Legal Concerns in Health-Related Artificial Intelligence:

A Scoping Review Protocol

Michael Da Silva, Tanya Horsley, Devin Singh, Emily Da Silva, Valentina Ly, Bryan Thomas, Ryan C. Daniel, Karni A. Chagal-Feferkorn, Samantha Iantomasi, Kelli White, Arianne Kent, Colleen M. Flood

[Note: This is an author's accepted manuscript for a forthcoming piece in Systematic Reviews.]

Abstract

Background: Medical innovations offer tremendous hope. Yet similar innovations in governance (law, policy, ethics) are likely necessary if society is to realize medical innovations' fruits and avoid their pitfalls. As innovations in Artificial Intelligence (AI) advance at a rapid pace, scholars across multiple disciplines are articulating concerns in health-related AI that likely require legal responses to ensure the requisite balance. These scholarly perspectives may provide critical insights into the most pressing challenges that will help shape and advance future regulatory reforms. Yet, to the best of our knowledge, there is no comprehensive summary of the literature examining legal concerns in relation to health-related AI exists. We thus aim to summarize and map the literature examining legal concerns in health-related AI using a scoping review approach.

Methods: The scoping review framework developed by Arksey and O'Malley (2005) and extended by Levac et al. (2010) and the Preferred Reporting Items for Systematic Reviews and Meta-Analysis extension for scoping reviews (PRISMA-ScR) guided our protocol development. In close consultation with trained librarians, we will develop a highly sensitive search for MEDLINE (OVID) and adapt it for multiple databases designed to comprehensively capture texts in law, medicine, nursing, pharmacy, other healthcare professions (e.g., dentistry, nutrition), public health, computer science, and engineering. English- and French-language records will be included if they examine health-related AI, describe or prioritize a legal concern in health-related AI or propose a solution thereto, and were published in 2012 or later. Eligibility assessment will be conducted independently and in duplicate at all review stages. Coded data will be analyzed along themes and stratified across discipline-specific literatures.

Discussion: This first-of-its-kind scoping review will summarize available literature examining, documenting, or prioritizing legal concerns in health-related AI to advance law and policy reform(s). The review may also reveal discipline-specific concerns, priorities, and proposed solutions to the concerns. It will thereby identify priority areas that should be the focus of future reforms and regulatory options available to stakeholders in reform processes.

Scoping Review Registration: This protocol was submitted to the Open Science Foundation registration database. See https://osf.io/zav7w.

Key Words: Health Law, Artificial Intelligence, Machine Learning, Health, Scoping Review

Background

Innovation in medicine offers tremendous hope but appropriate regulations are needed to realize this hope in healthcare and public health settings while minimizing attendant risks. The advent of health-related artificial intelligence (AI) provides an example of this need for regulation that fosters the requisite balance of potential benefits and pitfalls. Health-related AI is the topic of significant debate across multiple fields. Many portend that AI will improve healthcare systems by, e.g., increasing the accuracy of diagnoses, improving the efficiency of healthcare delivery, or mitigating human biases (e.g., Topol 2019ab). At the same time, there are understandable concerns that AI will undermine the patient-provider relationship, contribute to the deskilling of providers, undermine transparency, misdiagnose or inappropriately treat because of errors within AI decision-making that are hard to detect, exacerbate existing racial or societal biases, or introduce algorithmic bias that will be hard to detect (e.g., Gerke et al. 2020). Opinions are, of course, not starkly polarized. Many simultaneously recognize AI's promise and maintain concerns about its widespread adoption. Yet these considerations highlight the need for regulations or other forms of governance (law, policy, ethics) that help countries leverage AI's potential benefits while minimizing any attendant risks.

The use of health-related AI will intersect with the law in several ways. First, there are questions about whether existing laws will address identified concerns with health-related AI, such as:

i. whether medical device regulation, medical malpractice laws, product liability laws and professional self-regulation and accreditation will adequately attend to the possibility of error on the part of AI;

ii. whether existing rules concerning the attribution of liability for medical error are appropriate for when AI tool recommends – or even performs – a course of harmful treatment and how liability should be attributed as between healthcare professionals and AI developers and manufacturers;

iii. whether existing anti-discrimination and human rights laws can attend to the problem of algorithmic bias in which AI tools inappropriately produce different outcomes for historically disadvantaged groups;

iv. whether existing privacy laws sufficiently protect patients given AI's big data needs and the fact that Machine Learning (ML) tools, for example, will collect data in real time;

v. whether existing laws and policies relating to data governance are sufficient to enable AI innovators to have access to representative training data sets so as to appropriately include historically underrepresented populations; and

vi. whether existing laws of informed consent are sufficiently robust to protect patients when clinicians choose to use AI in diagnosis and treatment.

Satisfactory responses are pressing as AI use increases in healthcare and public health settings. AI tools can already read medical images and reports more quickly and accurately than human counterparts (Topol 2019a; 2019b). The U.S. Food and Drug Administration has approved AI tools for diagnosing strokes and brain bleeds, detecting atrial fibrillation, and interpreting brain MRIs; hospitals could also use other existing AI tools to assess patients' risk of readmission and respond accordingly (Topol 2019b). Health-related AI is also used in Canada, including for triage purposes and to respond to COVID-19 (CIFAR 2021a). The development of ML as a subset of AI that uses large data sets to "make predictions and solve problems … without being explicitly programmed" (Murphy et al. 2021) has produced increasingly accurate health-related applications that can 'learn' from real-world data over time and improve healthcare systems. As these technological innovations continue, there is a clear need to examine whether existing laws are up to the task of ensuring beneficial health-related AI tools can be deployed in real-world settings while minimizing legitimate concerns about, e.g., bias or privacy, that may arise.

Depending on one's conclusions regarding the sufficiency of existing regulations to address concerns associated with the use of health-related AI, the next question is what kind of regulatory reform is required. Overly onerous regulations could stifle innovation, minimizing AI's potential to realize health-related benefits. At the same time, a lack of regulation or inadequate regulations may lead to widespread harm resulting in a loss of trust in health-related AI (from patients themselves and healthcare professionals) and, perhaps, much more arduous regulation.

Despite the importance of law for the successful implementation of AI in health settings, there is no systematic overview of the legal concerns raised by the use of health-related AI. A preliminary search of MEDLINE, Cochrane Library, PROSPERO, and JBI Evidence Synthesis was conducted; no current or underway systematic reviews or scoping reviews on this topic were identified. Two

published scoping reviews were identified that survey ethical concerns raised by the use of health-related AI (Lillywhite & Wolbring 2019; Murphy et al. 2021) with one of the two taking a narrower focus on ethical issues concerning the disabled (Lillywhite & Wolbring 2019). Although ethical issues can overlap with legal issues in some cases, legal concerns and legal responses are important to understand in their own right. For example, law has an important claim to *democratic legitimacy* and failing to include legal issues means a failure to include concepts, such as patient autonomy, as articulated in sources of public and private law (e.g., Ontario's *Personal Health Information Protection Act*). To illustrate with this one example, some may think that a consideration of ethical issues will capture concerns relating to consent to treatment, but a legal requirement for informed consent (as in Ontario's *Health Care Consent Act*) can differ in its nature and impact. Further, although lawyers disagree on what the law should be, there is less room for disagreement than with respect to many ethical issues. Precedent (higher courts binding lower courts) helps create consistency in interpretation notwithstanding personal views on what the law should be.

A range of reports from different jurisdictions have identified areas where law reform may be needed to respond to problems in health-related AI (Academy of Medical Royal Colleges 2019; AI Forum of New Zealand 2019; CIFAR 2020; RCPSC 2020; CMA 2020; CIFAR 2021a; CIFAR 2021b; WHO 2021). Scholars have produced general overviews of the use of AI in healthcare that make claims with legal implications. For instance, Eric Topol (2019a; 2019b)'s well-known studies discuss concerns about the risk of AI tools causing harm to patients that implicate questions about AI safety and efficacy regulation and about both contractual and tort liability for harm. Other scholars have produced ethical analyses of health-related AI with possible legal implications. For instance, Alessandro Blasimme and Effy Vayena (2020) raise concerns regarding AI safety, bias, and informed consent that may require legal responses of some kind. Mark Henderson Arnold (2021) raises similar concerns, as well as questions about what to do when providers become too reliant on AI, which raise further liability questions. A relatively small number of works have articulated views on the most pressing legal concerns and possible best practices for approaching them (e.g., Kerr et al. 2017; Gerke et al. 2020). Yet there is no systematic overview of the legal concerns. As discussed above and further below, law and ethics overlap in important ways but legal and ethical issues can and should be distinguished where possible. A scoping review of paradigmatic and self-identified legal issues that identifies which concerns experts from different disciplines view as most important will enable better analysis of the adequacy of laws in various jurisdictions and what reforms are required (if any).

<u>Objective</u>

We aim to systematically map legal concerns that are identified in health-related AI and the extent to which they are prioritized across multiple relevant disciplines, namely law, medicine, nursing, pharmacy, other healthcare professions (dentistry, nutrition, etc.), public health, computer science, and engineering. In keeping with the central purposes of a scoping review approach, we aim to examine the extent, range and nature of research activity across the disciplines, to summarize and disseminate research findings to relevant stakeholders, and to identify research gaps in the existing literature. The scoping review will be conducted in accordance with the framework developed by Arksey and O'Malley (2005) and extended by Levac et al. (2010) and aligned protocol reporting to the PRISMA-P checklist (Moher et al. 2015). It will thus include 6 stages: 1. Identifying the Research Question(s), 2. Identifying Relevant Studies, 3. Study Selection, 4. Charting the Data, 5. Collating, Summarizing, and Reporting Results and, in this case, 6. Stakeholder Consultation.

Stage One: Identifying the Research Question(s)

Through iterative discussions, research team members met to discuss and refine the research question(s). To ensure diverse perspectives and representation, the team included expert clinicians, AI innovators, legal researchers, ethics experts, and a member with scoping review expertise. From these discussions, we decided that our aim to map the nature, extent and range of literature examining legal concerns in health-related AI. To this end, our research question(s) are:

Primary Question:

1. What is known from the literature regarding legal concerns in health-related AI?

Secondary questions include

2. Are the legal concerns identified explicitly prioritized?
3. Do different disciplines identify, represent, or prioritize legal concerns differently?

Stage Two: Identifying Relevant Studies

Guided by two expert information specialists, we will develop a highly sensitive search strategy to identify relevant records. A preliminary pilot search of MEDLINE® and HeinOnline was conducted to pilot test the draft search and its ability to identify key articles. Titles, abstracts, keywords, and Medical Subject Headings (MeSH®) in MEDLINE® of the key articles were analyzed to develop the final search strategy. See Table 1 for details. The MEDLINE® search was peer reviewed by an independent information specialist using the Peer Review of Electronic Search Strategies (PRESS) checklist (McGown et al. 2016). This original search strategy will be adapted to search additional medical, legal, and multidisciplinary databases as applicable. Searching will be conducted in the following electronic databases: MEDLINE® (Ovid), EMBASE (Ovid), HeinOnline Law Journal Library, Index to Foreign Legal Periodicals (HeinOnline), Index to Legal Periodicals and Books (EBSCO), Web of Science, Scopus, and IEEE Xplore. Given the complexity of searching (e.g., its multidisciplinary), the strategies will be augmented by hand-searching reference lists of all relevant, full-text records to identify additional sources (Horsley et. al. 2011). Search results will be exported into a proprietary review software program (Covidence®) to facilitate review processes and manage each stage of the review (e.g., de-duplication, eligibility assessment).

Searches will be conducted between 2012 and the start of the review. This decision was made through consultation with all team members and was made to account for more recent developments in the field of AI and the deployment of AI into healthcare settings. For example, the development of deep learning AI is viewed as a "paradigm shift" in the field (Bohr and Memarzadeh 2020). While deep learning AI was developed prior to 2012, the use of deep learning in the ImageNet Large Scale Visual Recognition Challenge 2012-winning AlexNet (Krizhevky et al. 2012) led to the technology's increased public recognition. Experts accordingly view 2012 as the year deep learning was "widely accepted as a viable form of" AI (Topol 2019b). Health-related AI has become more common since that date (Topol 2019b).

Several members of the research team were part of a previous Royal College of Physicians and Surgeons of Canada Task Force Report on Artificial Intelligence and Emerging Digital

Technologies (TH, DS, CF) that contained a bibliography on the grey literature (Royal College 2020). The team will use that bibliography as a starting point for identifying relevant grey literature, supplementing it by crowdsourcing information about subsequent grey literature among experts working on the research project and further hand-searching, including review of references for mentions of other reports by governments (e.g., US GAO 2020), non-governmental organizations (e.g., CIFAR 2021ab), and professional organizations (e.g., Royal College 2020). Initial examinations of such reports suggest that they are written from an interdisciplinary perspective and thus inappropriate for our task of mapping concerns across disciplines. However, we will summarize their statements on legal concerns and compare them with those in published records.

Stage Three: Study Selection

Eligibility Criteria:

Our search strategy intends to identify published literature. It will be designed to identify indexed records, including articles and book chapters. It will also include American law journals, which are academic journals that undergo a standard student review process. Publications will not be limited to works in law journals as the reviewers are interested in what people working in other fields view as key legal concerns (or if they are even discussing legal concerns). Articles in medical, social scientific, or computer science or engineering journals that raise legal concerns or solutions thereto are included in this review. We will not include abstracts and conference proceedings.

This review is centrally focused on identifying records that describe, articulate, analyze or prioritize legal concerns associated with health-related AI. It is specifically focused on identified concerns in health-related AI that may require a legal response. This could include records identifying concerns that require a clear legal solution (e.g., a legal gap), evaluation of existing laws to meet the concerns, and examining the nature, advantages and disadvantages of various law reform proposals. We will include records addressing various legal concerns (e.g., an omnibus of concerns requiring legal responses from bias to liability) as well as those addressing a single concern (e.g., privacy). We will exclude records that simply describe existing law(s) that apply to the use of health-related AI. We recognize that 'law' and 'ethics' are sometimes conflated and overlap in some respects. Yet ethical concerns are and should be distinguished from legal concerns: not all laws reflect our best moral standards, and not all ethical concerns implicate the law. Further, ethical debates about health-related AI have been the subject of a scoping review (Murphy et al. 2021). To this end, records that identify an ethical concern will only be included when there is clearly an articulated nexus to a legal issue (for example, where there is an ethical concern related to privacy and then whether privacy laws are sufficient to protect patients and providers when using ML devices).

Records will be eligible for inclusion if they are English- or French-language publications that describe a legal concern pertaining to the use of AI that is health-related. Working definitions for the key terms ('legal concerns,' 'AI,' and 'health-related') are summarized succinctly in Table 2. For the initial searches, we will not impose a language restriction. However, we will include only records published in English or French in the scoping review itself.

Eligibility Assessment

An eligibility assessment tool (for level 1 title and abstract screening and level 2 full-text screening) will be developed and pilot tested in collaboration with the entire review team. Using a proprietary review software program to manage records, duplicate assessment of each record to determine inclusion will be conducted independently. Agreement will be assessed and reported using a Kappa statistic (e.g., Kundel and Polansky 2004). Conflicted decisions will be resolved by one of the subject matter expert authors (MD, CF, DS). Eligibility assessment and flow of the review will be reported in accordance with the PRISMA flow diagram.

Stage Four: Charting the Data

A standardized form to extract and categorize data will be developed through a series of consultations with the team and pilot tested using known includes. The form will cover two primary areas of information: (1) record-level demographic characteristics and (2) content specific to legal concerns, priority areas, solutions, etc. that can be inferred from the records.

Basic record-level information will include year of publication, disciplinary area of focus (e.g., law, medicine), geographic information (e.g., country, region), publication type, study design, language of publication, and area of health-related focus. For example, we will identify the medical specialty in which AI and the legal concern are being confronted. To categorize, we will design a list of specialties (using as a base the list employed by the Royal College of Physicians and Surgeons of Canada but modifying it to, e.g., add family medicine specializations). When this list is insufficient to capture the area of focus (e.g., general medicine, medical administration, electronic health record), we will extract the information verbatim as described in the original record. Given that this is a scoping review and is meant to be iterative and that information gathered can inform future stages, any additional decisions on data to be extracted that occur beyond the protocol stage will be declared transparently in the distal publication.

When extracting data pertaining to any legal concern(s), we will also extract any specific text that expressly indicates prioritization of the concerns and proposed solutions (e.g., new or modified regulations), including new interpretations or expansions of private law responsibilities (e.g., in tort law or contract law), ethical reform (e.g., on the part of AI innovators), and educational/training reforms. When looking at the prioritization of legal issues, we will rely on explicit self-identified priority claims (describing issues as 'high priority,' 'most important,' 'most pressing,' etc.).

Stage Five: Collating, Summarizing, and Reporting Results

Once data has been extracted and accuracy verified through duplicate assessment, we will summarize all record-level characteristics and use the summaries to map publication patterns (who is publishing on the intersection of health related-AI and law) and thematically code legal concerns into categories (e.g., typology). We will categorize the legal concerns using our framework, which seeks to summarize the legal concerns, priorities, and proposed solutions to the concerns.

Once categorized, we will examine the extent to which we see different legal issues of concern to different disciplines (law, medicine, nursing, pharmacy, other health professions, public health, computer science, and engineering). For example, we may see that privacy is raised as a concern by those in the fields of computer science or engineering whereas liability may be more frequently raised as a concern by legal researchers. We will also examine whether we see differences in what

legal concerns are identified by specialty, country/region, and whether or not any author on the record is identifiable as an expert from a relevant discipline other than that of the primary author. To achieve the latter, we will categorize authors using record-level information pertaining to their institutional affiliation (e.g., School of Engineering) or listed credentials. It is, of course, possible that individuals trained in one discipline work in another discipline (e.g., a lawyer could work in bioethics within a faculty of medicine). However, departmental affiliation is a useful marker of disciplinary perspective and what they represented as a co-author on the work. Any normative recommendations in the article stemming from data collected in this scoping review will be based on our multidisciplinary team's views on any inferences that can be drawn from the descriptive findings. In accordance with the PRISMA-Scr (Tricco et al. 2018), we will not formally assess methodological quality of records.

Stage Six: Stakeholder Consultation

Understanding the complexity of searching and breadth of the review scope, we plan to engage key consumers and stakeholders to suggest additional records and sources when an initial full-text include list is generated. Given that this review is one component of a larger Canadian Institutes of Health Research-funded project, we plan to solicit input from the broader multidisciplinary team, which includes patient partners and international experts on law, AI, and medicine (CHLPE 2021). If gaps in expertise or perspectives are identified, the team will employ a snowball sampling technique whereby an existing contact will be asked to provide the name(s) of other relevant experts and the process can be repeated as needed. The process of selection and consultation will be detailed sufficiently to ensure transparency.

<u>Limitations</u>

This scoping review is mapping records from across several disciplines and aims to identify legal concerns that are discussed in health-related AI. Given the complexity of searching, and despite all attempts to be comprehensive, it is conceivable that our search strategy may miss relevant records. This may be due to the records using idiosyncratic language or irregular key word entries, or other indexing-related barriers. It is also possible that relevant legal concerns could be in records that predate 2012. This said, our team anticipates that important debates will be captured. Finally, it is possible that few works will make explicit claims about the priority of legal issues however we will be able to impute the importance of particular legal questions by looking at the number of times that specific legal issues are raised.

<u>Discussion</u>

Health-related AI offers numerous potential benefits but also several potential pitfalls. Apt regulation is necessary to maximize its potential benefits while minimizing its risks. Regulatory reforms are more likely to maintain the requisite balance when informed by an understanding of how experts prioritize the concerns and potential solutions they have identified. This scoping review accordingly aims to systematically, comprehensively, and cohesively map the extant literature on legal concerns in health-related AI. A comprehensive map of how different disciplines frame, prioritize and analyze legal concerns in health-related AI and how they propose to address those concerns will support cross-disciplinary understanding of perceived legal barriers or gaps. It will also aid in policymaking through a holistic view of legal concerns in health-related AI and

multiple perspectives on solutions. Our review will thus provide the bedrock from which to better develop domestic and international responses to legal challenges posed by health-related AI.

References

Academy of Medical Royal Colleges. Artificial Intelligence in Healthcare. London: Academy of Medical Royal Colleges; 2019.

AI Forum of New Zealand. Artificial Intelligence for Health in New Zealand. Northshore: AI Forum of New Zealand; 2019.

Arksey H, O'Malley L. Scoping Studies: Towards a Methodological Framework. International Journal of Social Research Methodology. 2005;8:19-32.

Blasimme A, Vayena E. The Ethics of AI in Biomedical Research, Medicine and Public Health. In: Dubber MD et al., editors. The Oxford Handbook of Ethics of AI Oxford: Oxford UP; 2020.

Bohr A, Memarzadeh K. The Rise of Artificial Intelligence in Healthcare Applications. In: Bohr B, Memarzadeh K, editors. Artificial Intelligence in Healthcare. Cambridge: Academic Press; 2020.

Canadian Institute for Advanced Research (CIFAR). Building a Learning Health System for Canadians: Report of the Task Force on Artificial Intelligence for Health (AI4Health). Toronto: CIFAR; 2020.

~~~. AI and Healthcare: A Fusion of Law and Science – An Introduction to the Issues. Toronto: CIFAR; 2021a.

~~~. AI and Healthcare: A Fusion of Law and Science – The Regulation of Medical Devices with AI Toronto: CIFAR; 2021b.

Canadian Medical Association (CMA). The Future of Technology in Health and Health Care: A Primer. Ottawa: CMA; 2020.

Gerke S et al. Ethical and Legal Challenges of Artificially Intelligence-Driven Healthcare. In: Bohr B, Memarzadeh K, editors. Artificial Intelligence in Healthcare. Cambridge: Academic Press; 2020.

Hage J. Foundations. In: Hage J, Akkermans B, editors. Introduction to Law. New York: Springer; 2014a.

Hage J. Sources of Law. In: Hage J, Akkermans B, editors. Introduction to Law. New York: Springer; 2014a.

Henderson Arnold M. Teasing out Artificial Intelligence in Medicine: An Ethical Critique of Artificial Intelligence and Machine Learning in Medicine. Journal of Bioethical Inquiry. 2021;18:121-139.

Horsley, T et al. Checking Reference Lists to Find Additional Studies for Systematic Reviews. Cochrane Database of Systematic Reviews. 2011;10(8): MR000026.

Kerr I et al. Robots and Artificial Intelligence in Health Care. In Erdman J et al., editors. Canadian Health Law & Policy, 5th ed. Toronto: LexisNexis Canada; 2017.

Krizhevky A et al. ImageNet Classification with Deep Convolutional Neural Networks. NIPS'12: Proceedings of the 25th International Conference on Neural Information Processing. Systems. 2012;1:1097-1105.

Kundel HL, M Polansky. Measurement of Observer Agreement. Radiology. 2003;228(2):303-308.

Levac, D et al. Scoping Studies: Advancing the Methodology. Implementation Science. 2010;5:art 69.

Lillywhite A, Wolbring G. Coverage of Ethics within the Artificial Intelligence and Machine Learning Academic Literature: The Case of Disabled People. Assistive Technology. 2019;Apr:1-17.

McGowan J et al. PRESS Peer Review of Electronic Search Strategies: 2015 Guideline Statement. J Clin Epidemiol. 2016;75:40-6.

Meriam-Webster. Health Care. 2021a. Online: <https://www.merriam-webster.com/dictionary/health%20care>. Accessed 10 December 2021.

—. Public Health. 2021b. Online: <https://www.merriam-webster.com/dictionary/public%20health>. Accessed 10 December 2021.

Moher D et al. Preferred Reporting Items for Systematic Review and Meta-Analysis Protocols (PRISMA-P) 2015 Statement. Systematic Reviews. 2015;4:art 1.

Murphy K et al. Artificial Intelligence for Good Health: A Scoping Review of the Ethics Literature. BMC Medical Ethics. 2021; 22 art. 14:1-17.

Royal College of Physicians and Surgeons of Canada (Royal College). Task Force Report on Artificial Intelligence and Emerging Digital Technologies. Ottawa: RCPSC; 2020.

Topol E. Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again. New York: Basic Books; 2019a.

Topol E. High-Performance Medicine: The Convergence of Human and Artificial Intelligence. Nature Medicine. 2019b;25:44-56.

Tricco AC et al. 2018. PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation. Ann Intern Med. 2018;169(7):467–73.

United States Government Accountability Office (US GAO). Artificial Intelligence in Health Care Benefits and Challenges of Technologies to Augment Patient Care. Washington: US GAO; 2020.

University of Ottawa Centre for Health Law, Policy and Ethics (CHLPE). Machine M.D. 2021. Online: https://www.ottawahealthlaw.ca/projects/machine-m-d. Accessed 10 December 2021.

World Health Organization (WHO). Ethics and Governance of Artificial Intelligence for Health. Geneva: WHO; 2021.