AI & HEALTH CARE:
A FUSION OF LAW & SCIENCE

AN INTRODUCTION TO THE ISSUES

FEBRUARY 2021

This report was first drafted by Michael Da Silva in collaboration with the participants of the AI & Society workshop for AI & Health Care: A Fusion of Law & Science

AI & HEALTH CARE:
A FUSION OF LAW & SCIENCE

—

CIFAR EVENT
9 OCTOBER 2020
(ONLINE VIA ZOOM)

Various scholars, commissions, and task forces predict that Artificial Intelligence (AI) will transform and democratize health care systems, improving quality, safety, and efficiency, and offer the tantalizing possibility of empowering patients in directing their own health care.[1] The use of AI has the potential to improve human health at the systems and service delivery levels. Well-designed algorithms could, for instance, correct presently high rates of avoidable medical errors.[2] Yet, at the same time, the use of AI in health care presents clear risks and implementation issues, such as those connected to discrimination, informed consent, safety/quality (and liability for harm), and privacy.[3] Failure to regulate to minimize these risks could undermine Canada's ability to leverage the potential benefits of AI in health care. For instance, AI is unlikely to minimize medical error if it mirrors unrepresentative (and thus biased) training data that caused such errors.[4] The use of poor quality AI in health care could, in turn, lead to public pushback and subsequent onerous regulations that stifle innovation. This "inconvenient truth" could result in lost opportunities for valuable AI developments likely to improve health care and the loss of financial and research efforts to unusable or underused AI.[5]

AI & Health Care: A Fusion of Law & Science assembled an interdisciplinary group of experts in AI, law, ethics, policy, and medicine to address the core regulatory issue raised by these complexities: How can Canada maximize the potential benefits of the use of AI in health care while minimizing potential dangers? More specifically, the experts began examining whether existing laws (constitutional protections, legislation/regulations, case law, and "soft" law (e.g. professional ethical and research ethics codes) are sufficient to protect patients from harms and maximize benefits that may emerge from AI technologies and the form possible reforms may take. The event focused on the challenge of addressing privacy- and safety/quality-based concerns in legal regulations and related ethical issues surrounding the use of AI in health care.

ORGANIZERS:

· **Colleen M. Flood**
Law, University of Ottawa

· **Joelle Pineau**
Computer Science, McGill University/MILA/Facebook AI Research Montreal

· **Céline Castets-Renard**
Law, University of Ottawa/ Université Toulouse Capitole en France[6]

CASE STUDIES

The event began with a discussion of two 'case studies' in the use of AI in health care. AI is heterogeneous. Different kinds of AI produce different risks in different health care settings. Case studies help focus discussion on issues likely to arise in discrete settings, moving beyond abstract concerns to highlight specific potential benefits and risks. Discussing multiple case studies can highlight how benefits and risks differ across contexts and how regulators must attend to these differences to maximize good outcomes. The emergency triage- and COVID-19-based case studies presented at the event thus framed subsequent discussions, though these case studies alone do not exhaust all issues relevant to assessing the use of AI in health care.

A. EMERGENCY ROOM TOOL FOR IDENTIFYING & FAST-TRACKING PEDIATRIC PATIENTS NEEDING X-RAYS (DR. DEVIN SINGH, MEDICINE/COMPUTER SCIENCE, THE HOSPITAL FOR SICK CHILDREN)

First, Dr. Devin Singh discussed his work on automated machine learning-based medical directives to improve patient flow through emergency departments. Emergency departments are often over-crowded with patients experiencing long wait times to receive routine care from health care providers. The use of nurse medical directives is already viewed as an appropriate method to minimize these wait times, but they have many limitations. They interrupt triage workflows, are subject to human practice variation, and require recurring investments in nurse training. Machine learning-based medical directives overcome these limitations and can autonomously order downstream testing "to improve patient flow through the ED."

Per Dr. Singh, machine learning can improve this flow by leveraging triage data collected when patients first arrive to identify when patients may require testing (e.g., forearm X-rays) and automatically order these tests to be completed prior to being seen by an emergency physician. Dr. Singh's system is designed to "optimize for high positive predictive value ... to minimize overtesting by ensuring very low false positive rates." Taking an approach of ensuring low false positives has the consequence of models potentially missing cases of patients who do require testing. Hence, his proposed approach is a two pronged 'pathway' system in which individuals identified by a machine learning model can have the appropriate downstream testing ordered autonomously and those who are not selected simply undergo the standard pathway

of waiting for a health care provider to assess them prior to investigations being ordered. It is estimated that care for approximately 25 percent of emergency department patients can be expedited by this approach, allowing for system wide efficiencies that could reduce wait times for all patients. Dr. Singh discussed how he designed it and some research results on its potential utilization at the Hospital for Sick Children.

B. AI FOR COVID-19 (JOELLE PINEAU, MEDICINE, MCGILL UNIVERSITY/MILA/FACEBOOK AI RESEARCH MONTREAL, HTTPS://WWW.CS.MCGILL.CA/~JPINEAU/)

Second, Joelle Pineau discussed her ongoing collaborative work on the use of machine learning to improve resource allocation in hospitals in the context of the COVID-19 pandemic. Pineau is part of a team of researchers at the Centre hospitalier de l'Université de Montréal (CHUM), Mila, and Facebook that are building machine learning models from patient clinical data to target missing data imputation, the prediction of adverse events, and causal treatment effects in order to improve resource allocation at CHUM and similar hospitals. The group is part of an international data sharing and collaboration arrangement with NYU Langone Health.

Pineau outlined four properties of responsible AI that inform her own work. Per Pineau, responsible AI should be reliable and safe (viz., perform as expected without causing harm); private and secure; fair and inclusive; accountable (e.g., providing clear definitions of different parties' roles and responsibilities); and transparent (viz., understandable and available for audit).

COMMENTARIES

The case studies were followed by overviews by legal experts on some basic regulatory issues.

A. PRIVACY (TERESA SCASSA, LAW, UNIVERSITY OF OTTAWA, HTTPS://COMMONLAW.UOTTAWA.CA/EN/PEOPLE/SCASSA-TERESA)

First, Teresa Scassa discussed the basic framework for privacy-related issues in health care and its potential application in the AI context. Scassa noted that the general approach is highlighted by a focus on individual control of one's private health information. This is operationalized through requiring consent to collection, use, and disclosure of personal health information, limiting improper access to health information, and providing patient access to their health information. These requirements are found in a combination of privacy laws and research ethics norms. Yet other forms of personal information can be used without consent subject to research ethics board approval and even anonymized personal health information can be used without consent under certain circumstances (e.g., where seeking consent is impracticable and there are appropriate safeguards in place that protect against the risk of reidentification).

'Big data,' AI, and machine learning require large data sets where consent may prove impracticable, raising questions about which safeguards are appropriate for protecting against reidentification and the improper sharing of personal information, including personal health information. Questions about the risk of individual harm through reidentification, control over access to data, etc. must be addressed in a way that respects privacy and its larger role in a matrix of social and technical conditions (e.g., the relationship between privacy and quality or bias). There are some models for protecting personal health information in the big data/AI context. For instance, the European Union's General Data Protection Regulation contains AI-related rights.[7]

B. SAFETY/LIABILITY (LORIAN HARDCASTLE, LAW, UNIVERSITY OF CALGARY, HTTPS://LAW.UCALGARY.CA/PROFILES/HARDCASTLE)

Second, Lorian Hardcastle discussed the basic liability issues concerning the use of AI in health care. Hardcastle noted that when using AI, the liability for the health care provider will be greater if they maintain greater discretion in the decision-making process. As AI moves from providing information to providing diagnoses, to providing treatment recommendations, to applying treatment (in the stages discussed earlier by Joelle Pineau), health care providers may have less discretion and thus less liability (though liability may not perfectly track the stages).

To the extent that health care provider discretion remains, health care providers will (at least as a general matter) be held to the private law 'reasonableness' standard for their exercise of that discretion where 'reasonableness' is understood in terms of what other competent professionals in the field would do. New practice guidelines on the use of AI in health care could help set standards for reasonableness. There is some ongoing work on what those standards should look like. Yet questions remain over, for instance, when physicians can or should go against an AI recommendation – and how private law should treat decisions to ignore a recommendation.

Where there is less or no discretion, then liability will more likely accrue to the designer or manufacturer of the AI tool. There can, however, be instances where both the health care professional and the designer or manufacturer face liability. This basic framework mirrors the treatment of new technologies in the health care sector, though AI could raise new questions.

BREAKOUT SESSIONS

The commentaries were followed by three simultaneous 'breakout sessions' on (1) law and privacy, (2) law and safety, and (3) ethics. Rapporteurs then summarized the findings as part of a debriefing session. Each breakout session covered a wide variety of issues. The following is a stylized representation of the findings that focuses on some core thematic concerns in each.

BREAKOUT #1:
LAW AND PRIVACY

PARTICIPANTS:

- **Céline Castets-Renard**
  Law, University of Ottawa

- **Alex Dearham**
  Advisor, Ethics and Professional Affairs,
  Canadian Medical Association

- **Elaine Gibson**
  Law, Dalhousie University

- **Jean-Michel Loubes**
  Mathematics, Université de Toulouse

- **Florian Martin-Bariteau**
  Law, University of Ottawa

- **Diane McLeod-McKay**
  Yukon Information and Privacy
  Commissioner

- **Teresa Scassa**
  Law, University of Ottawa

- **Pascal Thibeault**
  Scribe; Law, University of Ottawa

- **Ma'n Zawati**
  Chair/Rapporteur; Law/Policy/Bioethics/
  Medical Ethics, McGill University

The breakout session on law and privacy focused predominantly on the single complex issue of how to strike the appropriate balance between protecting privacy and the need for innovation.

One central concern was whether AI-specific privacy regulation was necessary or whether health privacy reform in general would be a better avenue for any necessary new regulations.

One related concern was that incorporating AI into health care in a way that protects privacy rights may require major changes to the health care system. Attention should be paid to whether such changes are justified. While ensuring the privacy of medical data is crucial, participants cautioned against adopting laws that would be unfair to innovators. Some were also worried about the drawbacks of using legal frameworks that are not designed to fit or function within our health care system. There is a need to strike the right balance between allowing for useful innovation to enter our health care system while ensuring patient data privacy.

Participants noted that there is a movement towards implementing a necessity and proportionality-based regulatory framework into Canadian law. This framework would be borrowed from the human rights context and would look at the necessity of data collection and the proportionality and the privacy impact of that collection, and the goal it is meant to serve.

Questions were raised about how this could be implemented, especially given expertise asymmetries

between potential regulators and innovators, leading to further questions about who is best placed to make decisions on what counts as 'necessary' or 'proportionate.' Participants also noted that even experts do not always know what will be strictly 'necessary' ex ante.

Participants noted that 'necessity' to a stated goal may not track 'necessity' for broader systemic goals. For instance, race-based data collection may not be necessary for achieving a certain health outcome using AI but may be necessary to audit the potential biases of an AI tool.

Participants noted that overly onerous rules may lead to a lack of sufficient data to support innovation. Concerns were raised about the European system and how the restrictions lead to data opacity.

Concerns were raised about the possibility or desirability of a mere consent-based system. Current informed consent processes are largely pro forma affairs, and some argued that consent may not be valuable where those asked for consent lack expertise about AI, health care, etc.

Public interest-based frameworks may also be helpful. Those concerned with patient autonomy should, moreover, look beyond consent alone and examine how having a right to access one's own personal health information could be empowering from a patient perspective.

The importance of clarity in terms of what can be collected was highlighted by practitioners.

BREAKOUT #2:
LAW AND SAFETY

PARTICIPANTS:

- **Andrew Beckett**
  Medicine, Canadian Forces/ St. Michael's Hospital/University of Toronto

- **Jennifer Chandler**
  Law, University of Ottawa

- **Michael Da Silva**
  Scribe; Law, McGill University

- **Colleen M. Flood**
  Law, University of Ottawa

- **Michael Froomkin**
  Law, University of Miami

- **Anna Goldenberg**
  Computer Science, The Hospital for Sick Children/University of Toronto/Vector Institute/CIFAR

- **Lorian Hardcastle**
  Law, University of Calgary

- **Tanya Horsley**
  Associate Director of the Research Unit, The Royal College of Physicians and Surgeons of Canada

- **Teddy Weinstein**
  Law, University of Ottawa

- **Kumanan Wilson**
  Chair/Rapporteur; Medicine, Ottawa Hospital Research Institute/University of Ottawa

The breakout session on law and safety focused predominantly on (i) the quality of AI technologies, (ii) the appropriate use of AI technologies in clinical practice, (iii) health care provider responsibility and liability, and (iv) regulatory competence and responsibility.

I. THE QUALITY OF AI TECHNOLOGIES

One central issue concerned the current quality of AI technologies in health care — and how we should judge the quality of these technologies in the present and in the longer-term.

There appeared to be broad recognition that many AI technologies can, for instance, make more accurate and efficient diagnoses than their human equivalents, but none were so accurate that human users should defer to AI in all cases. There was also broad recognition that any evaluation of the use of AI in health care should compare regimes in which it is used and regimes in which it is not used. For instance, diagnostic systems with and without AI will both produce false positives and negatives. We cannot assume that either the status quo or change will be perfect. They should be evaluated together — at the individual and systems levels. These benefits may not overlap. For instance, even if an individual does not get a direct health benefit from being sent through a system more quickly, the system benefits from freeing a bed sooner. What to do when technology benefits an individual and not a system and vice versa was left for future study.

It was noted that there is heterogeneity in different health care providers' treatments of the same phenomenon and that AI could help address this, better standardizing care. Yet pertinent issues remain concerning (a) how to avoid overfitting and (b) how to regulate AI systems given their dynamic nature and the need to (i) ensure that they continue to produce good outcomes over time and (ii) prepare for the development of emergent behaviour.

There was a lack of consensus on whether AI 'updates' raise unique issues. Some believed that software updates receive insufficient scrutiny under the U.S. F.D.A. regulatory model. Others thought that AI only raises unique issues when it continuously learns. There were related debates on whether continuous learning AI should be held to the same safety, etc. standards over time — and how to regulate to ensure this if it is desirable. One possible way to address emergent behaviour is pharma-style 'life cycle' regulation.

Questions were also raised about the use of real-world data and whether companies can be trusted to be candid about its use or if regulation is needed to ensure it.

II. THE APPROPRIATE USE OF AI TECHNOLOGIES IN CLINICAL PRACTICE

The question of whether and how AI should guide clinical practice was central. There were suggestions that existing frameworks, such as practice guidelines, could be analogous to AI. Guidelines change over time, suggesting that they err, but are still useful. The same could be true of AI and we should consider treating them analogously.

There were, however, questions about whether the opacity of many AI tools makes them disanalogous to practice guidelines in morally or legally relevant ways. It was noted that there is already heterogeneity in how people follow practice guidelines. The same is true of attitudes towards AI. Senior physicians are less likely to rely on AI than junior physicians. A paradigmatic senior physician may seek a reason for the AI's contrary position before *considering* it while a paradigmatic junior physician will seek validation. Whether this is an issue was less clear. On one hand, more senior physicians appear to have better pre-test probabilities of their decisions being right. On the other, decision-making outcomes scores appear to lower as physicians age due to their not following guidelines (and following them tends to produce better outcomes).

Questions were also raised about whether 'over-reliance' could hinder the development of clinical judgment or if AI could otherwise impact clinicians' psychological makeup. The consensus seemed to be that AI tools should be used as a valuable source of 'advice' and that it is important to protect against over-reliance on AI and under-use of beneficial AI.

There was also a sense that even one *very* bad outcome due to a poor-quality tool could hinder the use of beneficial AI by leading to over-regulation tailored to that one bad outcome. Widespread use of such a 'bad' tool could have population-level health effects.

Questions about the representativeness of data were also highlighted (and arguably also raise quality questions above). It was noted that AI tools often rely on data that is not representative of the patient group. This is sometimes unavoidable (e.g., when the number of possible patients is too small to create data as in the Space Agency or military cases). Yet it is also sometimes a function of arguably discriminatory research practices that do not sufficiently study how different interventions impact vulnerable groups (e.g., women, Indigenous groups).

Questions were raised about when and how an AI tool trained on data for one population should be extended to another, though there was also a suggestion that this is not an AI-specific issue. It was suggested that a good AI tool should refuse to make a judgment on a person who does not fit the data profile on which the AI system was trained. This could be legally mandated. Another suggestion was to require transparency in data sources. Still another was to require a declaration on the suitability of the AI for different populations (analogous to what is in drug regulation).

Concerns were also raised about clinicians' trust in biased AI leading to more biased action that validates the AI, creating a 'feedback loop' of discrimination. This too raised questions about disclosure. It also raised questions about whether health care providers should be treated as responsible for that bias or any discrimination that follows, raising the next issue.

III. HEALTH CARE PROVIDER RESPONSIBILITY AND LIABILITY

Questions were raised about the extent to which health care professionals can be held responsible for the use (or non-use) of AI and the role that different regulations (professional regulatory agency rules, tort laws, etc.) should play in responding to that. Moral and legal questions included the following: Should a doctor rely on AI? What should we do if they choose not to when the AI is right or is more likely to be? What about when the tool is very risky, but they choose to use it? Should warning labels play a role in these determinations?

Many of these questions address basic issues in liability that appear in other contexts. There were suggestions that analogies could be raised with the use of practice guidelines. But arguments can always be made that a practice guideline should or should not have been followed in a given case and it is not clear that there is a uniform view of their appropriateness that can help us here. Moreover, the 'standards' in the AI context are arguably less transparent than elsewhere.

One suggestion is that health care providers should be prepared to express reasons for decisions to disagree with an AI recommendation and there was a thought that health care providers might be more likely to chart these points of disagreement.

The general 'reasonableness'-based 'standard of care' in medical cases could be consistent with various uses of AI in health care, including following the recommendations of a low quality AI tool that is in wide use and refusing to follow the recommendations of a high quality AI tool in less use. Indeed, using the high-quality tool could be

shown to violate the standard of care. This creates the risk of widespread adoption of a dangerous tool not creating liability. But even some 'bad' tools may be better than the status quo, so this may not be so bad. Concerns were also raised about whether any regulation could address the heterogeneity of AI. In any case, AI is likely to be built into the 'standard of care' for liability under existing tort law.

There was also a question about whether doctors should be held responsible for AI decisions and if keeping doctors in decision-making will make them responsible for things they do not understand and for which they should not be held accountable. If AI is 'really' making the decision and a doctor is a rubber stamp, questions arise as to whether they should be fully responsible for outcomes. Current liability rules would likely hold them fully accountable.

IV. REGULATORY COMPETENCE AND RESPONSIBILITY

While some persons suggested leaving things to professional judgment and/or innovators' own regulations, the government and professional colleges likely have a role to play in regulation. One approach would draw on an analogy with pharmaceutical regulations for regulatory ideas. But even insofar as the 'life cycle' approach to regulation from the pharmaceutical sector may seem appropriate here, there may be disanalogies insofar as 'drugs' are tested pre-market in a way that software updates are not and there are clear concerns that existing regulatory agencies may be too weak to fulfill their current mandates, let alone address these complex regulatory issues.

Concerns were also raised about 'regulatory capture' (e.g., strong connections between the regulated and regulators leading to a self-reporting scheme that does not actually constrain any actions (as with medical devices in the U.S.A.)) and about regulators setting goals for the use of AI that cannot be justified (e.g., the promotion of 'through-put' (procedure/person/shift)).

Some regulation seems necessary, but issues remain over how to regulate AI in ways that do not stifle promising innovations but address the safety- and liability-related issues above.

BREAKOUT #3:
ETHICS

PARTICIPANTS:

- **Louise Belanger-Hardy**
  Law, University of Ottawa

- **Cécile Bensimon**
  Director, Ethics and Professional Affairs,
  Canadian Medical Association

- **Jennifer Gibson**
  Chair/Rapporteur; Bioethics,
  University of Toronto

- **Gagan Gill**
  Knowledge Mobilization/Artificial
  Intelligence, CIFAR

- **Jason Millar**
  Engineering/Philosophy,
  University of Ottawa

- **Vanessa Gruben**
  Law, University of Ottawa

- **Lisa Schwartz**
  Philosophy, McMaster University

- **Devin Singh**
  Medicine/Computer Science,
  The Hospital for Sick Children

- **Ian Stedman**
  Public Policy/Law, York University

- **Kelli White**
  Scribe; Law, University of Ottawa

The breakout session on ethics highlighted that legal regulation must attend to ethical concerns. The session focused predominantly on (i) algorithmic bias, (ii) professional discretion, (iii) the location in which AI is used, and (iv) informed consent. One theme was that governance and ethics intersect across all levels of the AI technology life cycle.

I. ALGORITHMIC BIAS IN HEALTH CARE

One central theme was the threat of algorithmic bias in health care. While AI can be used to improve health outcomes, it may not do so uniformly across populations and many AI tools may not be responsive to differences between groups. One part of this discussion focused on possible causes of bias. Concerns that automation could introduce bias were raised. Whether this could be addressed by training to identify bias or not was one point of contention in the discussion.

While algorithmic bias occurs across many markers of identity, one of the core elements of the discussion during this breakout session focused on the use of race-based data in AI. AI can produce a form of algorithmic bias due to miscommunication of training data or the specific use/misuse of existing information. There are questions about whether the collection of race-based data can help address these concerns (by e.g., ensuring that algorithms are based on data representative of the wider population for whom AI tools will be used) or will instead exacerbate or create new inequities (by e.g., enabling the use of race-specific data for discriminatory purposes). There is a mandate or a priority for institutions in Ontario to start collecting race-specific data, but federalism and privacy, for just two examples, present challenges. Absent a specific proposal on how to weigh these concerns, there was a general consensus on the need to develop governance around the rigorous collection of data around race and ethnicity.

There was one suggestion that any bias was likely imputed from those entering the data and that tracking the input of data and remaining aware of its potential could address some concerns.

A related discussion highlighted the importance of attending to whose values are driving and reflected in AI in health care. Participants suggested asking questions like the following: Who is driving decisions? Whose values are being represented? Who makes the choices and identifies the values that need to be included in AI applications?

II. PROFESSIONAL DISCRETION

Another concern, which was also raised in other sessions, concerned professional discretion. Questions were raised about how to allocate responsibility for decision-making between the AI and the professional decision-maker – and whether their judgments can be 'blended.' There was again an emphasis on the importance of addressing cases of disagreement. While everyone agreed that patient safety should be a core concern in allocating liability, questions about who should be responsible for what forms of safety regulations remained unaddressed.

Related issues focused on what health care providers should do when they lack expertise in AI and the level of expertise that should be expected. Concerns about automation were again raised here as AI could lead to a form of 'deskilling' that undermines attempts to create mechanisms for health care providers to serve as 'checks' on AI (by e.g., flagging concerns for investigation).

III. THE LOCATION IN WHICH AI IS USED

Participants also highlighted that different contexts of use raise different ethical issues. For instance, the AI tools that can be used in an emergency room differ from those that can be used at home. While alarms can be used in both emergency room monitoring tools and a smartwatch, they may raise different ethical issues in terms of safety and liability. Moreover, location can play a role in access and equity. This point is not exhausted by concerns about access to health care providers. Differential access to the Internet can lead to differential access to otherwise 'at-home' AI tools that could be relevant to the provision of health care.

IV. INFORMED CONSENT

Finally, questions were raised about informed consent to use of new and emerging AI technologies, including the scope of consent one can be expected to provide.

Where AI algorithms learn over time, questions arise as to whether one can know the full extent of that to which they are consenting. There are, moreover, concerns about whether people understand AI well enough to be 'informed' about consent even absent these uncertainties. Questions about whether vulnerable populations required specific protections were also raised. The use of data by third parties, with or without patient consent, was also raised as an issue here.

CONCLUSION

Even the examination of two case studies in the use of AI in health care highlights the many difficulties of regulating such use in a way that leverages its transformational potential to improve health care in Canada while minimizing the attendant risks of AI in that setting.

Our interdisciplinary groups highlighted a large number of potential regulatory issues, including

- whether a proportionality framework is appropriate for addressing privacy concerns
- whether a consent-based regime is desirable or possible in the AI context
- how to ensure that AI tools in health care settings meet safety and quality standards
- possible issues with health care providers' understanding of AI-based tools
- what to do when AI and health care providers disagree on an option
- how to allocate liability for the use — or non-use — of AI tools in health care
- the appropriate regulators and forms of regulation for different AI-related issues
- the difficulty of providing informed consent to the use of AI tools
- the threat of potential biased inputs in and discriminatory outputs of AI tools
- the appropriateness of different comparison cases when weighing regulatory options
- how to account for AI's heterogeneity in any form of regulation

This list is non-exhaustive, and some concerns were unique to particular breakout sessions. It is, however, notable that concerns with (i) allocating liability for the use of AI in health care, (ii) minimizing the threat of algorithmic bias, and (iii) protecting privacy without stifling innovation were consistent across groups. This is, perhaps, a function of our case study selection and our identifying some of these issues as central topics for the 'breakout sessions.' Yet, these areas also mirror some of the core regulatory challenges identified in other works.[8] Regulators should thus consider prioritizing these issues and can use the points highlighted above to guide their work.

Our interdisciplinary group(s) also began to highlight the complex regulatory environment in which AI-based health care tools are being introduced. While there is no bespoke legislation regulating AI in health care in Canada, many existing regulations do exist. A complex collection of federal and provincial laws, policies, and professional regulatory documents address at least parts of many of the issues above, albeit not in an explicitly coordinated manner. For instance, federal medical device regulations, federal and provincial privacy and anti-discrimination laws, existing rules on liability in private law, and professional regulations each address parts of (i)-(iii).[9] Yet it remains to be seen whether they are sufficient to minimize the risks of AI in health care identified above without unduly stifling innovation and the potential to improve our health.

These considerations highlight the need for closer scrutiny of existing regulations. Future work should examine whether existing laws are well-suited to address the potential regulatory issues above. Scholars should then study whether other regulatory options (from other domains or comparative law and/or policy) are better suited to address these complex regulatory issues. This research team will begin this work in its next CIFAR-sponsored session, which will focus on the regulation of AI in/as medical devices. That session will address safety and quality-related concerns but will also provide opportunities to further discuss privacy and discrimination.

ENDNOTES

1.   See e.g., scholars like Eric J Topol (Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again (New York: Basic Books, 2019) [Topol, Deep]; "High-Performance Medicine: The Convergence of Human and Artificial Intelligence" (2019) 25 Nature Medicine 44) or reports like The White House, "Artificial Intelligence for the American People" (11 February 2019) The White House, online: https://www.whitehouse.gov/ai/; European Union, "Connected for a Better Future: Digital Transformation. Empowered Citizens. Healthier Society." European Commission (18 September 2019), online: https://ec.europa.eu/digital-single-market/en/news/connected-healthy-future-brochure; Royal College of Physicians and Surgeons of Canada, Task Force Report on Artificial Intelligence and Emerging Digital Technologies (Ottawa: Royal College of Physicians and Surgeons of Canada, 2020); and CIFAR AI for Health Task Force, Building a Learning Health System for Canadians – Report of the Task Force on Artificial Intelligence for Health (AI4Health) (Toronto: CIFAR, 2020), online: https://cifar.ca/wp-content/uploads/2020/11/AI4Health-report-ENG-10-F.pdf.

2.   Topol, Deep, ibid makes this claim. On the reduction of medical errors using A.I., see D Douglas Miller & Eric W Brown, "Artificial Intelligence in Medical Practice: The Question to the Answer?" (2018) 131 The American Journal of Medicine 129. On Canada's own rates of medical errors, see G Ross Baker et al., "The Canadian Adverse Events Study: The Incidence of Adverse Events among Hospital Patients in Canada" (2004) 170(11) CMAJ 1678.

3.   See e.g., Sara Gerke, Timo Minssen & I Glenn Cohen, "Ethical and Legal Challenges of Artificially Intelligence-Driven Health care" in Adam Bohr and Kaveh Memarzadeh, eds, Artificial Intelligence in Health care (Cambridge: Academic Press, 2020) 295 (outlining the main ethical issues plausibly requiring a regulatory response).

4.   E.g., Melissa D McCradden et al., "Ethical Limitations of Algorithmic Fairness Solutions in Health Care Machine Learning" (2020) 2(5) The Lancet Digital Health E22; Carmen Neithammer, "AI Bias Could Put Women's Lives at Risk – A Challenge for Regulators" (2 March 2020) Forbes, online: <https://www.forbes.com/sites/carmenniethammer/2020/03/02/ai-bias-could-put-womens-lives-at-riska-challenge-for-regulators/#352fecf2534f>.

5.   Trishan Panch, Heather Mattie & Leo Anthony Celi, "The 'Inconvenient Truth' About AI in Health care" (2019) 2 npj Digit Med 77. See also e.g., Nathan Cortez, "Regulating Disruptive Innovation" (2014) 29 Berkeley Tech L J 176; Christian Lovis, "Unlocking the Power of Artificial Intelligence and Big Data in Medicine" (2019) 21(11) J Med Internet Res E16607.

6.   Nb. Professor Ian Kerr played a central role in organizing the event proposal prior to his death in 2019.

7.   See European Parliament and Council of European Union, (2016) Regulation (EU) 2016/679.

8.   Recall note 3.

9.   For an overview of existing laws, see e.g., Colleen M Flood & Catherine Régis, "AI & Health Law in Canada" in Florian Bariteau-Martin & Teresa Scassa, eds, Artificial Intelligence and the Law in Canada (Toronto: LexisNexis, Forthcoming) TBD.