CIFAR

MACHINE MD:
Law and Ethics of Health-Related AI Case Study 5:

The Intelligent Powered Wheelchair

Workshop held: February 3, 2023

Report published: December 4, 2023

This report was drafted by Sophie Nunnelley in collaboration with the participants of the Machine MD: Law and Ethics of Health-Related AI Case Study 5: The Intelligent Powered Wheelchair.

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

S. Nunnelley, A. Goldenberg, C. Régis, C. M. Flood, T. Scassa, F. Ferland, K. Dahlia, A. Panchea, C. T. Sheldon, L. Hardcastle, and the workshop participants, *Machine MD: Law and Ethics of Health-Related A.I. Case Study 5: The Intelligent Powered Wheelchair* (Toronto: CIFAR, 2023).

Table of Contents

Law and Ethics Case Studies in Health-Related AI

The CIHR-funded *Machine MD: How Should We Regulate AI in Health Care?* project is led by Colleen M Flood (Dean of Law, Queen's University); Anna Goldenberg (Senior Scientist, SickKids); Catherine Régis (Law, Université de Montréal); and Teresa Scassa (Law, University of Ottawa). The project is dedicated to investigating the legal and ethical issues raised by artificial intelligence (AI) in health care and to developing recommendations for their optimal governance.

Part of the Machine MD team's work includes examining real AI technologies, the practical issues they raise, and their current treatment in Canadian and foreign law. This approach moves 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.

In keeping with this aim, the Machine MD team has partnered with CIFAR to host a series of online case study events.[1] Each event assembles an interdisciplinary group of experts in AI, law, ethics, policy, and medicine to discuss the regulatory issues raised by a specific AI technology. The previous four case studies took place in the spring and fall of 2022, and concerned the "OR Black Box", the Suicide Artificial Intelligence Prediction Heuristic or "SAIPH", "digital twins" technology, and pediatric cardiac arrest prediction.[2] This report summarizes the findings of the fifth case study in the series, concerning an Intelligent Powered Wheelchair.

[1] The current series of 5 case studies built on a previous collaboration between CIFAR and the Machine MD team. See: "AI & Health Care: A Fusion of Law & Science (Part One): An Introduction to the Issues" (February 2021) and "AI & Health Care: A Fusion of Law & Science (Part Two): Regulation of medical devices with AI" (May 2021), online: <https://cifar.ca/ai/ai-and-society/>.
[2] Previous reports can be found online: <https://cifar.ca/ai/ai-and-society/publications/> and <https://www.ottawahealthlaw.ca/projects/machine-m-d>.

Case Study #5: The Intelligent Powered Wheelchair

The Intelligent Powered Wheelchair (IPW) integrates autonomous robotics and AI technologies to enhance the experience of wheelchair users. Focusing on users with motor impairments, the developers are exploring touchscreen-based input devices, autonomous navigation and path planning, and non-invasive Brain-Computer Interfaces. The project promises to significantly benefit wheelchair users; however, it also raises important legal and ethical questions. This event brought together approximately 36 experts, with a broad range of expert backgrounds and perspectives (e.g., in law, healthcare, health policy, and artificial intelligence and innovation) to examine and discuss these issues and the range of possible regulatory responses.

Presentation by François Ferland (Université de Sherbrooke), Dahlia Kairy (Université de Montréal), and Adina Panchea (Université de Sherbrooke)

The workshop began with a presentation from François Ferland (Université de Sherbrooke), one of the technology's innovators, with assistance from two colleagues on the project, Dahlia Kairy (Université de Montréal) and Adina Panchea (Université de Sherbrooke). Ferland explained that this project has seen the wheelchair's evolution from manual, to semi-autonomous, to autonomous. The project began in approximately 2006 when a team, under Joelle Pineau's leadership at the McGill Centre for Intelligent Machines, took an ordinary power wheelchair and sought to enable it to semi-autonomously navigate large physical environments.[3] For instance, in 2013 they tested the wheelchair's ability to navigate between stores in a Montreal shopping mall.[4] This *semi*-autonomous wheelchair was the first prototype.

In 2018 the project was moved to the IntRoLab at the Université de Sherbrooke, where the research team has been pursuing several advancements that offer new ways for users to interact with the wheelchair and enhance the IPW's abilities to *autonomously* navigate. Ferland explained that in its current form, the IPW can navigate autonomously, and with user input, by processing data from several sources: (i) An odometry sensor that detects wheel velocity to precisely measure speed and distance travelled; (ii) two 2D light detection and ranging (LiDAR) sensors that detect and measure objects on a single plane in a 360 degree range around the wheelchair; (iii) two forward facing RGB-D cameras that help with depth readings and obstacle detection[5]; and (iv) a physical joystick and touchscreen that capture user input. Together, these allow the IPW to autonomously detect obstacles in the wheelchair's path. The touchscreens make the interface

[3] J Pineau, R West, A Atrash, J Villemure & F Routhier, "On the Feasibility of Using a Standardized Test for Evaluating a Speech-Controlled Smart Wheelchair" (2011) 16 Int J Intell Control Syst 124.
[4] D Kahiry et al., "Users' perspectives of intelligent power wheelchair use for social participation" (Conference Proceeding RESNA 36th International Conference on Technology and Disability, Bellevue, 2013).
[5] See Jianwei Li et al., "High-quality indoor scene 3D reconstruction with RGB-D cameras: A brief review" (2022) 8:3 Computational Visual Media 369.

more comfortable for users with disabilities that make it difficult to use manual joysticks, e.g., because of fatigue or difficulties with motor control.

Ferland explained the team is also exploring additional technologies that would improve user interface with the IPW as well as its autonomous navigation capabilities. One technology under exploration is a “brain computer interface”. They developed a non-intrusive helmet with electroencephalogram (EEG) sensors, to see if the measured brain activity could modulate the commands already being given to the wheelchair, e.g., through the touchscreen. In testing they were able to associate different frequencies in brain activities with commands like “forward” or “backwards”. However, they were not successful in using the data to accurately command the wheelchair. Ferland explained they are now exploring other uses of this data that incorporate machine learning. These could involve training an AI model on data from a particular user, e.g., to slow the chair when the person’s brain activity indicates a rise in stress or fear.

Another possible innovation discussed by the team would use AI to make the IPW’s navigation more “human aware”, taking into account the social environment. A version of this technology can already be found in service robotics, where machines autonomously navigate inside an environment (usually a home), to detect people and their interactions. Ferland explained that human aware navigation would build on this technology, improving its ability to detect social nuance. For instance, where people are facing each other, the model would detect a possible social interaction (e.g., people talking) and navigate around them. The developers could also integrate other capabilities, e.g., to direct the IPW to “meet person A”, where the wheelchair would detect that individual and autonomously orient itself in front of the person.

Ferland discussed the data collection and storage required for the IPW in its current and possible future forms, as well as some possible risks, if proper data and privacy safeguards are not maintained. He described the wheelchair’s current data output as an occupancy grid – essentially a 2D floor map built from information from the LiDAR sensors and cameras through a process called Simultaneous Localization and Mapping (SLAM). User input is not collected (it is processed live) and the wheelchair does not currently employ a global positioning system (GPS). Additional information would also be collected through the new technologies under development. For instance, the proposed “brain computer interface” would require EEG data to train the wheelchair

for a particular user, and “human aware” navigation would use cameras capable of detecting a person’s identity, position, and orientation.

Ferland explained that, if improperly managed, this data could pose risks to privacy. For e.g., the raw sensor and camera data could be used to reproduce an accurate 3D model of any environment and its occupants. The data could be analyzed for patterns, e.g., relating to the wheelchair user’s movements within the home and, were GPS capabilities added, this would also be true outside the home. The EEG brain activity data collected by the brain-computer interface, while low quality, could be used to identify a person’s emotional responses to events, which might be of interest, e.g., to advertisers. Moreover, the “human aware” navigation module would collect information capable of identifying individuals and their movements.

Yet, Ferland emphasized, their use of this data is subject to safeguards. Most notably, the operation of the IPW itself does not require the transmission of any information out of the device – data processing is done locally. While some information, such as the training of the EEG inference model and manual adjustments to the floor map must be processed outside the IPW, this can be done offline. Program updates have sometimes been run remotely, however, this is not essential to the IPWs operation, and updates could be hardwired. They are also planning additional safeguards; for instance, the “human aware” navigation module will refer to people anonymously (e.g. “person A”). In response to questions about the device’s potential mapping of public or sensitive spaces, Ferland suggested the mapping function could potentially be turned off, so that it would rely only on ‘live’ processing to manoeuvre around obstacles, without recording any sensitive data.

In closing, Ferland emphasized that the IPW has the *potential* to be very intrusive, but that those risks can be mitigated. At the same time, he noted that this research is a “thought experiment on how these technologies could be exploited in a commercial project”; should the IPW be developed commercially, careful information management practices will be essential. Indeed, his colleague Kairy emphasized that at the research phase they are primarily concerned with what *can* be done. At the commercialization stage it will be necessary to reflect on what should be done, considering privacy among other considerations.

Commentaries

Following the innovators' presentation, the group heard from three legal experts. The legal scholars provided overviews of three legal issues that frequently arise with health-AI— privacy, informed consent, and liability – analysing their potential applications to the IPW and discussing potential questions and concerns.

A. Privacy (Teresa Scassa, University of Ottawa)

Professor Teresa Scassa opened the legal discussion with a presentation on privacy law. She began by explaining that, as a general matter, the AI context requires consideration of at least two kinds of data. One is *the data used to develop and train AI*; we must ask, Scassa explained, what are these data, were they sourced in a way that takes privacy into account, and how can we maintain the privacy and security of those data, e.g., ensuring there are no data breaches or attempts at reidentification? Second, Scassa explained, we must consider how privacy and security are addressed in the collection, use and sharing of *data from the technology's users*. For instance, we might ask whether the data are processed and stored locally (on the device) or remotely (in the cloud). She noted that the IPW developers have apparently taken a privacy protective approach by choosing to process most of the data locally.

Additional privacy issues come up with particular technologies, Scassa explained, for instance, where they involve some tracking or surveillance, intrusion into private spaces, or recording or monitoring (e.g., video capture of assistants or family members). These privacy risks can arise from the technology itself (e.g., where it includes tracking by design), or where it is placed on external, multi-purpose devices like tablets, which have location data enabled. Whether these uses of personal information are problematic will often depend on how the data is used and stored. More novel privacy issues may also arise with certain technologies, such as with the IPW's proposed "brain computer interface". Scassa noted that these should be closely considered for their privacy implications.

Scassa went on to discuss a specific area of concern – the risks that arise from the transition from university or hospital-based research and development to the commercialization phase. She

explained that AI training data, and data from a technology's users, are often gathered in university or hospital-based studies with research ethics boards (REBs) approval. This REB approval may require risk assessment, adequate informed consent, data sharing agreements, and data anonymization, creating a safe space for research. Moreover, REB approval is often seen as a proxy for legal compliance. Yet, many AI projects will go on to a commercialization stage, potentially raising unanticipated privacy issues. For instance, commercialization may involve the combining of technologies from different sources and / or the addition of new features, raising new or different privacy issues. Indeed, the collection and sharing of personal data may be part of the business model for the technology where, for e.g., the selling of data will provide a revenue stream. In other words, while research and development projects may be privacy compliant, those projects can change in nature when they are commercialized.

Scassa emphasized that this transition from the research to commercial phase is challenging and requires careful consideration. Some questions to consider are: Should we be incorporating *user* concerns about privacy and data security at the research and development stage? And second, *how* might early consideration of privacy impact both development and commercialization strategies? Scassa suggested that the careful incorporation of user privacy concerns might have real effects on who can commercialize or how that process takes place.

B. Informed Consent (Tess C. Sheldon, University of Windsor)

The group learned about informed consent principles from Professor Tess C. Sheldon. Sheldon emphasized that while emergent assistive technologies offer promise, e.g., in their potential to enhance personal mobility, give respite to exhausted caregivers, and help people to stay at home, there are important questions about its appropriate regulation. One such question is whether existing informed consent laws are sufficiently robust in the context of persons with disabilities using AI.

Sheldon began by setting out some important background considerations and principles. First, she emphasized that these are disability justice questions; emergent assistive devices are more likely to be used by people with disabilities and the absence of sufficient legislative and policy safeguards puts them at particular risk. She noted that the IPW might not be accessible to all persons with disabilities; some people face barriers, e.g., relating to dexterity, sensory awareness, vision, memory, or attention, that could make the technology difficult to use. The technology could

also pose risks to some persons with disabilities, for instance, that their movements will be "tracked" (e.g., the last time they went to the fridge), leading to unwanted intervention. Sheldon considered that these kinds of considerations lie in the background of discussions of informed consent. Indeed, she emphasized that surveillance of persons with disabilities is already taking place, e.g., where mental health information is shared with border control officers, or through medical devices that offer digital monitoring of medication compliance.[6] Even where such tools are framed as "optional", Sheldon explained, "consent" can be significantly undermined by power imbalance.

Sheldon also surfaced competing models of disability and views about the role of technologies in the lives of persons with disabilities. Since the 19th century, she explained, disability has been associated with abnormality; people with disabilities have encountered pity and mourning, and have had to change themselves, e.g., through "cures" and prosthetics, to appear "normal". AI and other technology can align problematically with this view, fitting with a medical model of "solutionism" – the idea that disability is a problem that must be "fixed". She explained that a Critical Disability Theory lens shows us that technology sometimes exacerbates and obscures inequalities. A social model of disability, on the other hand, tells us humans are universally variable – there are no aberrations – and it is societal structures, not the body or brain, that are disabling. A focus on environment also forces us to consider, Sheldon emphasized, how the built environment, including many long-term care institutions, remain wheelchair inaccessible. She urged us question the point of technology like the IPW if we don't have the ramps, curb cuts, automatic door openers, and so on, which are necessary to make environments suitable for wheelchair use.

Sheldon then turned to the law of capacity and informed consent. In Ontario these are governed by a complex statutory framework centered around two statutes, the *Substitute Decisions Act* (SDA) and *Health Care Consent Act* (HCCA).[7] At the core of these is the principle that there is no treatment without consent. Sheldon suggested that intelligent wheelchairs are a therapeutic or medical device and, as such, their use requires informed consent. The consent must come from

[6] Sheldon offered the example of Abilify, a psychiatric medication that embeds a trackable sensor inside an ingestible pill. See e.g., Susan Scutti, "FDA approves pill with digital tracking device you swallow", *CNN Health* (15 November 2017) online: <https://www.cnn.com/2017/11/14/health/fda-digital-pill-abilify/index.html>.
[7] *Substitute Decisions Act*, 1992, SO 1992, c 30; *Health Care Consent Act*, 1996, SO 1996, c 2, Sch A (HCCA).

either the person themselves or, where they are deemed to lack legal capacity, from a substitute decision-maker.

Yet, Sheldon noted that applying these capacity and consent principles to the IPW can bring to light some key challenges and questions. For instance, what information is required for consent to be "informed"? Sheldon took the position that users of AI-powered wheelchairs (or their substitute decision-makers) would need information not just about the device's potential risks, benefits, and possible alternatives as they relate to "treatment", but also risks relating to data security and privacy. She noted, moreover, that such information could be difficult to obtain and convey in the case of emergent assistive technologies, where even healthcare providers may not understand the underlying AI, or there might be limited evidence regarding risk, especially as it relates to persons with disabilities.

Substitute decision-makers might also face challenges, Sheldon explained. When a person in Ontario is found to lack legal capacity to make a medical decision, the substitute decision-maker must follow a "substitute judgement" approach that involves putting themselves in the person's shoes and following the person's values, beliefs, and prior capable wishes.[8] Yet, it may be impossible for a substitute decision-maker to know a person's values and beliefs regarding a new technology. A similar challenge applies in the case of advance directives; people cannot express their future wishes regarding a technology that does not yet exist. Sheldon suggested these challenges may complicate the process of obtaining meaningful informed consent to IPW-use.

Sheldon also troubled the legal capacity principles that determine when we can make our own medical decisions. She noted that historically, legal capacity was treated as a single state or condition, and persons with disabilities were often presumed to lack capacity. This changed in the 1980's and 1990's with the reforms that led to the *SDA* and *HCCA*. Legal capacity is now understood to be a legal (not clinical) construct. Under the law, everyone is presumed to have legal capacity, and capacity can fluctuate across time and context. Importantly, this standard does not require compliance with treatment recommendations, nor does it turn on perceived best interests.[9] The legal capacity test generally requires that the person be able to "understand" the

[8] HCCA, *ibid*, s 21.
[9] *Starson v Swayze*, 2003 SCC 32, at paras 79, 91.

information relevant to the decision and to "appreciate" the reasonably foreseeable consequences of the decision or lack of decision.[10]

Yet, Sheldon emphasized, many people who appear to lack capacity *can* make decisions with supports, such as clearer information, a support person, or a more comfortable environment. The Ontario Human Rights Commission has said that "[b]efore determining that a person lacks capacity, an organization, assessment body, evaluator, etc., has a duty to explore accommodation options to the point of undue hardship".[11] And the UN Special Rapporteur on the Rights of Persons with Disabilities has discussed the importance of ensuring legal capacity and decision-making supports in the in the AI context specifically.[12] Sheldon emphasized that at the core of these statements is a principle of inclusion; every stage of AI development, regulation, and use, must take place with the meaningful involvement and consent of persons with disabilities.

C. Liability (Lorian Hardcastle, University of Calgary)

Professor Lorian Hardcastle took up our third legal issue, liability. She began by setting out the four conditions that must be met in any successful negligence claim. First, the person must owe a legal duty of care to the patient who suffered harm. Hardcastle noted this is usually easy to establish in healthcare as health professionals, device developers, and manufacturers, all owe duties of care to patients. The second requirement – a violation of the standard of care – is harder to meet. The key question here is whether the health professional or other individual met the standard we would reasonably expect such a professional to meet. A third requirement is a loss or injury; "near misses" are not grounds for liability. In the context of the IPW, Hardcastle explained, the injury could be to the wheelchair user (e.g., where they ended up in traffic) or to a third party (e.g., collision with a pedestrian). A person could also suffer moral injuries relating, for instance, to a lack of adequate consent to use the technology, or to being wrongfully excluded from using the device. The fourth requirement for a successful negligence claim is that the injury is causally linked to the professional's failure to meet the standard of care.

---

[10] See e.g. the Ontario *HCCA*, *supra* note 7, s 4; *Starson, ibid*.

[11] Ontario Human Rights Commission, *Policy on preventing discrimination based on mental health disabilities and addictions*, "16. Consent and capacity" (Toronto, 2014), online: <https://www.ohrc.on.ca/en/policy-preventing-discrimination-based-mental-health-disabilities-and-addictions/16-consent-and-capacity>.

[12] *Artificial intelligence and the rights of persons with disabilities, Report of the Special Rapporteur on the rights of persons with disabilities*, A/HRC/49/52 (28 December 2021), online: <https://www.ohchr.org/en/documents/thematic-reports/ahrc4952-artificial-intelligence-and-rights-persons-disabilities-report>.

Hardcastle also considered *who* might face liability claims relating to the IPW, discussing four groups. First, those who develop tools for assessing and training IPW users, as well as health professionals like occupational and physical therapists who apply those tools, could face claims. However, Hardcastle suggested their risk of liability is likely low. For guideline developers, there is a general understanding in the literature that guidelines do not precisely govern conduct and health professionals remain responsible for their application in practice. Health professionals, in turn, can often point to guidelines compliance as evidence of having met the standard of care. At the same time, Hardcastle noted, there could be cases where a particular application of clinical guidelines will fail to meet the standard of care.

Second, developers and manufacturers could face allegations of negligence. For instance, where the IPW is a retrofit of a traditional wheelchair, an injured person might claim the "wrong" base model was chosen (e.g., too old, not technically compatible), leading to harm. The developers of AI algorithms could also face claims but, Hardcastle noted, negligence law is difficult to apply in this context; where medical devices are typically assessed for safety at a single moment in time, AI algorithms may evolve. A good precautionary rule of thumb for AI developers, Hardcastle noted, is to build in as many redundant safety features as possible.

A third group discussed by Hardcastle was users and their caregivers. Wheelchair users could theoretically bear some liability for harm through the legal principle of "contributory negligence", e.g., if they used an override feature to ignore a safety warning, or used the device in a way that was contraindicated). Caregivers, on the other hand, are unlikely to face significant liability risk as they don't have the same duties and standards of care as, say, health professionals. Finally, Hardcastle discussed potential claims against government regulators for failure to adequately regulate the technology, noting that courts have generally been unreceptive to these kinds of claims, e.g., in cases relating to jaw implants and breast implants.[13] She explained courts typically view regulators as too distant from the injuries suffered and that they worry about opening the "floodgates" of claims against government.

[13] See *Attis v Canada (Health)*, 2008 ONCA 660, leave to appeal to SCC denied, *S Joyce Attis and A Tesluk v Her Majesty the Queen in Right of Canada as represented by the Minister of Health, Attorney General of Canada, Regulatory Institution 1, Regulatory Institution 2, John Doe and Jane Doe and Dow Corning Corporation; Drady v Canada (Health)*, 2008 ONCA 659.

In conclusion, Hardcastle emphasized that the law rightly requires reasonableness, not perfection. Requiring perfection could lead to unintended consequences, e.g., making developers overly risk averse and leading to the withdrawal of assistive technologies from persons (e.g., those with more significant impairments) who might benefit from it most. The key principle should be one of balance between ensuring beneficial access and protecting against undue risk.

Breakout Sessions

The presentations were followed by breakout sessions on each of privacy, informed consent, and liability. The purpose of these was to allow members to deeply engage with the issues in small interdisciplinary groups, allowing for rich discussion of challenges and solutions. At the conclusion of the sessions a rapporteur from each group summarized their discussion and findings for the full group. Some of their core thematic concerns are set out below.

Breakout #1: Informed Consent

Attendees: Ian Stedman (group rapporteur), Nicole Davidson (scribe), Sophie Nunnelley, Jason Millar, Adina Panchea, Catherine Frazee, Jake Okechukwu Effoduh, Megan Fultz, Michael Froomkin, Tess Sheldon

The participants of the informed consent breakout group began by reflecting on the importance of centering wheelchair users at every stage, including development. As an overarching matter, they cautioned against “over-medicalizing” the IPW. They emphasized that this is an essential mobility device and calling it “medical device” (which triggers a range of legal requirements) could hinder access.

The remainder of the group’s discussion troubled and explored the requirements of informed consent (i) in research settings, and (ii) at the commercial use stage. On the first, participants noted that informed consent requirements are more stringent at the research phase, due to the role of ethics boards / Institutional Review Boards, with a range of implications. On the one hand, these rightly aim to protect the most vulnerable people. On the other, these standards can result in some people being excluded from research, creating problematic gaps in data. The group recognized the challenges in this tension. They noted, for instance, that it can be challenging for

researchers to navigate consent with users with cognitive impairment, and research ethics boards are conservative, not wanting to assume any risk. At the same time, people who are perceived to lack legal capacity should be included in research, and they can often participate with the right supports.

The group also discussed the benefits of a proportional approach that looks at the *actual* risks and benefits of participating in a study and adjusts the informed consent process accordingly. They wondered, moreover, whether we focus too much on perceived risk in the AI context (wanting to know, e.g., exactly how machine learning works and exactly what the camera will see). They considered that in a world with Roomba vacuums, smartphones, and other smart devices that frequently record us, we might need a more nuanced idea of "risk". For instance, rather than ask what *is* being recorded, we could perhaps ask what should *not* be recorded, and why? The group also noted the potential value in partnering with patient and wheelchair user communities in determining how best to understand informed consent in this context. Participants also acknowledged the importance of continuous consent to research, especially where participant data is being used for multiple purposes or being shared with third parties.

This group also discussed the less stringent informed consent requirements that apply when products move from the research stage to open market, describing this consent as closer to implied. They noted that people frequently purchase technology and agree to terms and conditions that they have not read. Participants considered that what is required for consent to use the IPW will depend on how the technology develops, for instance, whether it starts collecting and transmitting data, bringing it under privacy legislation. It also turns on whether we characterize the IPW as a medical device – prescribed by a health professional – or as a mobility device that can be purchased on the open market. One participant raised the implications of a user bringing the IPW to another jurisdiction that permits more surveillance and data collection, and wondered whether those functions could be limited to protect users in those places. The group also highlighted the need for discussion about how to better support users deciding about IPW use, rather than move to substitute decision-making on their behalf.

Breakout #2: Privacy

Attendees: Elaine Gibson (rapporteur), Saly Sadek (scribe), Geneviève Lavertu, Sarah Grieve, Teresa Scassa, Colleen Flood, Justine Gauthier, Dahlia Kairy, Rona Fleming

The members of the privacy breakout group began by talking about the mapping of sensitive private and commercial spaces, such as homes, changerooms, banks, and airports. Like the informed consent breakout group, the group drew a comparison to the Roomba vacuum, noting that concerns have been raised about that device's collection of personal information that could, in theory, be sold.[14] Members considered the importance of having an easy way to disconnect the mapping functionality in these environments. However, they also raised an important point about discrimination; an IPW user might have disabled recording in, say, a bank, but bank personnel might still refuse access, not trusting that the mapping technology is off.

Members of this group also discussed who should have the responsibility to ensure appropriate data and privacy protections, and to communication those protections. For instance, they considered that where a health professional is recommending use of the IPW, potential users might erroneously believe they can trust that their information will be appropriately managed. This risk is perhaps compounded by the tendency for patients to think positively about new devices, downplaying the risks. Yet information can easily get out of hand. The group suggested that 'getting this right' should not be the sole responsibility of individual healthcare providers and wheelchair users navigating informed consent. They considered that regulation may be required to set out clear standards and responsibilities relating to data and privacy, especially on the part of developers. One member offered the caveat that patients already access information from various sources (e.g., the internet) and we should not assume naivety. The group agreed a balance must be struck between allowing some individual autonomy, including the voluntary assumption of risk, while also protecting against undue harm.

[14] Maggie Astor, "Your Roomba May Be Mapping Your Home, Collecting Data That Could Be Shared", *The New York Times* (25 July 2017), online: <https://www.nytimes.com/2017/07/25/technology/roomba-irobot-data-privacy.html>.

Breakout #3: Liability

Attendees: Lara Khoury (rapporteur), Michelle Rodrigues (scribe), Lorian Hardcastle, François Ferland, Lindsay Thompson, Bryan Thomas, Tanya Horsley, Catherine Régis, Cécile Bensimon, Jennifer Chandler.

A broad theme taken up by the liability breakout group was the tolerability of some risk. Group members discussed how aspiring to perfect safety can delay access to beneficial technology. They noted that this issue intersects with informed consent, as users may voluntarily accept certain risks in order to benefit from IPW technology. Indeed, healthcare routinely involves such trade-offs, where patients give informed consent to undergo even risky procedures and treatments. Like members of other groups, they also discussed the risks of overmedicalizing the IPW, asking: Is this a medical device, a mode of transportation, or a hybrid of the two?

The group also discussed the responsibility for ensuring adequate information regarding a technology's risks and benefits. They emphasized that developers must be responsible for ensuring sufficient information is conveyed to users and their care teams, along with the value of having wheelchair users involved to inform both the understanding of risk, and its effective communication to users and their caregivers. One participant also noted the value in having an interdisciplinary support team, including a person with technological expertise, around the wheelchair user to help translate technical information.

Members of this group also talked about the allocation of responsibility in cases of user or third-party harm. They considered that where insurance is required for, say, self-driving cars, it is unclear what role insurance would play in the IPW context, or who would pay for any insurance. More broadly, they discussed the critical responsibility of regulators, and the need for regulation that is appropriately protective, yet not "overzealous" such that it impedes access to beneficial technology.

Discussion

As in our past case study events, some noteworthy themes emerged across the discussions and analyses. We set out some of these below.

A. How we characterize a technology has important implications for regulation and access

The importance of preliminary categorization choices – to call a device "medical" nor not – emerged as a significant theme throughout the workshop. Participants raised these questions at numerous junctures asking, as AI and other technologies are increasingly integrated into our lives, what makes the IPW "medical"? They discussed a range of possible analogies, comparing the IPW to driverless cars, Roomba vacuums, motorized scooters, and even bicycles, recognizing the significant implications of these characterization decisions.

The general tension and question that emerged was, what categorization will best respect patient autonomy and choice while also protecting against undue risk? Some cautioned that "over medicalizing" the IPW will impede access to an essential mobility device by requiring potential users to go through a healthcare provider. Yet others noted that medical devices are sometimes accessible under public funding programs; they worried that putting the IPW on the open market could widen existing gaps in who can access emergent technologies. Related questions relate, for instance, to whether we should require users to have liability insurance (as we do with motor vehicles), and at whose expense?

Participants also wrestled with other potential implications of a de-medicalized, direct-to-consumer process that does not keep a health professional "in the loop". For instance, they asked, without healthcare's legal requirement of informed consent, how would we ensure that IPW users are fully informed of the devices benefits and risks? That responsibility would presumably shift to developers and device providers, however, what regulation would best ensure that disclosure is effective and meaningful? Yet, having health professionals as gatekeepers could pose other risks – e.g., that their involvement will erroneously be seen to imply a trusted process for data and privacy protection. There is a confusing blurriness, in some health AI contexts, between the

information regarding a *treatment's* risks and benefits (which clinicians must disclose to obtain informed consent to healthcare) and information about the *data security and privacy risks* that might flow from health AI use. Where these risks intertwine, they may be more difficult to address through our traditional legal structures.

B. Technology users should be involved at every stage of development and discussion

A second prominent theme was the importance of user involvement at every stage of health AI development and regulation. Participants noted that involving users early, at the development stage, could shape the very understanding of "risk", as users are best placed to identify which wheelchair features feel risky, especially as public perceptions of technology can change rapidly.[15] For instance, a member of the innovation team noted that some wheelchair users have expressed concerns that overreliance on IPW technology will lead to a loss of ability. Having user perspectives integrated early can help to ensure responsive design and regulation.

Early and consistent user involvement could also help with the appropriate management and communication of risks. Participants emphasized, for instance, that wheelchair users should have "a say" in balancing a technology's risks and benefits, and in determining what information is required to do that balancing. For example, they queried whether meaningful informed consent requires that users be informed about the nature of AI or machine learning. While this issue was not resolved, the essential point, repeatedly raised in discussions, is that users must be included in these discussions and decisions, from the research and development stages, through to commercialization and regulation.

C. The risks of commercialization warrant early and careful consideration

A third key theme was the importance of thinking about regulation in an integrated way, from research and development, through to the commercialization and consumer stages. Participants worried that researcher concerns and priorities – e.g., for the protection of personal information – could get lost at the commercialization stage where data can be monetized. They questioned, for instance, how realistic it is to assume that commercial providers will maintain rigorous safeguards on data collection and use. This came up especially in discussion of privacy.

[15] For instance, Kairy explained that in early consultations, wheelchair users were asked how they would feel about a wheelchair moving independently – e.g., would it be feared? Yet now, as technologies like self-driving cars have become better known, this is unlikely to be a concern.

However, it was also raised in conversations about safety – e.g., commercial providers might be motivated to collect and share sensitive personal information relating to user accidents or device error to improve device safety and quality.

This theme also intersects with the previous two, relating to device categorization and user involvement. If devices like the IPW are deemed private consumer devices rather than medical ones, it arguably renders more urgent the need for regulatory safeguards. Indeed, a participant aptly noted that even if the IPW itself stays off the open market, we might see the commercial production and sale of third-party add-on devices, which people could put on their wheelchairs. Where technologies like smart watches that measure heartrates are already blurring the line between private consumer and medical devices, it suggests both researchers / developers and regulators should be considering these commercialization activities from the outset.

This focus on the commercial context also highlights, once more, the central importance of incorporating both legal and user perspectives early and often. Participants reflected on how changing the culture of lab research – to ensure routine but meaningful consideration of matters such as privacy, informed consent, bias, and user experience – could bear fruit at the commercial development stage. Where these matters are integrated early on – "socializing technologists into the ethical and legal context" as put by a participant – they can become meaningfully integrated into AI research and innovation. Moreover, participants emphasized that if such concerns become part of "how we build AI research", this will also influence industry and regulation.

Conclusion

In this workshop series we take up specific applications of health AI, recognizing the enormous diversity of AI technologies and the importance of particularity. The interdisciplinarity of our expert participants – from law, medicine, AI innovation, patient experience, and ethics, among other perspectives – adds another important layer of nuance. The form of conversation generated by these workshops is critical. At this moment in which leading AI figures are calling for a halt to AI development to allow for the apt regulation of risks, we require interdisciplinary discussion and collaboration. The five case studies in this series have provided opportunities for precisely such discussion, furthering the goal of responsive AI regulation.

The IPW considered at this fifth case study event represents a potentially important advancement in assistive technology. Yet, like the other four technologies in the series, it raises potential issues, e.g., relating to safety, privacy, and informed consent, which require careful consideration. Specific issues emerged in this case study, for instance, the difficulty and importance of determining whether a device is “medical”. Examining these issues with sufficient particularity helps to ensure responsive law reform. Looking at the case study series as a whole also helps by revealing cross-cutting issues. For instance, we see across all five case studies a concern that we cannot only address risks *ex ante*, after harms have occurred. There is an apparent need for proactive regulation to ensure we can collectively benefit from health AI’s tremendous potential while also being protected from undue risk. This case study, with the others in the series, offers guidance to AI researchers, developers, and lawmakers in that important task of regulatory reform.

Workshop Participants:

Cécile Bensimon
Marc Bilodeau
Jennifer Chandler
Nicole Davidson
François Ferland
Rona Fleming
Colleen Flood
Catherine Frazee
Michael Froomkin
Megan Fultz
Justine Gauthier
Elaine Gibson
Anna Goldenberg
Sarah Grieve
Pierre Guerette
Lorian Hardcastle
Tanya Horsley
Dahlia Kairy
Lara Khoury
Joanne Kim
Geneviève Lavertu
Ravi Malhotra
Caroline Mercer
Jason Millar
Sophie Nunnelley
Jake Okechukwu Effoduh
Adina Panchea
Catherine Régis
Michelle Rodrigues
Saly Sadek
Teresa Scassa
Tess Sheldon
Ian Stedman
Bryan Thomas
Lindsay Thompson
Mark Walker

CIFAR