes in mortgage application review; ³⁵
- Social cohesion, through social media engagement only with the like-minded, and privacy erosion. ³⁶

There is strong potential for inequality to be exacerbated through a *layering effect* of biased and/or inequitable AI systems that further compounds inequality over time through the SDOH. Use of AI combined with existing institutional inequities could further entrench and expand healthcare inequality around the globe, even without ill intent, if proactive measures are not taken to dismantle the existing biases and discriminatory systems that continue to drive unequal outcomes. ³⁷

Brief Overview of Global AI Regulatory Frameworks and Other Approaches

The global race to develop AI technology set off a parallel global race for governments to regulate AI. This section examines various representative regulatory frameworks and other approaches from around the world to provide a global picture of the numerous ethical considerations, perceived risks, and legal and policy issues currently under consideration as governments

contemplate the regulation of this fast-growing and uncertain space. The enormous potential of AI and its global healthcare implications, across countries with widely varied resources and healthcare access, necessitates an approach with acknowledgement of, and appreciation for, the various motivations behind proposed legislation and governance models. The potential for negative consequences and inequality through the use of, and access to AI is paramount to the challenges we face; to truly understand the potential impact, it has been argued that regulatory *insight* must precede regulatory *oversight*. (38) Regulators must develop an understanding of the challenges and opportunities presented by AI that is informed by stakeholders at all levels of society both locally and globally, given AI's reach in today's interconnected world. Regulation of AI without such insight can create serious harms, the likes of which we do not yet understand.

The global landscape currently consists of a patchwork of various approaches in this nascent field, with many governments taking a wait-and-see approach, while others, such as the European Union, Australia, and China, taking active steps to regulate AI.

Setting the stage for a global overview, the World Health Organization (WHO) provides a unique vantage point as a global actor that understands the significant benefits and risks that AI/ML may bring. (39) In its *Ethics and Governance of Artificial Intelligence for Health*, WHO lists six key principles for the use of AI in healthcare:

- 1 Protect autonomy of human decision making in healthcare systems and medical decisions;
- 2 Promote human well-being, human safety, and the public interest: AI technologies should not harm people and should satisfy all requirements for safety, accuracy, and efficacy before use;
- 3 Ensure transparency, explainability, and intelligibility for AI technologies;
- 4 Foster responsibility and accountability, such that patient harm does not go unaddressed [faultless/collective responsibility model recommended];
- 5 Ensure inclusiveness and equity by developing and monitoring AI technologies through as many diverse lenses as possible and sharing open-source software and/or source codes as widely as possible; and
- 6 Promote AI that is responsive and sustainable through ongoing impact assessments to users and the broader environment. (40)

Additional global networks and consortiums continue to develop, across industries, to support a comprehensive global AI approach. One such network is the Global Partnership for Artificial Intelligence (GPAI). ⁴¹ GPAI formed in 2020 through conversations within the G7 and hosted at the Organisation for Economic Cooperation and Development (OECD); it aims to:

- “[S]upport and guide the responsible development, use and adoption of AI that is human-centric and grounded in human rights, inclusion, diversity and innovation, while encouraging sustainable economic growth;
- [F]acilitate international collaboration in a multistakeholder manner; and
- [M]onitor and draw on work being done domestically and internationally to identify knowledge gaps, maximise coordination, and facilitate international collaboration on AI.”

⁴²

While some of the WHO principles have been incorporated into the various countries that are examined below, there is also variation in how countries view their risk/reward calculus based on their economic circumstances, among other factors, which will undoubtedly affect regulatory and enforcement priorities. With the work of organizations such as GPAI, there is hope that more robust information and analysis can inform the development of regulations that will help make AI beneficial to all, including in healthcare.

Developing Country Perspectives – Africa

Africa is the second largest continent in the world and has the world’s largest share of developing countries (See Figure 1). As a result of this economic backdrop, the benefits of AI for Africa have already been estimated to be “modest . . . due [to] the much lower rate of adoption of AI technologies expected.” ⁴³ The global vision for AI must take into account the expected modest rate of adoption in Africa, a stark contrast to what is simultaneously viewed as a modern-day gold rush by many in developed countries. ⁴⁴

■ Low income
 ■ Lower middle income
 ■ Upper middle income
 ■ High income

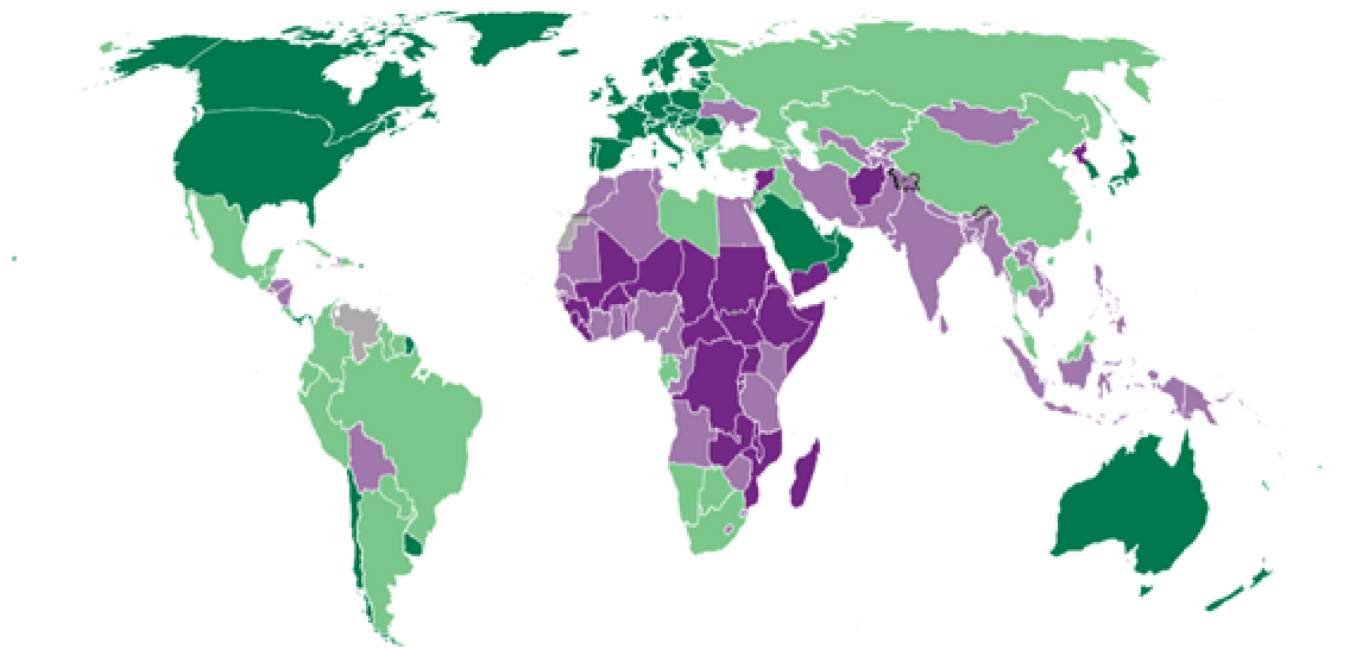


Figure 1: The World by Income 2021

Source: The World Bank, available at: <https://datatopics.worldbank.org/world-development-indicators/the-world-by-income-and-region.html>

Like most other countries around the world, African countries are in the initial stage of understanding the costs and benefits of AI, the regulatory landscape, and evaluating future regulations. Smart Africa, an alliance of 36 African countries, developed the *2021 Blueprint: Artificial Intelligence for Africa*. The blueprint claims a bold vision for Africa's opportunities and the ability to address challenges by focusing on five distinct "Framework Pillars," which are:

- 1 "Human capital, underscoring the importance of educational development and enhancing the proficiencies, competencies and understanding of individuals who use and develop artificial intelligence solutions;
- 2 Lab to Market initiatives that foster research, development, innovation, and commercialization;
- 3 Networking, cooperation, and collaboration, in pursuit of joint partnerships across private and/or public sectors to favorably impact the uptake of AI among all;
- 4 Infrastructure investments that will foster the development of digital and telecommunication systems which support efficient data collection and usage; and

5 Regulation that is effective, infused with an ethics that support equality, and international best practices.” (45)

Of import is the fundamental recognition that AI can only be as intelligent as the humans who power it, the human data that drives it, the human need for innovation that brings forth the right questions, and the requirement that our ethical concerns center on the need to support equity in the use of AI.

Safety/Patient Centered Approaches – Canada, Australia/New Zealand

While Canada seeks a position of influence in the world of AI, there is recognition in both the private and government sectors that a normative framework is necessary. (46) Like in many other countries, this has manifested initially in the adoption and consideration of various laws and regulations to protect privacy, with a focus on the future and crafting an appropriate balance between innovation and safety. The government summarizes its approach as one which must:

- “Understand and measure the impact of using AI by developing and sharing tools and approaches;
- Be transparent about how and when we are using AI, starting with a clear user need and public benefit;
- Provide meaningful explanations about AI decision making, while also offering opportunities to review results and challenge these decisions;
- Be as open as [possible] by sharing source code, training data, and other relevant information, all while protecting personal information, system integration, and national security and defen[s]e; and
- Provide sufficient training so that government employees developing and using AI solutions have the responsible design, function, and implementation skills needed to make AI-based public services better.” (47)

This approach is reflected in Canada’s Directive on Automated Decision-Making, which seeks a sensible, middle-of-the-road approach in its administrative AI use “that reduces risks to Canadians and federal institutions, and leads to more efficient, accurate, consistent, and interpretable

decisions made pursuant to Canadian law.” ⁴⁸ Canada’s sensibilities are also reflected in its Algorithmic Impact Assessment tool, which aims to root out bias and inequality on the front end of automated decision-making systems by providing human and ethical measures for developers to build in and test against in order to gain government approval. ⁴⁹ These frameworks and tools can be used in conjunction with existing laws to provide some protection against increasing inequality and harm in the healthcare space. Existing laws include: The Personal Information Protection and Electronic Documents Act, The Canadian Consumer Product Safety Act, The Food and Drugs Act, The Motor Vehicle Safety Act, The Bank Act, The Canadian Human Rights Act and provincial human rights laws, and The Criminal Code. ⁵⁰

Canada has also proposed amendments to its regulations for medical devices that will allow the minister to:

“[A]t any time, impose terms and conditions on a medical device license, or amend those terms and conditions, after considering the following factors:

- (a) whether there are uncertainties relating to the benefits or risks associated with the device;
- (b) whether the requirements under the Act are sufficient to
 - (i) maintain the safety and effectiveness of the device,
 - (ii) optimize the benefits and manage the risks associated with the device, and
 - (iii) identify changes and manage uncertainties relating to the benefits and risks associated with the device;
- (c) whether the proposed terms and conditions may contribute to meeting the objectives set out in subparagraphs (b)(i) to (iii);
- (d) whether compliance with the proposed terms and conditions is technically feasible; and
- (e) whether there are less burdensome ways to meet the objectives of the proposed terms and conditions.” ⁵¹

Health Canada enforcement priorities are risk-based and stand upon existing departmental policies. Violation of the terms and conditions imposed by the minister could ultimately result in prosecution. (52)

Finally, to close any gaps left by the various laws and regulations that will apply in Canada, there is the proposed Artificial Intelligence and Data Act (AIDA). AIDA, which if passed, is expected to come into force no sooner than 2025, was designed to align with “evolving international norms in the AI space,” noting the regulations proposed in the European Union (EU), the United Kingdom (UK), and the United States as supporting Canada’s need to adopt “a corresponding framework to enable citizen trust, encourage responsible innovation, and remain interoperable with international markets.” (53)

In a healthcare-related example tied to other countries that appear to be taking a patient-centered approach, The Royal Australian and New Zealand College of Radiologists (RANZCR) felt compelled to act on AI years ago, as the benefits to their profession felt both imminent, and at the same time, “getting this wrong for the population in question gives significant potential for harm.” (54)

Whether the “population” referred to were those whose lungs were to be compared to algorithms developed based on a population with a higher incidence of smoking, tuberculosis, or opportunistic lung infections, the details of the data matter, and without context, “[t]he unlabeled data on their own... are meaningless, no matter how voluminous the amount of information.”

(55) Given then the potential for both enormous benefit and great harm, RANZCR proactively moved to define an approach to AI in the clinical radiology and radiation oncology space in Australia and New Zealand that would first and foremost, be safe for patients. RANZCR’s Ethical Principles for Artificial Intelligence in Medicine (56) call for:

“Principle 1: Safety. The first and foremost consideration in the development, deployment or utilization of ML or AI must be patient safety and quality of care, with the evidence base to support this.

Principle 2: Privacy and Protection of Data. A patient’s data must be stored securely and in line with relevant laws and best practice.

Principle 3: Avoidance of Bias. To minimize bias, the same standard of evidence used for other clinical interventions must be applied when regulating ML and AI, and their limitations must be transparently stated.

Principle 4: Transparency and Explainability. When designing or implementing ML or AI, consideration must be given to how a result that can impact patient care can be understood and explained by a discerning medical practitioner.

Principle 5: Application of Human Values. The doctor must apply humanitarian values (from their training and the ethical framework in which they operate) to any circumstances in which ML or AI is used in medicine, but they also must consider the personal values and preferences of their patient in this situation.

Principle 6: Decision-Making on Diagnosis and Treatment. While ML and AI can enhance decision-making capability, final decisions about care are made after a discussion between the doctor and patient, taking into account the patient's presentation, history, options and preferences.

Principle 7: Teamwork. To deliver the best care for patients, each team member must understand the role and contribution of their colleagues and leverage them through collaboration.

Principle 8: Responsibility for Decisions Made. The potential for shared responsibility when using ML or AI must be identified, recognized by the relevant party and recorded upfront when researching or implementing ML or AI.

Principle 9: Governance. A hospital or practice using or developing ML or AI for patient care applications must have accountable governance to oversee implementation and monitoring of performance and use, to ensure practice is compliant with ethical principles, standards and legal requirements.” (57)

While noting the benefits of this framework, it is acknowledged that ethical principles, without a framework for accountability, education and training, opportunities for deeper process alignment, and a continuous improvement process are of limited value. To truly harness the power of AI, there must be a continuous focus on and respect for a patient's autonomy/beneficence, a commitment to non-maleficence, and an equally strong commitment to equity. (58)

The concerns about patient health and safety driving the approaches in more cautious jurisdictions appear to be, in large part, driven by a need to address those issues that most urgently impact patient experiences and satisfaction with AI-related experiences: the potential for

bias, opacity and incontestability, and erosion of privacy. (59) AI systems must be designed, from the start, with attention, to and an understanding of, bias, how it currently permeates our healthcare systems (thereby affecting our potential data sets and algorithmic inputs), and the continued need to monitor for and root out bias whenever and wherever it is found in AI systems.

(60)

Patchwork of Laws, Guidance, and Initiatives – United States of America

The United States currently has a patchwork of laws, guidance and executive orders from the White House, guidance from various federal agencies, and a mix of state laws that attempt to regulate AI use cases. Like Canada and other countries, many existing laws related to data privacy and consumer protection apply to AI. However, the U.S. has not yet created a comprehensive framework of laws and regulations on AI. A more comprehensive regulatory framework may be created in the near future if Congress passes such legislation. However, any such proposed legislation has not yet picked up steam. While not an exhaustive review of all U.S. laws and guidance on AI, following is a select sampling of such information to provide an overview of the existing U.S. AI healthcare regulatory framework.

The U.S. Congress passed the National Artificial Intelligence Initiative Act (NAIIA) of 2020 on January 1, 2021. (61) This law provides for a coordinated program across the federal government to accelerate AI research and application for economic prosperity and national security. Following passage of the NAIIA, the National Artificial Intelligence Initiative was created with the main purposes of ensuring continued U.S. leadership in AI research and development; leading the world in the development and use of trustworthy AI systems in public and private sectors; preparing the present and future U.S. workforce for the integration of artificial intelligence systems across all sectors of the economy and society; and coordinating ongoing AI activities across all federal agencies. (62) Under implementation of this initiative, a new office was created under the White House Office of Science and Technology Policy (OSTP) that directs the president in reaching the aforementioned goals and supports the NAIIA. This law and initiative, while significant, focus on coordination and strategy at the federal agency level primarily in relation to national security and the economy.

While the NAIIA provides direction and coordination across the federal government on AI, the U.S. does not currently have a framework that consists of a set of laws and regulations that govern the development and use of AI nor one that addresses AI in healthcare broadly. Medical devices,

including AI/ML-enabled medical devices, are regulated by the Department of Health and Human Services (HHS) U.S. Food and Drug Administration (FDA) in accordance with the Federal Food, Drug, and Cosmetic Act. ⁽⁶³⁾ However, the FDA's traditional regulatory paradigm for medical device regulation is not well-suited for adaptive AI and ML technologies, ⁽⁶⁴⁾ and accordingly, the FDA has issued a discussion paper on a proposed regulatory framework ⁽⁶⁵⁾ and a number of other related guidance on the topic, ⁽⁶⁶⁾ seeking stakeholder input. A more comprehensive AI/ML in medical devices regulatory framework may be forthcoming.

The HHS Office of the National Coordinator for Health Information Technology (ONC) Health IT Certification Program (Certification Program) is a voluntary certification program established by the ONC to provide for the certification of health information technology. ⁽⁶⁷⁾ Requirements for certification are established by standards, implementation specifications, and certification criteria adopted by the Department of Health and Human Services (HHS). The ONC recently issued a proposed rule, published on April 18, 2023, on health data, technology, and interoperability. ⁽⁶⁸⁾ The preamble recognizes that “the U.S. healthcare industry does not have universally applicable, consistently applied framework(s), best practices, or norms for transparency about technical and performance aspects and organizational competencies (e.g., model risk management) in place for [decision support interventions].” ⁽⁶⁹⁾ In the proposed rule, ONC proposes to rename the existing “clinical decision support” (CDS) certification criterion to “decisions support interventions” (DSIs) and introduce transparency requirements under this criterion. The proposal introduces “information transparency to address uncertainty regarding the quality of predictive DSIs that certificated Health IT Modules enable or interface with, so that potential users have sufficient information about how a predictive DSI was designed, developed, trained, and evaluated to determine whether it is trustworthy.” ⁽⁷⁰⁾ ONC also proposed requirements that would enable users to know when a DSI uses demographic, social determinants of health assessment data. While the voluntary certification requirements apply only to health information technology, and not all AI/ML-enabled healthcare products, health information technology is nonetheless widely used. The proposed DSI requirements may thus have broad impact, albeit on a subset of AI/ML technology in healthcare, if this proposal is finalized.

Furthermore, on the regulatory front, the HHS Office for Civil Rights (OCR) recently published a notice of proposed rulemaking (NPRM) to revise its regulations on nondiscrimination in health programs and activities. ⁽⁷¹⁾ OCR issued this proposed rule regarding Section 1557 of the Affordable Care Act (ACA), which prohibits discrimination on the basis of race, color, national

origin, sex, age, or disability in certain health programs and activities. If finalized, the proposal would make explicit that covered entities are prohibited from discriminating through the use of clinical algorithms on the bases prohibited by Section 1557. OCR sought comment on whether to limit this provision to clinical algorithms or to include other forms of automated or augmented decision-making tools or models, such as AI/ML. OCR is expected to respond to public comments and determine in future rulemaking whether and how to expand the nondiscrimination protections to AI/ML decision-making tools.

In October 2022, the White House released the *Blueprint for an AI Bill of Rights—Making Automated Systems Work for the American People*. ⁷² This publication was a signal to the industry and to Congress that additional consumer protections and safeguards against the harms of AI are needed now. It sets forth a framework for consumer protections and considerations, described in more detail below. The White House has provided other guidance and directives as well, including release of an Executive Order (EO) in February 2023 entitled “Further Advancing Racial Equity and Support for Underserved Communities Through the Federal Government.”

⁷³ This EO directs federal agencies to “promote equity in science and root out bias in the design and use of new technologies, such as artificial intelligence.” It further states that “[w]hen designing, developing, acquiring, and using artificial intelligence and automated systems in the Federal Government, agencies shall do so, consistent with applicable law, in a manner that advances equity.” Through these actions, the administration is guiding the federal government on how to responsibly use AI.

The AI Bill of Rights applies to AI across sectors, providing a national values statement and toolkit to help build protections into technological design processes and to inform policy decisions. The *Blueprint for an AI Bill of Rights* outlines five principles to govern automated systems:

- 1 Safe and Effective Systems: Individuals and communities should be protected from unsafe or ineffective systems; such systems should be developed with consultation from diverse stakeholders and experts and should undergo pre-deployment testing, risk identification and mitigation, and ongoing monitoring. Individuals should be protected from inappropriate or irrelevant data use in the design, development, and deployment of automated systems as well as from the compounded harm of its reuse.
- 2 Algorithmic Discrimination Protections: Algorithms and systems should be designed in an equitable way and should not disfavor individuals based on classifications protected by law

(e.g., race, color, ethnicity, sex, religion, age, national origin, disability, veteran status, or genetic information). AI system developers should use proactive and continuous measures to guard against algorithmic discrimination, including equity assessments and algorithmic impact assessments featuring both independent evaluation and plain language reporting. Healthcare clinical algorithms that are used by physicians to guide clinical decisions may include sociodemographic variables that adjust or “correct” the algorithm’s output on the basis of a patient’s race or ethnicity, which can otherwise lead to race-based health inequities.

- 3 Data Privacy: Individuals should be protected from abusive data practices via built-in protections, including ensuring that data collection conforms to reasonable expectations and that only data strictly necessary for the specific context is collected. Automated systems developers are encouraged to seek consent before using personal data. Consent should only be used to justify data collection in cases where it can be “appropriately and meaningfully given.” If it is not possible to obtain consent in advance, developers are encouraged to implement privacy by design safeguards. Data in sensitive domains, including healthcare-related data, should be subject to enhanced protections and restrictions.
- 4 Notice and Explanation: AI system developers should provide timely and accessible descriptions in plain language to describe overall system functioning and the role automation plays, notice that automated systems are in use, the individual or organization responsible for the AI system, and explanations of outcomes. Automated systems should provide explanations that are technically valid, meaningful, and useful to operators of the system.
- 5 Human Alternatives, Consideration, and Fallback: Individuals should be able to opt out from automated systems in favor of human alternatives, where appropriate or required by law. Appropriateness should be determined based on reasonable expectations in a given context in addition to ensuring broad accessibility and protecting the public from especially harmful impacts. Automated systems with an intended use within sensitive domains should be tailored to the purpose, provide meaningful access for oversight, include training for any people interacting with the system, and incorporate human consideration for adverse or high-risk decisions.

This Blueprint provides a framework for future legislation and regulations on AI and the safeguards that should be in place to protect consumers’ privacy, freedom of choice, and protect consumers against biases. The Blueprint rightly calls for heightened consumer protections in circumstances where AI is used in healthcare.

While regulatory gaps remain at the federal level, state and local law is beginning to step in to fill those gaps. For example, New York City passed Local Law 144, which would require employers and employment agencies to conduct a bias audit on any automated employment decision tools they intend to use. ⁷⁴ A California assembly member recently introduced a bill to combat algorithmic discrimination by automated tools that make consequential decisions. ⁷⁵ A number of other states have recently proposed similar legislation. ⁷⁶ Unless and until the U.S. Congress enacts comprehensive legislation, we may continue to see states take steps to fill in gaps in AI laws and regulations.

In contrast to the United States, the European Union is well on its way to passage of comprehensive AI regulatory regime.

Emerging Leadership – European Union

The European Union (EU) is on the verge of creating a comprehensive, far-reaching regulatory regime for AI through approval of the proposed Artificial Intelligence Act (AI Act), a proposed law over two years in the making. On June 14, 2023, the EU took yet one more step toward passage of the sweeping legislation: the European Parliament, a main legislative branch of the EU, passed a draft of the AI Act. As with comprehensive data privacy protections the EU passed under the General Data Protection Regulation (GDPR), the AI Act is thorough and is a proactive and unified effort by the member states of the EU to shape the industry and create corporate accountability. The AI Act is expected to pass in late 2023 and as the first of its kind, may set the standard on AI regulation on a global scale. The earliest the law would likely apply is the second half of 2024.

⁷⁷

The law assigns applications of AI to four risk categories: unacceptable risk, high risk, limited risk, and minimal risk. ⁷⁸ AI systems considered to be a clear threat to safety and the livelihood and rights of people will be deemed to have unacceptable risk and will be banned. An example of such AI is social scoring by governments (i.e., classifying individuals based on behavior, socioeconomic status, or personal characteristics). ⁷⁹ High-risk AI includes critical infrastructures that could put the life and health of citizens at risk, education training that may determine access to education, and safety components of products, such as robot-assisted surgery. High-risk AI systems will be subject to strict obligations before they can go to market.

Limited-risk AI refers to systems with specific transparency obligations that would allow users to make informed decisions. ⁽⁸⁰⁾ For such technology, users must be informed that they are interacting with a machine to be able to make an informed decision on whether or not to use the system. Lastly, minimal- or no-risk AI allows the free and largely unregulated use of AI. Examples of minimal- or no-risk AI include AI-enabled video games or spam filters.

The four risk tiers are each subject to different constraints and requirements. Developers can generally satisfy the requirements by complying with the technical standards that are currently being formulated by European standards-setting bodies.

The proposed AI Act “focuses primarily on strengthening rules around data quality, transparency, human oversight, and accountability. It also aims to address ethical questions and implementation challenges in various sectors ranging from healthcare and education to finance and energy.” ⁽⁸¹⁾ The scope of the AI Act is expansive and extraterritorial: it applies to providers and users of AI outside of the EU when the system output is used in the EU.

Perhaps most importantly, the AI Act has teeth. As it currently stands, it contains strikingly high fines—the greater of up to €40 million or 7% of the company’s total worldwide annual turnover for the preceding financial year. This large scope and penalty system will shape behavior outside of the EU and impact companies worldwide. Its strong penalties may curb or slow development in AI. However, given the potential for AI’s exponential growth, a slow start may be prudent, with adjustments along the way.

Patchwork Approaches – United Kingdom, China

With the 2020 withdrawal of the United Kingdom (UK) from the EU, the UK is working on its own regulatory approach to AI. On March 29, 2023, the UK government’s Department for Science, Innovation, and Technology and Office for Artificial Intelligence released a white paper detailing its plan for implementing a pro-innovation approach to AI regulation and seeking input through consultation. ⁽⁸²⁾ It states that it seeks to be a leader in this area. The approach is underpinned by five principles intended to guide how regulators approach risk:

- 1 Safety, security, and robustness
- 2 Appropriate transparency and explainability

3 Fairness

4 Accountability and governance

5 Contestability and redress

The white paper notes that currently, AI technologies in the UK are regulated by a “complex patchwork of legal requirements.” The creation of an approach to AI regulation was prompted in part by a concern that the absence of cross-cutting AI regulation creates uncertainty and inconsistency, which can undermine business and consumer confidence in AI, stifling innovation.

The existing patchwork of laws in the UK that provide some coverage of AI issues include the Equality Act of 2010, which provides consumer protections against discrimination. Medical device laws similar to those in the U.S. exist in the UK and regulate some products that include integrated AI. Consumer rights laws may offer protection to consumers where they have entered into sales contracts for AI-based products and services.

The framework sets out to engage industry, the public sector, regulators, and other stakeholders. Among other things, the government will work to design and publish an AI Regulation Roadmap with plans for establishing the central functions, including monitoring and coordinating implementation of the principles. The UK approach comes in stark contrast to the EU approach: in the UK’s press release for its white paper, the government makes clear that it “will avoid heavy-handed legislation which could stifle innovation and take an adaptable approach to regulating AI. Instead of giving responsibility for AI governance to a new single regulator, the government will empower existing regulators [...] to come up with tailored, context-specific approaches that suit the way AI is actually being used in their sectors.” ⁸³ The UK’s regulatory approach to AI, while lighter than that of the EU, may ultimately be subsumed by the EU’s AI Act given the law’s potential reach.

Asia does not have a singular, unified approach to regulating AI. However, China is taking an active role to regulate specific types of AI algorithms and capabilities and is doing so on a rolling basis. It is one of the first countries in the world to do so. In contrast to the EU’s horizontal approach to AI regulation that uses a single piece of legislation to regulate across the industry, China’s approach, led by the Cyberspace Administration of China, is vertical, applying to common or risky use cases. The first set of regulations target generative AI, focusing on algorithms that make recommendations and use deep synthesis technology. ⁸⁴

Finalized on July 13, 2023, as “interim” measures, and set to go into effect on August 15, 2023, ⁽⁸⁵⁾ China’s AI regulations were altered from the draft to soften their impact and demonstrate support for innovation by, in part:

- only requiring those who are developing public facing products to submit security assessments (i.e., companies working on enterprise/internal facing products would not have the same hurdles);
- removing language that required a “three-month waiting period for ‘improving model training and other methods to prevent recurrence’ of content that violates the guidelines”;
- removing draft fines of up to 100,000 yuan (approximately \$14,000.00 USD); and
- providing exemptions for companies in China that want to provide generative AI products to markets outside of China, while ensuring that foreign companies wishing to provide generative AI products in China are subject to the regulations.

Regulations that remained intact from the draft version released in April include:

- in processes such as algorithm design, selecting training data, and model generation and model optimization, measures are required to be in place to prevent discrimination on the basis of race, ethnicity, religion, and nationality;
- content generated through the use of generative AI is required to be true, accurate, and measures must be adopted to prevent generation of false information;
- consent is required for use of personal data for generation of AI product pre-training and optimization training; ⁽⁸⁶⁾ and
- developers must register their algorithms, allowing regulators to review the algorithms and information such as the training data used and security risks.

While China has mandated that generative AI products must adhere to “core socialist values,” the final regulations also clearly reflect the government’s goal to help Chinese companies gain an advantage in the global technological AI race.

Global Perspective

AI regulation across the globe is in its infancy, as shown in the above overview. AI technology is also in its infancy but, by its very nature, is positioned to take off exponentially at a rate that is likely to quickly outpace the development of laws and regulations. The US, China, Australia, and EU, among other governmental bodies, have positioned themselves to lead this space, both technologically and in the development of a regulatory framework. The development and evolution of these frameworks will have significant consequences on both technological innovation and on consumer populations across the world.

The European Union and China have moved quickly to develop contemporary laws and regulations in this space. The EU's proposed regulatory framework is likely to have the strongest impact regionally and globally, given the comprehensiveness, scope, and scale of its proposed law. Countries in less developed regions appear to be taking a wait-and-see approach, allowing the developed countries to move forward first in order to assess the effectiveness, and implications, of their regulatory actions. Whether AI is developed in a particular country or not, all countries should consider implementing regulations that will protect their citizens against the harms of AI used and sold within its borders. Louder, more frequent calls for global regulation have begun, and global leaders are taking notice. 87

The Intersection of AI and Health Equity

Of the various AI frameworks, principles, and proposed regulations examined, what then are the mechanisms likely to bolster political will, ensure accountability, and create space for marginalized voices in the healthcare arena? What are the best ways to engage people who are rightfully concerned about bias, transparency, and privacy? What actual protections can be relied upon in the face of permissive “blueprint” documents that have no enforcement mechanisms? Undoubtedly, there will be a plethora of new regulations, guidelines, principles, and frameworks to come as AI evolves. The extent to which they advance health equity will depend upon the extent to which they regulate AI and address known inequities in healthcare. This section will offer suggestions through the gateway of data to discuss concerns identified in this article.

Representative Big Data – The Gateway to Equitable AI

While there are many pieces of the AI puzzle that will require insight, innovation, regulation, and continuous improvement, perhaps none is more important at the outset than data. “Big data” is a term that describes “large amounts of data that is unmanageable using traditional software or

internet-based platforms..., which surpasses the traditionally used amount of storage, processing, and analytical power.” (88) The term involves data that has a high volume, is generated with great velocity, and contains many varieties—attributes that all apply to healthcare information today, especially in HIC. (89) To produce more accurate, unbiased, and representative information, AI/ML tools must be trained with high-quality, representative data collected from across all demographics. *The quality of the output is limited by the quality of the input.* How can stakeholders support data interoperability (the ability of different information systems, devices, and applications to access, exchange, integrate, and cooperatively use data in a coordinated manner) (90) not only within a country, but across countries, while at the same time supporting new methods of data collection and engagement that allow for the collection of data across all populations of people? Particularly in LMIC, there is a need not only for data, but also for the technological infrastructure to support its collection, storage, management, and safety.

One novel example of a way to collect and share healthcare data is the NIH's All of Us program, which is designed to amass diverse healthcare data for use in research, in collaboration with public and private partners. (91) Core values include:

- “Participation is open to all. Enrollment is open to all eligible adults who live in the United States. People of every race, ethnicity, sex, gender, and sexual orientation are welcome. No health insurance is required. You can be healthy or have health issues. You can sign up directly through JoinAllofUs.org or through participating healthcare provider organizations. In the future, children will be able to join.
- Participants reflect the rich diversity of the United States. To develop individualized plans for disease prevention and treatment, researchers need more data about the differences that make each of us unique. Having a diverse group of participants can lead to important breakthroughs. These discoveries may help make healthcare better for everyone.
- Participants are partners. Participants shape the program with their input and contribute to a project that may improve the health of future generations. They may also learn about their own health.
- Transparency earns trust. We inform participants about how their data are used, accessed, and shared. Participants can choose how much information to share.

- Participants have access to their information. All of Us lets participants see their own information and records.
- Data are broadly accessible for research purposes. All of Us makes information about participants as a group available in a public database. Everyone can explore the database or use it to make discoveries. Data from individual participants are also available, but only for researchers who apply and are approved. Any personal information that identifies a participant, such as name or address, is removed from data that researchers can access.
- Security and privacy are of highest importance. Data are stored in a secure, cloud-based database. All systems meet the requirements of the Federal Information Security Management Act. Ongoing security tests help protect participant data. Learn more about how the All of Us Research Program protects data and privacy.
- The program will be a catalyst for positive change in research. Working together, All of Us researchers, partners, and participants can build a better future for health research and care.”

92

All of Us is designed to collect data that can positively affect the SDOH by integrating biological data with environmental and lifestyle data to provide researchers a more meaningful and appropriate dataset upon which to build not only precision medicine solutions that treat cancer, for example, but many other diseases as well. Further, it is hoped that the program can support research insights into healthier living generally, without reference to disease or treatment. 93

All of Us is nearly halfway to its goal of a million participants, with over 409,000 participants included as of February 2023. 94

How can individuals be incentivized to join this or similar programs in order to ensure the data is representative, and what benefits will participants see from their participation? How can we support program participation by those who have fewer resources? How can resources be marshalled to support the collective development of these types of initiatives globally?

Strengthening Political Will

Political will, “the process of generating resources to carry out policies and programs... based on public understanding and support,” 95 will be key in driving the development of, and access to, the data that will drive AI innovations in healthcare. In this sense, forming a comprehensive,

insightful, and impactful AI approach is best understood, at the outset, as an educational exercise that must create understanding and empower all stakeholders in the system to play their role in a way that supports health equity. Without the proper foundational understanding of what is to be regulated, the potential risk and rewards, and an appreciation for the many ethical dilemmas that will arise, there will be little political support to shepherd stakeholders through the necessary process to build an appropriate regulatory framework. Building political will starts with the education of all: policymakers, legislatures, lobbyists, and government; researchers and students; all communities, but especially marginalized communities; influencers and champions for accountability; regulators; service providers; healthcare systems; and other support organizations that stand to benefit from the use of AI in healthcare. What is clear is that political will is intentionally developed over time and is not the province of any particular stakeholder—it is a journey that all must take together if there are to be beneficial results.

How can the general population be educated about AI and its effect on healthcare? What can major stakeholders do to expand opportunities for programmers, health monitors, researchers, data scientists, and algorithmic developers who will be needed in droves, to ignite the political will for equitable AI in healthcare the world over? Broad support for programs such as All of Us, that reach populations across all demographics and regions, will hopefully support new research breakthroughs, strengthening political will by demonstrating the benefit of a more inclusive and data-driven approach to developing healthcare research data and creating more equitable health outcomes.

Accountability and Transparency

Accountability must be examined anew to create a comprehensive approach to review, assess, and implement change management with respect to AI use in the healthcare space. Regulation cannot, as is often the case, lag far behind the pace of technological advancement, leaving the public, especially those who are traditionally marginalized, left behind and at risk. As the backbone of accountability is transparency, developing best practices and standards around transparency will be crucial. Although there is sure to be variation from jurisdiction to jurisdiction, transparency measures/metrics over time—especially those that can be standardized—will drive a more cohesive understanding of the technology itself, which will support accountability and political will among stakeholders.

Elevation of Marginalized Voices

At the most basic level, investment in AI is truly about investment in people. All people, and especially historically marginalized groups, need to contribute to and benefit from the use of AI in healthcare. To promote the equalization of playing fields, a portion of resources in LMIC, perhaps from investment from HICs and non-profits, should be directed toward the education system, allowing for the youth to become trained in programming and engineering, allowing these countries to be active participants in AI's development, uses, and thus benefits. This includes support for the education of internal and external data monitors, scientists, and researchers, as well as government investment in the education of developers, programmers, and coders in LMIC. Developers across countries, and particularly in HICs, should be educated on responsible and ethical design of AI and implications of biases and unrepresentative data.

Who Owns Data?

Consideration of two very different paradigms on the use of data as a society could offer ideas on how to regulate the fast-growing world of AI/ML in healthcare. Pursuing AI in healthcare as a common good, from inception to use and beyond, could provide wide access to developers and others with the least amount of government intervention or regulation. ⁽⁹⁶⁾ Consideration of data as a public good that can be utilized with broad public consent may be a reasonable approach to data protection, depending on how governments define the “public good,” whether the definition changes over time, how the public is educated, and how well data anonymization can deflate privacy concerns. ⁽⁹⁷⁾ Conversely, healthcare data could be viewed as an individual asset that belongs to the individual from whom the data was derived. Under such a paradigm, individuals would own, control, and potentially monetize their own data, readjusting incentive structures and dynamics. While individual ownership and control could more equitably distribute economic power and control to individuals as it relates to personal data, that model should be premised upon individual knowledge, access, and accountability to support truly equitable outcomes.

Certainly, there are pitfalls with either approach: a common public good is dependent on trust and corporate responsibility as individuals do not play a direct role in how their data is shared under such a model; meanwhile, an individualized approach is far removed from our current system, and the political feasibility of implementing such an approach may be low as a result. Both models have positive attributes from which future regulations can be drawn, but any decisions should be contextualized and implemented within the larger context of our existing norms and systems in order to be sound. No matter the scheme, privacy-by-design systems should be

employed in all healthcare applications where patient data is involved to provide the greatest level of trust at the outset.

Conclusion

Ready or not, AI has taken off exponentially, and perhaps may ultimately surpass human intelligence. Global AI regulation, which is in its infancy—must develop at an unprecedented pace, with global collaboration and alignment, to appropriately grapple with numerous ethical, legal, business, and policy implications of AI's promise and potential peril. To support the advancement of health equity to its fullest potential, all stakeholders must work together; deidentified health data must be readily available for analysis by the public and private sectors; such data must be representative of all populations, with particular attention and efforts made to gather data on marginalized populations; and such data must be transparently managed and protected, with the fruits of the data's analysis distributed and accessible to all. AI necessitates collective collaboration and perhaps, a reorientation of healthcare and education as common goods, to fulfill our collective highest potential and avoid widening existing economic and health disparities across the globe.

Acknowledgement: The authors would like to acknowledge the research assistance of St. Jude Children's Research Hospital Spring 2023 interns Callaghan Basil (The University of Mississippi School of Law, JD, May 2023) and Elizabeth Hamman (Mississippi College School of Law, JD expected May 2024).

Endnotes



Authors



Roma Sharma

Crowell & Moring, Washington, DC

Roma Sharma is counsel in Crowell & Moring's Washington, DC, office and is a director in Crowell Health Solutions, a strategic consulting firm affiliated with Crowell & Moring. Her practice entails counseling healthcare entities, including hospitals, group practices, technology companies, accountable care organizations, and managed care organizations in various regulatory, transactional, and investigational matters. Ms. Sharma counsels clients on fraud and abuse compliance, certificate of need, reimbursement, practice of medicine, telemedicine, licensing, and federal and state demonstration

programs. She also counsels clients on the ethical guardrails of AI. She can be reached at RSharma@crowell.com.

Tienne Anderson

St. Jude Children's Research Hospital, Inc., Memphis, TN

Tienne Anderson is managing counsel at St. Jude Children's Research Hospital in Memphis, TN. She has practiced in the healthcare space for nearly a decade, as in-house counsel with both for- and non-profit organizations. As part of St. Jude's Department of Global Pediatric Medicine, Ms. Anderson supports its mission to improve survival rates of children with cancer and other catastrophic diseases around the world. In addition, she co-chairs the Office of Legal Service's International Affairs & Compliance Practice Group and is intrigued by all questions at the intersection of philosophy and the law. She can be reached at Tienne.Anderson@stjude.org.

ENTITY:**HEALTH LAW SECTION****TOPIC:****HEALTH**

The material in all ABA publications is copyrighted and may be reprinted by permission only. Request reprint permission [here](#).

ABA American Bar Association |

/content/aba-cms-dotorg/en/groups/health_law/publications/health_lawyer_home/august-2023/the-promise-and-peril-of-advancing-health-equity-through-ai