CIFAR

MACHINE MD:
Law and Ethics of Health-Related AI Case Study 1:

The OR Black Box

Workshop held: March 4, 2022

Report published: March 10, 2023

This report was drafted by Nicole Davidson and Sophie Nunnelley in collaboration with the participants of the Machine MD: Law and Ethics Case Study on OR Black Box.

Acknowledgments

This event was co-hosted by CIFAR and the Canadian Institutes of Health Research (CIHR)-funded *Machine MD: How Should We Regulate AI in Health Care?* project, with support from the Alex Trebek Forum for Dialogue. It is part of CIFAR's AI & Society Program. CIFAR's leadership of the Pan-Canadian AI Strategy is funded by the Government of Canada, with support from Facebook and the RBC Foundation. The organizers thank CIFAR, CIHR, and the Alex Trebek Forum for Dialogue for their support.

Citation

N. Davidson, S. Nunnelley, A. Goldenberg, C. Régis, C. M. Flood, T. Scassa, F. Rudzicz, N. Cortez, I. Stedman, F. Martin-Bariteau and the workshop participants, Machine MD: Law and Ethics of Health-Related A.I. Case Study 1: The OR Black Box (Toronto: CIFAR, 2023).

Table of Contents

Law and Ethics Case Studies in Health-Related AI

The Machine MD project is committed to the use of case study analyses to explore the law and ethics of health-related artificial intelligence (AI). The team seeks to identify and analyze the legal issues associated with AI in healthcare by looking at real technologies, identifying any issues they raise, and analyzing how they are treated in Canadian and foreign law. The objective of these case studies is to move beyond abstract concerns into concrete realities, helping to inform law reform with a better understanding of real-world applications. The goal is to support beneficial AI technology innovation, while minimizing associated risks through appropriate legal governance.

The Machine MD team and CIFAR are partnering to host events dedicated to these case studies. Each event assembles an interdisciplinary group of experts in AI, law, ethics, policy, and medicine to discuss regulatory issues raised by a specific AI technology. These events follow earlier AI & Health Care: A Fusion of Law & Science collaborations.[1] This report summarizes the findings of the first of three online case study events in the spring of 2022. The two later events discussed the Suicide Artificial Intelligence Prediction Heuristic or "SAIPH" (March 11, 2022) and "digital twins" technology (April 1, 2022).

[1] See: *AI & Health Care: A Fusion of Law & Science – An Introduction to the Issues*, drafted by Michael Da Silva in collaboration with the participants of the AI & Society workshop for AI & Health Care: A Fusion of Law & Science (Toronto: Canadian Institute for Advanced Research, 2021), online: <https://uploads-ssl.webflow.com/5e94a26db210b579bca67e7c/60b15d4338d77688f3056d12_604140e8419b84275713ca86_CIFAR%20AI%20Report%20(Final).pdf>; *AI & Health Care: A Fusion of Law & Science – Regulation of Medical Devices with AI,* drafted by Michael Da Silva in collaboration with the participants of the second AI & Society workshop for AI & Health Care: A Fusion of Law & Science (Toronto: Canadian Institute for Advanced Research, 2021), online: <https://uploads-ssl.webflow.com/5e94a26db210b579bca67e7c/60b159da764a10e80837a75a_AI-Healthcare-A-Fusion-of-Law-Science-II.pdf>.

Case Study #1: The OR Black Box

4 March 2022 (Online via Zoom)

The OR Black Box, which is produced by Surgical Safety Technologies, Inc., is an AI-enabled tool that seeks to improve patient safety and reduce preventable adverse events in the Canadian healthcare system and worldwide. The OR Black Box records audio and video in the operating room (OR) and laparoscopically. These recordings are reviewed by AI and human analysts to identify distractions, instances of error, and to inform best practices.[2] Its primary goal is to minimize adverse events and maximize patient safety by providing healthcare teams with a quantitative, anonymized analysis of the surgery, providing a quality assurance and educational tool to clinicians.

The OR Black Box currently provides teams with nonpunitive feedback about what went well, and what could be improved, to ensure efficiency and positive patient outcomes during surgical procedures. In the future it could be used to support surgical teams in real-time, permitting them to 'course-correct' to prevent mistakes from being made. However, the tool also presents a series of risks, for instance, relating to patient and provider privacy and informed consent, which may challenge existing legal regulation and require appropriate reform. This event examined the OR Black Box's potential benefits, risks, and associated regulatory challenges, through a presentation by one of its developers, commentaries by legal scholars on three legal issues raised by the tool, and breakout sessions where participants sought to better understand – and help resolve – problems.

[2] Jessica Jue et al, "An Interdisciplinary Review of Surgical Data Recording Technology Features and Legal Considerations" (2020) 27(2) Surgical Innovation 220 at 224.

Presentation by Frank Rudzicz, Surgical Safety Technologies Inc.

Frank Rudzicz, one of the developers of the OR Black Box, began the discussion by explaining the Black Box and its affiliated products. He first explained the use of black box technology in different applications. For example, a black box is used on aircrafts to record flight data and provide an objective analysis of crashes or other occurrences. He then focused his discussion on the OR Black Box, which relies on a collection of sensors, video cameras, and microphones in the OR to capture intraoperative details and the OR environment. These details are then analyzed by both an AI platform and human analysts to, for instance, (1) detect risks and hazards; (2) detect adverse events (such as active bleeding or injury); (3) mitigate risks; and (4) track interactions between team members during surgery. The developers of the OR Black Box hope the tool will empower healthcare organizations by providing clinically-relevant analytics from surgeries, thereby improving surgical quality and efficiency, and reducing avoidable adverse events by capturing what works well (and what does not work well).

Rudzicz also explained the OR Black Box's potential educational and quality improvement benefits. Surgeons traditionally use Objective Structured Assessment of Technical Skills (OSATs), which are a quasi-subjective human analysis of what went well (or not) during a surgery. However, this approach often does not provide meaningful feedback. The OR Black Box offers a quantitative analysis of technical skills over time, which can then be used by program directors to see how trainees are progressing. A dashboard program called "Black Box Insights" also provides customers with a range of analytics relating, for instance, to how many cases are performed on time; whether time is being used effectively; and how to optimize one's tray (for example, it provides data on which tools are to be used during a given surgery and which should be used in the future). Anonymized laparoscopic videos can also be uploaded to the Black Box Explorer platform and graded by peers and ranked for performance.

Videos from surgeries are stored in a double encrypted, secure cloud, and are accessible by the healthcare team and Black Box analysts. The OR Black Box uses the "Swiss Cheese Model" to de-identify videos and analyses that are reported to customers, to protect the privacy of both

patients and staff.[3] De-identifying is initially done by AI, which blurs faces, bodies, shoes, and items that might identify the room the surgery took place in (wall clocks, etc.). Human analysts then review the videos and complete any de-identification that the AI missed. Privacy of both patients and staff is paramount to ensure that staff feel safe and that the purpose of the tool remains quality improvement and education, rather than apportionment of liability.

In response to questions, Rudzicz noted that his team intends for human analysts to remain involved in the evaluation process for as long as possible despite OR Black Box use. Human intuition can better draw novel inferences that AI might not understand, such as noticing a link between a healthcare provider breaking sterility (e.g., not changing their gloves) and a subsequent adverse event.

[3] For a description of this model see e.g. "Swiss Cheese Model" (undated) The Decision Lab, online: <https://thedecisionlab.com/reference-guide/management/swiss-cheese-model>.

Commentaries

The legal commentaries focused on three issues that have been discussed at previous CIFAR events and that were pre-identified by planners as raising potential issues for the OR Black Box.

A. Liability (Nathan Cortez, SMU Dedman School of Law)

Nathan Cortez delivered a presentation on the potential liability issues raised by the OR Black Box, focusing on how OR Black Box information might be used in medical malpractice litigation. He explained that traditional medical malpractice litigation requires that significant time be devoted to reconstructing what happened during a surgery. AI technology can mitigate this issue by offering video footage of the surgery, providing objective insight into what happened. Yet, this also presents concerns, for instance, about which data collected by OR Black Box is privileged, and what might be discoverable and subject to litigation.

Cortez discussed the potential application of American peer review privilege laws, which protect certain hospital data from disclosure in litigation.[4] He explained that most American states require hospitals to have internal peer review procedures, which are nonpunitive and intended to improve quality and ensure staff are meeting the standard of care. Records and communications generated through these review procedures are generally protected from discovery and litigation. Moreover, the records are confidential in most states, meaning they cannot be disclosed to third parties outside of the hospital. Cortez explained this depends on the state, however; some states provide narrower privileges that only protect records generated for the purpose of peer review. In these states a plaintiff could argue that raw video of their surgery generated by OR Black Box is a key component of their medical record, which would subject it to discovery. In broader privilege states, peer review records and any records triggering a peer review are protected.[5] Cortez noted that in these states it might be harder for patients to argue that OR Black Box recordings fall outside of the scope of the peer review process.

Cortez noted the limited American jurisprudence regarding the discoverability of surgical videos

[4] Nathan Cortez, "A Black Box for Patient Safety?" (2019) 68 DePaul L Rev 239.
[5] *Ibid* at 246.

(in both narrow and broad privilege states). In the existing jurisprudence, however, courts seem to consider the purpose for which the video was made to be indicative of whether it will be admissible in court. Videos made for educational or quality improvement purposes have been held inadmissible. However, a video made for the purpose of forming part of the patient record could be subject to discovery. Cortez emphasized that these inferences are based on limited jurisprudence, and that it is impossible to conclusively predict how the courts will treat AI data in medical malpractice cases going forward.

Cortez also noted the uncertainty about how liability will be apportioned when things go wrong while using the OR Black Box. Currently, OR Black Box provides clinical decision support, leaving the ultimate decisions about how to provide care to the team performing the surgery. Given that decision making remains with the surgical team, Cortez predicted that questions of liability would be unchanged from the current standard, which focuses on the liability of the healthcare provider, not the device. However, as AI health tools become more autonomous, it is possible that less liability will rest on the physician, and more will be distributed to the hospital that decided to adopt the particular technology.

B. Informed Consent (Ian Stedman, York University)

Ian Stedman delivered a presentation on informed consent, emphasising the important links between informed consent and patient autonomy, and how AI data collection changes the consent conversation. Stedman noted that the law in Ontario and similar provinces views consent as an ongoing process, rather than a discrete event. The patient must be informed about the procedure itself as well as its potential benefits and risks and the benefits and risks of similar procedures.[6] He suggested that, in the context of AI, securing genuine informed consent to treatment might also require informing patients that their data is being collected, the purposes of such data collection, any de-identification processes, and any risks of re-identification.

Stedman explained that the informed consent structure will vary with the use of the technology. If AI tools are used in hospitals for quality improvement purposes, risks may not need to be disclosed at all under current law. However, if the OR Black Box or other AI health technologies begin providing real time feedback to surgeons, Stedman suggested, consent for the use of AI will likely be required. Patients will need to understand how this feedback informs decision making

[6] *Health Care Consent Act*, SO 1996, c 2, Schedule A, ss 10, 11.

during their surgery and should be permitted to opt out of AI use. Stedman noted that providing this information in a way that patients understand may prove difficult. Conversations about, and forms used in, consent processes are long and difficult for the average patient to understand. He emphasised the need to carefully incorporate information about AI-specific factors without overwhelming patients.

Stedman also discussed the possibility that, with increased use of AI, the content of the required consent procedure may change. Under current law, healthcare providers must consider what the "reasonable patient" would need to know about the tool in order to be considered informed. He suggested that as the public gains a stronger understanding of machine learning, less AI-specific information might be required in a given case.

C. Privacy (Florian Martin-Bariteau, University of Ottawa)

Florian Martin-Bariteau outlined the data protection and privacy framework in the private and public sectors. He explained that personal information is defined in the *Personal Information Protection and Electronic Documents Act* (*PIPEDA*) as any information that can be used to identify someone.[7] However, most privacy laws in Canada are antiquated and do not contemplate the challenges that AI pose to privacy. Martin-Bariteau noted, for example, that de-identified or anonymized information is not considered "personal information", and therefore is not protected by statute. While bills have been introduced at the House of Commons to address risks of re-identification, there is ongoing debate about whether personal information can ever be fully de-identified.[8]

Martin-Bariteau also discussed best practices in the context of AI and privacy. He noted, first, that using appropriate security protocols when storing and accessing data is crucial. This includes storing data in a cloud within the same jurisdiction. Organizations should also have protocols in place to mitigate data breaches. Second, Martin-Bariteau argued privacy protection should be "by

[7] *Personal Information Protection and Electronic Documents Act,* SC 2000, c5, s 2(1).
[8] The federal government's proposed consumer privacy protection legislation was first set out in Bill C-11, *An Act to enact the Consumer Privacy Protection Act and the Personal Information and Data Protection Tribunal Act and to make consequential and related amendments to other Acts,* 2nd Sess, 43rd Parl, 2020 (first reading completed 17 Nov 2020), online: <https://www.parl.ca/LegisInfo/en/bill/43-2/c-11>. After this case study event a revised version of the bill was introduced, Bill C-27, *An Act to enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act and to make consequential and related amendments to other Acts*, 1st Sess, 44th Parl, 2022 (introduction and first reading, 16 June 2022), online: <https://www.parl.ca/legisinfo/en/bill/44-1/c-27>.

design" – that is, innovators should consider best privacy practices from the development stage through to implementation, to ensure it is intrinsic to the design and not an afterthought. Indeed, he observed that privacy is less of a roadblock to innovation when it is approached in this way. Finally, Martin-Bariteau emphasized the principle that limits collection to the smallest amount of data possible to achieve the intended objective. Other key principles for collecting private health information include obtaining consent for the collection, use, and disclosure of data, and ensuring it is collected in a secure manner.

Martin-Bariteau also noted that consent can be a roadblock to innovation in the health sector, since AI algorithms rely on a diverse and vast collection of patient information in order to learn and improve. He discussed possible solutions, for instance, having legislation that alters the requirement for consent with respect to healthcare improvements. He suggested that, in general, a balance must be struck between promoting machine learning and preserving patient autonomy.

Breakout Sessions

The commentaries were followed by breakout sessions on liability, privacy, and informed consent. Rapporteurs then summarized the findings during a debriefing session. The core thematic concerns arising in each session are summarized below.

Breakout #1: Liability

Attendees: Lorian Hardcastle (Rapporteur), Nicole Davidson (Scribe), Anna Goldenberg, Christian Blouin, Frank Rudzicz, Justine Gauthier, Michael Froomkin, Nathan Cortez, Sana Tonekaboni

The breakout session on liability focused predominantly on (i) apportionment of liability when AI is informing decisions, and (ii) how AI advancements will affect the adjudication of medical malpractice matters.

I. Apportioning Liability

One concern raised by the group was how AI decision making autonomy might affect physician liability. Participants were generally of the view that this concern is conjecture, given that AI autonomy is not yet a reality, and that the discussion should focus on AI tools that provide clinical decision support without driving decision making. This is consistent with current tort law, which focuses on the liability of individual actors such as physicians. However, the group noted that as AI becomes more autonomous, and informs decision making to a greater degree, there could be a shift towards enterprise (hospital or manufacturer) liability.

Participants also noted that the OR Black Box is not yet classified as a medical device by the FDA. Rather, it is classified as a clinical decision support tool, where the ultimate decision making remains with the healthcare team. They noted this protects developers from liability in medical negligence, but that this regulatory classification could eventually change.

II. Using AI Data in Litigation

Participants discussed how the standard of care and causation of harm are determined in court and how the OR Black Box data might affect these processes. Currently, the courts rely on testimony from competing expert witnesses who have reviewed the available records, to draw retrospective inferences regarding the probable cause of harm to the patient. Participants noted the OR Black Box's potential to provide a clearer and more objective assessment of what happened, which could help resolve the "battle of the expert witnesses" issue. However, questions arose regarding whether the courts would rely on just the video and audio recordings of a surgery, or whether the insight reports produced by the OR Black Box would also be admissible.

Participants also discussed how best to protect the interests of patients in medical malpractice matters. For instance, they raised the benefits of a no-fault compensation system. In such a system, rather than seeking compensation for injuries through medical malpractice suits, injured patients could receive government-funded compensation and relinquish their right to sue for damages. Participants suggested this would benefit the patient, who could receive compensation for minor injuries without having to navigate the court system, and could alleviate physician concerns about liability arising from using new technologies. They also emphasized that the OR Black Box should ultimately improve patient safety, so the number of medical malpractice suits should decrease as a result of adopting this technology.

III. Standard of Care

Participants discussed the fact that tort law does not apply a 'standard of perfection' to physicians and suggested such a standard should not be imposed on technology either. Supporting this approach, they referred to the report *To Err is Human: Building a Safer Health System*, which shifts the focus from pointing fingers for honest mistakes to encouraging quality improvement.[9] Group members noted that, consistent with this report, the OR Black Box focuses on maximizing quality. They argued it should remain a tool to support learning and improvement in the OR rather than become a tool to punish human error.

[9] Linda T Kohn, Janet M Corrigan & Molla S Donaldson, "To Err is Human: Building a Safer Health System" (Washington DC: National Academies Press, 2000) online: <https://pubmed.ncbi.nlm.nih.gov/25077248/>.

Participants also noted that while AI technology in the healthcare setting is currently novel, it could eventually become expected practice and a component of the standard of care. Participants cited *Good Machine Learning Practice for Medical Device Development: Guiding Principles*, created by the FDA, Health Canada, and the United Kingdom's Medicines and Healthcare products Regulatory Agency.[10] The document sets out principles that are not determinative of the standard of care, but that could help inform courts' positions about what is required of practitioners who rely on machine learning technologies in providing care.

[10] US Food and Drug Administration, Health Canada, Medicines and Healthcare Products Regulatory Agency, *Good Machine Learning Practice for Medical Device Development: Guiding Principles* (October 2021), online: <https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/good-machine-learning-practice-medical-device-development.html>.

Breakout #2: Privacy

Attendees: Ira Parghi (Rapporteur), Michael Da Silva (Scribe), Gagan Gill, Florian Martin-Bariteau, Teresa Scassa, Elissa Strome, Pascal Thibeault

The breakout session on privacy focused on three issues: (1) whether OR Black Box data is information necessary for care and the implications of this distinction for consent, (2) employee privacy, and (3) the relationship between privacy and liability.

I. Information Necessary to the Provision of Care

Participants noted that the collection of identifiable personal health information can generally only be collected with patient consent in Ontario and provinces with similar privacy laws. While there is an exception for information necessary for the provision of care, participants debated whether data collected by the OR Black Box would fall under this exception.

Where the tool is primarily used for quality improvement and risk management purposes, some participants felt that the resulting data is not necessary for care. However, participants also noted that what is 'necessary for care' can change over time; if the OR Black Box becomes standard practice, its use could become necessary for care and no longer require consent. The opt-out framework for consent already used in facilities where the OR Black Box has been deployed suggests it may already be becoming standard practice. Whether this is appropriate remains debatable.

II. Employee Privacy

Participants also discussed employee privacy issues that arise from being recorded. They noted that footage of employees working in their professional capacity is generally not considered personal information, meaning that employees would not have a legal privacy right in respect of such footage. Yet questions arose about whether this changes if employees are being recorded for multiple hours. They asked, for instance, whether such a recording could be deemed a form of surveillance, and if so, whether employees would have a reasonable expectation of privacy upon which the tool might infringe. While arguably this is analogous to the constant recording on airplanes (where recordings also pick up personal information shared during long work shifts), participants noted that black box recordings from airplanes are generally only consulted after an

incident occurs. This is different from the present case where all recordings feed into quality assurance frameworks and other information can be picked up in the process. Participants also discussed the potential for problematic usage of the recordings, for instance, for employment-related purposes (contract renewal decisions, sanctions, etc.), noting this could limit providers' willingness to use the tool. They also discussed the potential for surveillance to increase anxiety in employees and affect their relationships with their employers and co-workers. A related concern was whether the tool's efficiency gains might come at the expense of basic human interests, such as finding meaning and interpersonal connections in one's work.

III. Privacy and Liability

Finally, participants noted intersections between privacy and liability concerns. If the OR Black Box is consistently used as part of a quality assurance framework, its data may remain privileged and shielded from judicial scrutiny in cases where there has been patient harm, and the OR Black Box may have captured relevant information. If the OR Black Box begins influencing clinical decision making, however, it will become harder to keep the recorded material privileged. Employees might feel more comfortable using the OR Black Box if the information obtained through its operation is not discoverable in litigation. This could improve uptake of the OR Black Box tool, which would be valuable if the tool does improve care. But questions arose as to whether it would be appropriate to shield evidence of violations of the standard of care from review. Strong privacy requirements may limit patients' ability to recover when they are harmed.

Breakout #3: Informed Consent

Attendees: Jennifer Gibson (Rapporteur), Caroline Mercer (Scribe), Colleen M. Flood, Bryan Thomas, Cécile Bensimon, Ian Stedman, Lindsay Thompson, Sylvain Bédard, Vanessa Gruben.

The breakout session on informed consent focused on what it means to obtain informed consent with the rise of AI technology and associated data collection in healthcare.

I. What are Patients Consenting to?

Participants focused on articulating what, specifically, patients are consenting to when it comes to the use of OR Black Box technology. Distinctions were raised between consenting to (i) the OR Black Box's video recording (data collection), (ii) the use of an AI tool for quality assurance or

educational purposes, and (iii) the possible eventual use of the OR Black Box to inform treatment (affecting the surgical procedure itself). They emphasized that each of these raises different consent considerations.

For instance, participants agreed that transparency and patient autonomy are paramount considerations but noted informed consent to use the OR Black Box may not be required if it is used only for quality improvement and educational purposes. If the OR Black Box begins driving clinician decisions, on the other hand, they agreed it would be necessary to have a thorough informed consent discussion, and to fully answer patient questions, in advance of surgery.

With respect to data collection, participants discussed how the OR Black Box technology differs from other information collection practices that are already commonplace in medical settings, such as the collection of personal information when patients log into hospital Wi-Fi. They reasoned that data collected by the OR Black Box differs because the AI technology may eventually make real-time decisions altering the course of the surgery, which would require a detailed informed consent procedure. Another analogy employed was that of providers obtaining consent to create a video recording of a virtual medical visit. One participant noted that OR Black Box technology may seem more intrusive because personal details of the surgery would be visually recorded. Participants noted that patients would likely want to know their surgery was being recorded and whether their data was being shared with a third party outside the hospital. The group emphasized that transparency is crucial to maintaining the public's trust in the healthcare system.

II. Obtaining Informed Consent: Disclosing the Risks and Benefits

Participants discussed the risks and benefits of using – and of not using – the technology, noting both should be discussed with patients to ensure consent is informed. They noted that should the technology become the standard of care, it will be particularly important to detail the risks of forgoing its use. As for risks, participants highlighted the possibility of re-identification, which is especially concerning in rare disease cases, which can be unique to a few patients in a given location or hospital setting. In terms of benefits, the group discussed the value in contributing personal data to machine learning with the goal of improving the quality of health care in the future. Participants also emphasized the importance of looking at the benefits of the technology – especially given the inherent danger of surgery – and being careful to not overstate its risks.

Finally, participants noted that transparency is key to consent discussions, especially when using recording devices in surgical settings. They emphasized that patients should be aware that the surgery is being recorded, and informed of any sharing of their data with third parties.

Conclusion

This case study highlights the significant benefits that OR Black Box technology may provide in relation to quality assurance in the surgical setting. However, it also raises questions about associated risks and their appropriate regulation.

Some of the following themes emerged during the presentations and breakout sessions:

- Whether OR Black Box data collected for the purposes of quality improvement and education could be discoverable in litigation, and if so, which data would be admissible
- Whether legislation should privilege such data, excluding it from disclosure in litigation, and if so, which components should be privileged
- Whether the data would form part of the patient record, and if so, whether it would be limited to the raw video and audio, or extend to insight reports
- How to apportion liability as the OR Black Box becomes more autonomous and / or is used to inform real-time treatment decisions
- The difficulty of obtaining informed consent with respect to the use of AI technology and data collection
- How the informed consent discussion needs to evolve as AI health tools become more autonomous
- The need for legislation contemplating the risk of anonymized data being re-identified, and whether data is ever fully de-identified if this risk exists
- Whether obtaining patient consent for quality assurance tools is a roadblock to innovation
- Whether the standard of care will evolve to require the use of AI tools such as the OR Black Box
- Whether an opt-out consent system is sufficient and informed
- Whether employees will have a right to privacy when they are subjected to what could be described as surveillance

This is a non-exhaustive list, and some concerns were unique to particular breakout sessions. However, concerns regarding (i) apportionment of liability for patient harm, especially as the OR Black Box becomes more autonomous and begins to inform treatment decisions, (ii) when and how informed consent principles should apply to the OR Black Box in its various possible usages, and (iii) privacy risks, including the risk of data re-identification, were common across groups.

Moreover, participants generally agreed that at least some of these issues may require new legislative frameworks. On the issue of liability and regulating the allocation of risk, some suggested legislators could look to the jurisprudence and regulation of driverless cars to inform the regulation of increasingly autonomous AI. On other issues, such as privacy, participants noted the more immediate need for reform. They noted that Canadian legislation currently does not recognize anonymized data as personal information, despite the risk of re-identification, leaving Canadians vulnerable to the consequences of a data breach. They also discussed the need for legislation that provides more control to individuals over their personal data and levies larger sanctions on companies that violate privacy. At the same time, participants recognized the need for balance so that companies can continue to innovate and allow their algorithms to learn from a diverse dataset. This case study identifies questions and issues that regulators, legislators, and policy makers should consider prioritizing going forward.

Workshop Participants:[11]

Frank Rudzicz	Michael Da Silva	Sylvain Bédard
Nathan Cortez	Gagan Gill	Vanessa Gruben
Ian Stedman	Teresa Scassa	Catherine Régis
Florian Martin-Bariteau	Elissa Strome	Rosario Cartagena
Lorian Hardcastle	Pascal Thibeault	Jodie Al-Mqbali
Nicole Davidson	Jennifer Gibson	Michael Brian Lang
Anna Goldenberg	Caroline Mercer	Genevieve Lavertu
Christian Blouin	Colleen M. Flood	Cindy Lu
Justine Gauthier	Bryan Thomas	Heba Roble
Michael Froomkin	Cécile Bensimon	Manik Saini
Sana Tonekaboni	Lindsay Thompson	
Ira Parghi		

[11] The following people participated in this workshop but were not part of a breakout group: Catherine Régis; Rosario Cartagena; Jodie Al-Mqbali; Michael Brian Lang; Genevieve Lavertu; Cindy Lu; Heba Roble; Manik Saini.

CIFAR

cifar.ca/ai