



HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

Standing Committee on Access to Information, Privacy and Ethics

ETHI • NUMBER 157 • 1st SESSION • 42nd PARLIAMENT

EVIDENCE

Thursday, June 6, 2019

—
Chair

Mr. Bob Zimmer

Standing Committee on Access to Information, Privacy and Ethics

Thursday, June 6, 2019

• (1545)

[English]

The Chair (Mr. Bob Zimmer (Prince George—Peace River—Northern Rockies, CPC)): We'll call the meeting to order. This is the Standing Committee on Access to Information, Privacy and Ethics, meeting 157. Today's topic is the ethical aspects of artificial intelligence and algorithms.

We have with us today, as an individual, Brent Mittelstadt, research fellow, Oxford Internet Institute, University of Oxford, by teleconference. From the Canadian Association of Radiologists, we have Nicholas Neuheimer, chief executive officer; and An Tang, chair, artificial intelligence working group. From the Information Technology Association of Canada, we have André Leduc, vice-president, government relations and policy.

We do have some business to follow, so we're going to try to get this done as soon as we can. We're giving you a full hour. We just had votes. You have our apologies for that, but it's something out of our control.

We'll start off right away with Mr. Mittelstadt. Go ahead for 10 minutes.

Dr. Brent Mittelstadt (Research Fellow, Oxford Internet Institute, University of Oxford, As an Individual): Thank you so much for inviting me.

I have been researching the ethical challenges of algorithms and AI for nearly half a decade. What's become apparent to me in that time is that the promise of AI largely owes to its apparent capacity to replace or augment any type of human expertise. The fact that it's so malleable in that sense means that the technology inevitably becomes entangled in the ethical and political dimensions of the jobs, the practices and the organizations in which it's embedded. The ethical challenges of AI are effectively a microcosm of the political and ethical challenges that we face in society, so recognizing that and solving them is certainly no easy task.

I know, from witnesses in your previous sessions, that you've heard quite a bit about the challenges of AI, dealing with things such as accountability, bias, discrimination, fairness, transparency, privacy and numerous others. All those are extremely important and complex challenges that deserve your attention, and really the attention of policy-makers worldwide, but in my 10 minutes I want to focus less on the nature and extent of the ethical challenges of AI and more on the strategies and tools we have for solving them.

You've heard also quite a bit about the tools available to address these ethical challenges, using things such as algorithmic and social scientific auditing, multidisciplinary research, public-private partnerships, and participatory design processes and regulations. All of those sorts of solutions are essential, but my concern is that we're perhaps broadly using the wrong strategy or at least an incomplete strategy for the ethical and legal governance of AI. As a result, we may be expecting too much from our current efforts to ensure AI is developed and used in an ethically acceptable manner.

In the rest of my statement, what I want to address are the significant shortcomings that I see in current efforts to govern AI, specifically through data protection and privacy law on the one hand and through principled self-governance on the other. My principal concern here is that these strategies too often conceive of the ethical challenges of AI in an individualistic sense, when in fact they are collective challenges that require collective solutions.

To start with data protection and privacy law, responsibility far too often falls on the shoulders of individuals to protect their vital interests, or their privacy, autonomy, reputation and those sorts of things. Data protection law too often ends up protecting data rather than the people the data represents. That shortcoming can be seen in several areas of law globally. The core concepts of data protection and privacy law—personal data, personally identifiable information and so forth—are typically defined in relation to an identifiable individual, which means that the data must be able to be linked to an individual in order to fall within agreement of the law and thus to be protected by the law.

The emphasis on the individual is really mismatched with capabilities of AI. We're excited by AI precisely because of its ability to find small patterns between people and group them in meaningful ways, and to create generalizable knowledge from individual records or individual data. In the modern data analytics that drive so many of the technologies we think of as AI, the individual doesn't really matter. AI is interested not in what makes a person uniquely identifiable but rather in what makes that person similar to other people. AI has transformed privacy from an individual concern to a collective challenge, yet relatively little attention is actually paid in existing legal frameworks to collective or group aspects of privacy. I see that as something that really needs to change.

That shortcoming itself extends to the sorts of legal protections that we quite often see in data protection and privacy law that are offered to individuals and to their data. These protections are still fundamentally based on the idea that individuals can make informed decisions about how they produce data, how that data is collected and used, and when it should not be used. The burden is really placed on individuals to be well informed and to make a meaningful choice about how their data is collected and used.

As is suggested by the name, informed consent only works if meaningful, well-informed choice is actually possible. Again, we're excited about AI precisely because it can process so much data so quickly, because it can identify novel and unintuitive patterns within the data and because it can produce knowledge from them. We're excited because the data analytics that drive AI are so big, so fast and so unpredictable, but the voracious appetite that AI has for personal data, combined with the seemingly limitless and unpredictable reusability of the data, means that even if you're a particularly motivated individual, a well-informed choice about how your data is collected and used is typically impossible. Under those conditions, consent no longer offers meaningful protection or allows individuals to control how their data is collected and used.

• (1550)

Moving forward, in terms of data protection and privacy law in particular, we need to think more about how to shift a fair share of the ethical responsibility to companies, public bodies and other sorts of collectives. Some of the ethical burden that's normally placed on individuals should be placed on these entities, requiring them, for example, to justify their data collection and processing before the fact, rather than leaving it up to individuals to proactively protect their own interests.

The second government strategy I want to address has seen unprecedented uptake globally. To date, no fewer than 63 public-private initiatives have formed to determine how to address the ethical challenges of AI. Seemingly every major AI company has been involved in one or more of these initiatives and has partnered with universities, civil society organizations, non-profits and other sorts of bodies. More often than not, these initiatives produce frameworks of high-level ethical principles, values or tenets meant to drive the development and usage of AI.

The strategy seems to be that the ethical challenges of AI are best addressed through a top-down approach, in which these high-level principles are translated into practical requirements that will act as a

guide for developers, users and regulators. The ethical challenges of AI are more often than not presented as problems to be solved through technical solutions and changes to the design process. The rationale seems to be that insufficient consideration of ethics leads to poor design decisions, which create systems that harm people and society.

These initiatives are essentially producing self-regulatory frameworks that are not yet binding, in any meaningful sense. It seems as though the blame for unethical AI tends to fall, again, on the individuals, or individual developers and researchers, who have somehow behaved badly, as opposed to any sort of collective failure of the institutions, businesses or other types of organizations driving development in the first place.

With that in mind, I'm not entirely sure why we assume that top-down principles and codes of ethics will actually make AI, and the organizations that create it and use it, more ethical or trustworthy. Using principles and ethics is nothing new. We have lots of well-established professions, such as medicine and law, that have used principles for a very long time to define their ethical values and responsibilities, and to govern the behaviour of the professionals and organizations that employ them.

If we can think of AI development as a profession, it very quickly becomes apparent that it lacks several characteristics necessary to make a principled approach actually work in practice.

In the first place, AI development lacks common aims and fiduciary duties to users and individuals. Take medicine as a counter example: AI development doesn't serve the public interest in the first instance, in the same sense. Developers don't have fiduciary duties toward their users or people affected by AI, because AI is quite often developed in a commercial environment where fiduciary duty is owed to the company's shareholders. As a result, you can have these principles that are intended to protect the interests of users and the public coming into conflict with commercial interests. It's not clear how those are going to be resolved in practice.

Second, AI development has a relatively short professional history and it lacks well-established and well-tested best practices. There are professional bodies for software engineering and codes of ethics, but because it's not a legally recognized or licensed profession, professional bodies exercise very little power over their members, in practice. The codes of ethics they do have tend to be more high-level and relatively brief in comparison to other professions.

The third characteristic that AI development is seemingly lacking is proven methods to translate these high-level principles into practical requirements. The methods we do have available tend to exist or have been tested only in academic environments and not in commercial environments. Moving from high-level principles to practical requirements is a very difficult process. The outputs we've seen from AI ethics initiatives thus far have almost universally relied on vague, contested concepts like fairness, dignity and accountability. There's very little offered in the way of practical guidance.

Disagreements over what those concepts mean only come out when the time comes to actually apply them. The huge amount of work we've seen to develop these top-down approaches to AI ethics have accomplished very little in practice. Most of the work remains to be done.

• (1555)

What I would conclude with is that essentially ethics is not meant to be easy or formulaic. Right now we too often think of ethics purely in terms of technical fixes or checklists or impact assessments, when really we should be looking for and celebrating these normative disagreements because they represent, essentially, taking ethical challenges seriously in the plurality of opinion that we should expect in democratic societies.

The difficult work that remains for us in AI ethics is to move from high-level principles down to practical requirements. It's really only in doing that and in supporting that sort of work that we'll really come to understand the ethical challenges of AI in practice.

Thank you, and I look forward to your questions later.

The Chair: Thank you, Mr. Mittelstadt.

Next up is Mr. Tang. Go ahead for 10 minutes.

Dr. An Tang (Chair, Artificial Intelligence Working Group, Canadian Association of Radiologists): Thank you, Mr. Chair.

[Translation]

Thank you, Mr. Chair and members of the Standing Committee on Access to Information, Privacy and Ethics, for giving me the opportunity to speak with you today about artificial intelligence in radiology, specifically in relation to ethical and legal issues in the implementation of this technology in medical imaging.

My name is Dr. An Tang and I am here representing the Canadian Association of Radiologists (CAR), as chair of the Artificial Intelligence Committee within the CAR.

[English]

The CAR AI working group is composed of more than 50 members who have a keen interest in technology advancement in radiology as it pertains to AI. The composition of this working group is varied, from predominantly radiologists to physicists, computer scientists and researchers. It also includes a philosopher specialized in the ethics of AI and an academic lawyer.

Under the CAR board of directors' leadership we have been entrusted with taking a global look at AI and the impact it will have on radiology and patient care in Canada.

I believe I speak for most of my colleagues in thinking that this is a good-news story and that AI can dramatically impact the way radiologists practise, in a positive way. Through the collection of data and simulation, using mathematical algorithms, we can help reduce wait times for patients, thus expediting diagnosis and positively affecting patient outcomes.

AI software analyzing medical images is becoming increasingly prevalent. Unlike early generations of AI software, which relied on expert knowledge to identify image features, machine learning techniques can automatically learn to recognize these features with the use of training datasets.

AI can be used for the purpose of detecting disease, establishing diagnosis and optimizing treatment selection. However, for this to be performed accurately, access to large quantities of medical data from patients will be required. This, of course, brings the privacy question into the equation. How do we collect this data while still guaranteeing we are collecting this information in an ethical way that protects the privacy of our patients?

Because of the transition from film to digital imaging that occurred two decades ago in radiology, and because of the availability of digital records for each imaging examination, radiology is well positioned to lead the development and implementation of AI and to manage associated ethical and legal challenges.

• (1600)

[Translation]

CAR believes that the benefits of AI can outweigh risks when institutional protocols and technical considerations are appropriately implemented to safeguard or remove the individually identifiable components of medical imaging data.

[English]

Technology advancements are occurring so quickly that they are outpacing current radiology procedures. We need to establish regulations pertaining to data collection and ownership to ensure that we are safeguarding patients and not infringing on ethical or privacy guidelines.

The CAR is advocating for the federal government to take a leadership role in the implementation of an ethical and legal framework for AI in Canada. Despite health care being a provincial priority, AI is a global issue. We feel the government is well positioned to lead the provinces in the regulation of the implementation of such a framework. Similar examples are the federal government's leadership in the national medical imaging equipment fund in the early 2000s.

The CAR can help, and the AI working group, under the CAR board's leadership, has published two white papers on AI, the first published in 2018 on AI in radiology, and a general overview of machine learning and implementation in radiology. This second paper, published in May 2019, focused on ethical and legal issues related to AI in radiology.

We have provided copies of the white papers, with our recommendations, for each of you. For the purpose of the discussion, I would like to highlight the more prevalent ones as they relate to the federal government's role in this capacity.

The first is the implementation of a public awareness campaign regarding consent and patient sharing of anonymized health data and harm reduction strategies. This information is essential for helping to identify disease and treatment for future AI applications.

Second is the general adoption of broad consent by default, with the right to opt out.

Third is developing a system for ensuring data security and anonymization of radiology data for secondary use, and implementing system standards to ensure that this criterion is being met.

Fourth, train radiology data custodians and establish clear guidelines for their role in the implementation of data sharing agreements for common AI-related scenarios and third parties.

The CAR has to work with the federal government and provincial ministries of health, including the Canadian Medical Protective Association, or CMPA, to develop guidelines for appropriate deployment of AI assistive tools in hospitals and clinics, while looking at minimizing harm and liability for malpractice for errors involving AI. We need to educate radiologists and other health care professions on the limitations of AI and reiterate the use of the tool in supplementing the work rather than replacing radiologists.

AI is not going away. Sharing medical data is a complex issue that balances individual privacy rights versus collective societal benefits. Given the potential of AI in helping to improve patient care and medical outcomes, I believe we will start to see a paradigm shift from patient's rights to near absolute data privacy through the sharing of anonymized data for the good of society.

• (1605)

[Translation]

We need to work together to implement a framework to ensure that we can move forward with this technology, while respecting the patient's anonymity and privacy. AI in healthcare is going to happen sooner or later; let's make sure it is implemented in an ethical way.

[English]

Thank you for your time. I'm happy to answer your questions in either French or English.

The Chair: Thank you for your testimony.

We'll move on next to Mr. Leduc. Go ahead for 10 minutes.

[Translation]

Mr. André Leduc (Vice-President, Government Relations and Policy, Information Technology Association of Canada): Thank you Mr. Chair and members of the committee.

It is a privilege to be here today to present the industry's perspective on behalf of the Information Technology Association of Canada. ITAC is the national voice of the telecommunications and Internet technology industry. We have more than 300 members, including more than 200 small and medium-sized businesses.

[English]

As already noted by the other speakers, there's a lot of promise and opportunity behind artificial intelligence to support economic growth and societal improvements, and the opportunities are seemingly boundless. From human mobility by automating vehicles to precision health care, many of our forthcoming solutions will be powered by artificial intelligence.

To realize the full benefits of artificial intelligence, we'll need to create systems that people trust. I've provided a brief outline of the slides that I'll present here today, including our industry's obligations and where our industry is already going; a call on our government to lead in terms of developing an ongoing dialogue via public-private partnership; the types of impacts that this will have on our workforce and the need for re-skilling, upskilling and training; and the recommendations in order to build trust in artificial intelligence.

Canada has been recognized as a global leader in artificial intelligence research and development. We are attracting global talent to universities across Canada to study in this field. We're already experiencing the benefits of AI in a number of fields, from start-ups and SMEs to larger global tech companies, all of which have developed AI systems to help solve businesses' or some of society's most pressing problems. Many others are using AI to improve supply chain efficiencies, to advance public services and to advance groundbreaking research. By leveraging large datasets, increased computing power and ingenuity, AI-driven solutions can address any number of societal or business problems, from precision or predictive health care to automated and connected vehicles improving human mobility and decreasing traffic, having an exponential impact on our environment.

AI systems need to leverage vast amounts of data. The availability of robust and representative data, often de-identified or anonymized, is required for building and improving AI and machine learning systems. We can't overstate this enough: Having access to broad and vast amounts of data is the key to advancing our artificial intelligence capabilities in Canada.

That said, the AI ecosystem is global. It's very competitive and it's multi-faceted. Our association welcomes a multi-stakeholder engagement approach to artificial intelligence, one that encourages Canada to bolster global engagement on AI policy to ensure we are all prospering from the potential benefits for our societies.

I'll note six key factors for the committee to consider.

First, traditional industries are already seizing and leading in AI opportunities. From oil and gas to mining, forestry and agriculture, they are embracing this technology to drive efficiencies and compete on a global scale. They are developing new services and new products based on the information being analyzed and leveraging artificial intelligence.

Second, AI is a journey. This isn't going to be an end state. This is going to be something that continues to evolve over the forthcoming decades.

Third, central to any economy's digital transformation is cultural transformation, and misinformation in this space will kill consumer and citizen trust in new technology and artificial intelligence.

Fourth, there will be workforce disruption, but based on historical factors, we believe new technologies including AI will create more job opportunities than it will kill.

- (1610)

Fifth, we need partnerships for workforce development, including the re-skilling and upskilling of existing people who may face disruption, based on their current roles.

Sixth, next-generation policies are needed. These are next-generation technologies. It's time for us to start thinking outside the box.

When I first joined government in 1999, one of the first jobs I had was working to support the development of PIPEDA. I was also one of the lead architects of Canada's anti-spam legislation. I did my master's thesis on why SMEs struggle to comply with CASL and PIPEDA, so I've been working on this for the better part of the last 17 or 18 years. Interestingly, we never foresaw the impact that data would have on the legislative frameworks we have today. We couldn't foresee, when developing PIPEDA or CASL, the types of data-driven businesses that have come our way to date.

Next, I want to talk about industry's obligation to promote responsible development and use of artificial intelligence.

First, we recognize our responsibility to integrate principles and values into the design of AI technologies, beyond compliance with existing laws. While the potential benefits to people in society are amazing, AI researchers, subject-matter experts and stakeholders should continue to spend a great deal of time working to ensure the responsible design and deployment of AI systems, including addressing safety and controllability mechanisms, the use of robust and representative data, enabling greater interpretability and recognizing that solutions must be tailored to the unique risks presented by the specific context in which a particular system operates.

Second, in terms of safety, security, controllability and reliability, we believe technologists have a responsibility to ensure the safe design of AI systems. Autonomous AI agents must treat the safety of users and third parties as a paramount concern, and AI technology should strive to reduce risks to humans. Furthermore, the development of autonomous AI systems must have safeguards to ensure controllability of the AI systems by humans, tailored to the specific context in which a particular system operates.

Third is robust and representative data, with a specific focus on mitigating bias. To promote the responsible use of data and to assure its integrity at every stage, industry has a responsibility to understand the parameters and characteristics of the data, to demonstrate the recognition of potentially harmful bias and to test for potential bias before and throughout the deployment of AI systems.

AI systems need to leverage large datasets. The availability of robust and representative data for building and improving AI and machine learning systems is of utmost importance.

By the way, this could be a significant competitive advantage for Canada. We have a globally representative population, including indigenous communities. It would be a wonderful target for medical testing and AI testing in the medical field.

In terms of interpretability, we should leverage public-private partnerships to find ways to better mitigate bias, inequity and other potential harms in automated decision-making systems. Our approach to finding such solutions should be tailored to the unique risks presented by the specific context in which a particular system operates.

Finally, the use of AI to make autonomous consequential decisions about people informed by, but often replacing, decisions made by humans has led to concerns about liability. Acknowledging existing legal and regulatory frameworks, our industry is committed to partnering with relevant stakeholders to form a reasonable accountability framework for all entities in the context of automated systems.

- (1615)

We believe we should leverage and build a public-private partnership that can expedite AI R and D, democratize access to data, prioritize diversity and inclusion and prepare our workforce for the jobs of the future. ITAC members also believe that we need to prioritize an effective and balanced liability regime via the continued engagement of multi-stakeholder expert groups. The right solution is only going to come from an open exchange with all actors in the AI supply chain.

If the value favours only certain incumbent entities, there's a risk of exacerbating existing wage income and wealth gaps. In this scenario, this isn't "us versus them", "private versus public". It's just "us". There should be increased partnership to explore how to develop a safer and more secure and trusted data-driven digital economy.

There is a concern that AI will result in job change, job loss and worker displacement. While these concerns may be understandable, it should be noted that most emerging AI technologies are designed to perform a specific task or to assist and augment a human's capacity rather than to replace a human employee. This type of augmented intelligence means that a portion—most likely not all—of an employee's job could be replaced or made easier by AI.

Leveraging AI to complete an employee's menial tasks is a way to increase their productivity by freeing up time to engage in customer service and interaction or more value-added job functions. Nevertheless, while the full impact of AI on jobs is not yet fully known in terms of both jobs created and jobs displaced, an ability to adapt to rapid technological change is critical. We should leverage traditional human-centred resources as well as career educational models, and newly developed AI technologies should assist in developing both the existing workforce and the future workforce to help Canadians navigate through career transitions.

The Chair: Mr. Leduc, you're two minutes over. We had 10 minutes for your presentation.

Mr. André Leduc: I'll just run quickly. I'll go to the next slide, in which you can see our recommendations onscreen around prioritizing Canada's competitiveness, promoting innovation and ethical AI practices, leveraging global standards, investing in AI R and D, and using a balanced and flexible regulatory approach. I think this creates an opportunity for us to marry privacy and cybersecurity.

In summary, many if not all of the uses of AI are going to rely on data—in certain circumstances personal data—and responsible use of that data is key. Burdensome regulation or reporting will limit the pace of AI innovations. We have to get this balance right. Industry will follow the key principles of responsible use of personal data in AI. We believe these principles are echoed in Canada's first-ever digital charter, which we support as a foundational framework for launching AI that is trustworthy, secure, ethical and safe for Canadians.

Thank you.

The Chair: Thank you, Mr. Leduc. I apologize for constraining AI into 10 minutes. It's more than challenging, I can imagine. You did a pretty good job.

Anyway, now we have Mr. Saini for seven minutes.

Mr. Raj Saini (Kitchener Centre, Lib.): Good afternoon, gentlemen. Thank you very much for coming here today.

Mr. Tang, I'm going to start with you, because very rarely do we have a medical practitioner, and as a pharmacist I thought I would start with you first.

You talked about the white paper you produced. One thing I found interesting, which I found even in my own practice, is the translational research component, which you've termed the “valley of death” because there's a lack of resources.

How do we get beyond that problem? We might be able to create a great piece of equipment or a great piece of software, but the transition to actually seeing it used clinically usually is very difficult.

What do you propose, on a medical basis, whereby we can get the benefit of all this technology but actually apply it usefully to help patients?

Dr. An Tang: Thank you for offering me the opportunity to answer.

Serendipity has it that the federal government recently awarded a strategic innovation fund grant to a consortium led by Imagia Cybernetics, a Canadian start-up specialized in AI in oncology, along with the Terry Fox Foundation and academic radiology departments across the country, in partnership with the four top computer science labs specializing in artificial intelligence. The goal, over a three- to five-year period, is to make sure that we harness the imaging data we have and create new applications that can be used in academic departments prior to commercialization of these products down the road.

● (1620)

Mr. Raj Saini: My second question is for Mr. Leduc.

If we look at the advent of artificial intelligence or the advent of technology, we are now into a different phase of human progress. Automation was created to do repetitive tasks, tasks for which intellectual capacity was not necessarily required because the tasks were repetitive. We are now entering another phase, whether you want to call it industrialization or another phase in our economic growth, whereby artificial intelligence now has the ability to do intellectual tasks.

Now, because of automation of repetitive tasks, you're creating algorithms and creating artificial intelligence through machine learning such that the decision-making is getting better.

How do we deal with potentially having underemployment of a class of people who are educated or trained to apply their intelligence to any task?

Mr. André Leduc: There are a few things built into this.

One is that I think we're going to need to embrace lifelong learning. I think we are going to see the automation of a lot of menial and repetitive tasks and of some human decision-making. You will see, and we've seen it going back in history over time, that when we created the automobile the first time, we didn't need stable boys or stables for the horse and buggy anymore. In our own sector, we replaced the operators who used to walk in front of the switchboard and switch everything. We replaced them with a router and switch.

We believe that the opportunity for these types of technologies to create more employment is going to outpace the disruption. That said, people who are in menial-task fields are often the most vulnerable, and I think we need to embrace programming for re-skilling and upskilling of people who are going to face displacement based on these new technologies.

This is going to come. It isn't an option. Businesses will strive to become more efficient. They need to compete globally. If they can leverage AI to complete tasks within the enterprise, they will choose that route, because it will cost less.

Mr. Raj Saini: Here is one other question I want to ask,

I'm sure you're all aware of the term "singularity". Singularity can be construed as a science fiction term for cases in which eventually you may have overlords of machinery that control human beings. Let's, however, take one step back from that point.

The basis of artificial intelligence and machine learning is for them to be smarter, more efficient and more capable than human thinking. Ultimately, though, there still has to be a human component. If you look at technology the way it is, you can program it within a certain narrative, but there's also the human dimension that makes calculations as you go along. One thing I've read is that if you program autonomous cars to go at the speed limit and human beings don't always go at the speed limit, how do you compensate for that?

If we look at singularity as the end point, how do we make sure that the human dimension is still involved? We want the advantages. We want the resources that AI and machine learning can provide us, but how do we make sure that there's still a human component to ensure that decisions are still being made in the human interest or with human interest involved?

It's an easy question. Take 20 seconds.

Voices: Oh, oh!

Mr. André Leduc: I highlighted in our presentation that there still needs to be human control over all artificial intelligence. It has to be enabled by human control. At the end of the day, there's a lot of fear of the unknown in this space but I think that allowing industry to place the standards and, wherever there are market failures, creating the right legislative and regulatory frameworks to address those market failures is going to be important.

What we as an industry want to see is an ongoing dialogue and a balanced approach to legislation and regulation. We don't say there is no need for it. We understand that going through our own standard setting is not always the be-all and end-all. Sometimes there will be market failures that will require legislative or regulatory action. What we're suggesting is that it has to be done in a dialogue and be balanced so that we don't impede our access to innovation and our ability to do R and D in this field.

•(1625)

Mr. Raj Saini: Here is a final question. You guys can comment on the other question, but I want to make sure I get this question in.

There has been some discussion philosophically that—because of the global race for AI and machine learning, some commentaries have suggested—maybe we should take it one step at a time, because the research far surpasses any legislative ability or any human comprehension of how to deal with the moral and ethical implications of AI.

Would you suggest that we should have some framework whereby, as we hit certain milestones in the progress of AI, we

should take a step back and regroup to think about how we're going to manage the next phase of development?

Mr. André Leduc: I'll go again on that one.

The problem you run into is that not everybody's going to play by the same set of rules. This is like a global space race. Now we're in an AI race. Our country has invested a significant amount of time, money and effort into being leading researchers in research and development of algorithms and artificial intelligence, and we'll need to be able to commercialize that R and D and promote the use of our capabilities and capacities in AI.

If we took the time to take a step back and review and took a couple of years to do it, we'd essentially just be putting up a roadblock vis-à-vis our global competition in this space. That's why we say we think this needs to be an ongoing dialogue. I think it's wonderful that you guys have brought this issue for study, but I think these types of issues around data and leveraging of data, privacy and what frameworks actually work now, and what the issues are around consent.... We've been going through that. Through this committee, you've been doing this for the better part of the last 20 years. To take the time to stop and review will impede our competitiveness.

The Chair: You're way past time. Thank you.

I have to go on to Mr. Kent for seven minutes.

Hon. Peter Kent (Thornhill, CPC): Thank you, Chair.

Thank you all for adding a couple of new dimensions to the study we've been doing on digital government, digital threats, privacy issues and so forth.

I'd like to start with Professor Mittelstadt on the area of the vulnerability of massive amounts of highly personal data across society—medicine, health, so forth, business—and liability and regulation.

I'd like your comment on exactly how bringing in the GDPR, the new spectrum of regulation in Europe, with significant penalties for breaches or improper use of privacy, changed the development of artificial intelligence in its various applications, but also the precautions that have been taken in various industries, such as the health industry, or such as social media on the other hand.

Dr. Brent Mittelstadt: I can say a few things. Largely, the impact of the GDPR is still uncertain because so much of it is vague or the actual requirements it imposes are not entirely clear at this moment. Many complaints have been filed at the member state level that are still being worked through by national data protection authorities.

We'll get some more clarity from those and also as cases are brought in front of national courts and European courts as well. There are very large fines. Data protection authorities are starting to use them, so I think we'll start to see what the actual impact is over the next two or three years in particular.

In terms of how it has actually impacted the development of AI, one effect I would say it's had—although arguably the 1995 data protection directive had this effect as well—has been to encourage developers to anonymize or de-identify data before doing anything of interest with it, because as soon as that has happened, essentially the GDPR no longer applies. It applies to the de-identification process, it applies if you re-link the knowledge that you create back to individuals, but it doesn't apply to anything you do in the in-between stage.

That's one negative, I would say, that it's had. On a positive note, I would say it has encouraged more developers to consider how humans can actually be put into the loop of automated decision-making, because there are several rights that kick in for solely automated processes—essentially, AI that does not have a human in the loop to help make a decision or with the ability to intervene in a decision.

• (1630)

Hon. Peter Kent: The overriding element of consent would touch all of this, presumably.

Dr. Brent Mittelstadt: Certainly. My comments on consent definitely apply here. There are limitations on how data can be repurposed, but again these apply only to identifiable data, so they are limited in their applicability.

Hon. Peter Kent: Monsieur Leduc, how would the Information Technology Association of Canada feel about regulations similar to the GDPR—not necessarily exactly the same but much more regulation than exists today in Canada?

Mr. André Leduc: We follow this issue very closely. We've been reviewing the impacts the GDPR has had, both positive and negative, in the European context. The positive is in terms of improving privacy rights for European citizens. The negative side, as my colleague pointed out, is that there's a lack of transparency and a lack of clear and simple guidance around how to comply.

This is particularly impactful not upon the largest organizations, who have teams of lawyers to filter through the legislation and figure out how to comply. Although it's a cost burden on them, it's particularly impactful upon small and medium-sized enterprises. We've seen a significant impact in the EU.

It's not that we wouldn't welcome GDPR-like principles brought into our digital charter and Canada's data strategy and welcome improvements made to PIPEDA, but I would caution about just turning a light-switch on for GDPR exactly as is. For multinationals, that might make compliance a little bit easier, because they're already GDPR-compliant, but for SMEs, taking the leap from what is today PIPEDA and moving into a GDPR-like framework without clear and simple guidance about how to comply with the law would have a significant negative impact on smaller and medium-sized enterprises.

Hon. Peter Kent: Do you mean in terms of discouraging research and development?

Mr. André Leduc: I mean not just discouraging research and development. It's the cost of compliance. We've seen reports coming out of the EU that the average cost is around \$100,000 U.S. to comply with GDPR. If you take that into account, it's not a lot of money for a very large organization, but for a small business of 10

people that has potentially a million dollars in revenue, it would eat up essentially their entire profit margin at 10%.

It is something I studied in depth over the course of about a year and a half. The lack of simple tools and compliance guidance for SMEs cripples their capability to comply with privacy legislation.

Hon. Peter Kent: I'll go to Mr. Tang and Mr. Neuheimer now.

I consumed with interest your “Who Are Radiologists...?” page. With regard to job loss and job transition as a result of the development of AI in your field, radiology, it would seem that you have some concern about the fourth stage, in which the radiologist does analysis today and refers that analysis to a physician.

If AI progresses to the extent that we're told it will one day, the radiologist's job may be—and you tell me, but it would seem to be—reduced to the first interaction with the patient, taking those images, and then you'd be out of the loop because the doctor would be able to use AI to make the diagnosis and recommend treatment.

Dr. An Tang: To paraphrase Mark Twain, I would say that rumours of our demise are vastly exaggerated at this point.

At the CAR, we've approached this question conceptually. It also addresses the previous question: What are the various levels of autonomy of software? We make an analogy with the self-driving car and we create a scale ranging from zero to five, in which zero indicates no automation at all and then proceeds to physician assistance, partial automation, high automation and full automation. We don't see on our radar anything that will replace the entire work that's accomplished by radiologists.

• (1635)

Hon. Peter Kent: No.

Dr. An Tang: However, we see many helpful applications for specific tasks that are repetitive and mundane and that would free up time for us to perform more meaningful tasks, such as communication, explaining procedures to patients or even performing these procedures and attending tumour boards. This would be much more productive.

The Chair: Thank you, Mr. Kent. We're well past time.

Hon. Peter Kent: Thank you.

The Chair: Mr. Boulerice is next, for seven minutes.

[Translation]

Mr. Alexandre Boulerice (Rosemont—La Petite-Patrie, NDP): Thank you, Mr. Chair.

I thank everyone for being here today for this study and the important questions it raises. We live in a world where artificial intelligence will take up more and more space. It will be given more responsibilities. It will make increasingly complex decisions because its algorithms will be able to process countless amounts of data at a speed faster than any human brain.

I want to put my first questions to Mr. Leduc and Mr. Mittelstadt, and they concern the ethics of artificial intelligence.

An algorithm or supercomputer is in itself incapable of displaying discrimination or bias. On the other hand, the human being who programs the algorithms is capable of doing so at different stages: during data collection, during processing, or during the preparation of questions the algorithm will try to answer.

In your opinion, how, at these different stages, can we avoid these normal human prejudices, which could lead to discriminatory results? Which one of you two wants to dive into this easy question?

[English]

Mr. André Leduc: Perhaps Dr. Mittelstadt could begin and I will follow up.

[Translation]

Mr. Alexandre Boulerice: Fine.

Mr. Mittelstadt, did you want to speak?

[English]

Dr. Brent Mittelstadt: I'm happy to answer this.

It's a very important question, how we both identify bias as it's picked up by algorithms and then mitigate it once we know that it's there. I think the simplest way to explain the problem is that we live in a biased world and we're training algorithms and AI with data about the world, so it's inevitable that they pick up these biases and can end up replicating them or even creating new biases that we're not aware of.

We tend to think of bias in terms of protected attributes—things such as ethnicity, gender or religion, things that are historically protected for very good reasons. What's interesting about AI is that it can create entirely new sets of biases that don't map onto those characteristics or even characteristics that are humanly interpretable or humanly comprehensible. Detecting those sorts of biases in particular is very difficult and requires looking essentially at the set of decisions or outputs of an algorithmic system to try to identify when there is disparate impact upon particular groups, even if they are not legally protected groups.

Besides that, there is quite a bit of research, and methods are being developed to detect gaps in the representativeness of data and also to detect proxies for protected attributes that may or may not be known in the training phase. For example, postal code is a very strong proxy for ethnicity in some cases. It's discovering more sorts of proxies like that.

Again, there are many types of testing—automated methods, auditing methods—whereby essentially you are doing some sort of analysis of training data of the algorithm while it's performing processing and of the sets of decisions that it produces.

There is, then, no simple answer to how you do it, but there are methods available at all stages.

[Translation]

Mr. Alexandre Boulerice: Thank you.

Mr. Leduc, it's your turn.

Mr. André Leduc: As Mr. Mittelstadt said, in artificial intelligence, there are frequent opportunities where biases are created in the data itself and in the codes generated in relation to artificial intelligence. We suggest that, in the industry, there should be a review at each step, whether an algorithm is being developed, databases are being used or data analysis is being done. The aim is to ensure that the results from artificial intelligence processes are not biased.

• (1640)

Mr. Alexandre Boulerice: I have a supplementary question, along the same lines. It is addressed to all of you.

Artificial intelligence algorithms will make decisions that will have an impact on people's lives. They will be used for facial recognition, identification, police investigations, probably, and credit investigations. They will be able to guide decisions regarding the granting of a mortgage or a loan to a business, or hiring decisions. These algorithms will be asked to make decisions that can be considered fair and equitable.

Since the very principle of what is fair and equitable changes with history, culture and ideology, how can we ensure that we get fair and equitable decisions?

[English]

The Chair: Would anybody like to tackle that question?

[Translation]

Mr. Alexandre Boulerice: Perhaps Mr. Mittelstadt would like to respond.

[English]

Dr. Brent Mittelstadt: Yes, I'd be happy to give that a go.

I think, as I was alluding to at the end of my statement, what is going to be determined as fair or ethical is going to be extremely context-dependent. Maybe the highest level we could go to in terms of having guidelines for what constitutes an ethical or fair decision would be at a sectoral level, at which you have existing regulation that gives you some restrictions concerning what is considered permissible or discriminatory, because these things will vary across different sectors.

Really, it's something that can only be answered at that contextual level. I think maybe we have a head start in AI that will be used in professions that are already licensed or legally recognized as professions, where they have fiduciary duties to the people they serve, because they have these very long histories where they've developed best practices, guidelines, principles and lower-level norms, basically, to define what is a good behaviour and what is a good decision.

It's a difficult question, but I think that's how we start.

[*Translation*]

Mr. Alexandre Boulerice: Thank you.

Mr. Tang, did you want to speak, briefly?

Dr. An Tang: Yes, I will venture to answer both your questions, because I had some time to think about it.

On the issue of bias, I would say that one of the strategies to minimize it in the medical field would be to use a large amount of data to reflect the target population, particularly in terms of gender, ethnicity or age group.

As far as discrimination is concerned, I think the best way to minimize it is to keep the human element in the equation and involve a doctor or other member of the care team. Indeed, in the end, health care is highly personalized and deeply affects privacy. Beyond the recommendation established by the algorithm on the basis of a large demographic sample, the decision made by the patient and physician will remain.

[*English*]

The Chair: I'm just going to explain that we were going to do some committee business in about five minutes, but we've talked with the vice-chairs and all parties, and we're going to push that committee business to next Tuesday, if you can stick around until five o'clock.

Is that something you can do? Okay. We'll take it right to five with questions.

We'll go with Nate for seven minutes.

Mr. Nathaniel Erskine-Smith (Beaches—East York, Lib.): Thanks very much.

I want to start with you, Mr. Mittelstadt, and talk first about AI, risk assessments and algorithmic transparency. At a government level, there are now rules that the Treasury Board has put in place for government agencies and departments. It's a risk assessment, and then, depending upon how they answer the 85 questions, they're categorized from stages one to four. Depending on where they slot in, there are mitigation measures that are then required.

Perhaps you can explain the usefulness of that, if you think it's useful, and the deficiencies, if you think there are deficiencies, and how we can improve upon that, potentially, and what else might be required.

Dr. Brent Mittelstadt: I tend to say that there is no single silver bullet for appropriate governance of these systems, so risk assessments can be a very good starting point.

They're very good in the sense of catching problems in the pre-deployment or procurement stage. Their shortcoming is that they're only as good as the people or the organizations that complete them, so they do require a certain level of expertise and, potentially, training—essentially, people who are aware of potential ethical issues and can flag them up while actually going through the questionnaire.

We've seen that with other sorts of impact assessments such as privacy impact assessments, environmental impact assessments and, now, data protection impact assessments in Europe. There really has to be a renewed focus on the training or the expertise of the people who will be filling those out.

They are useful in a pre-deployment sense, but as I was suggesting before with biases, problems can emerge after a system has been designed. We can test a system in the design phase and during the training phase and say that it seems to be fair, it seems to be non-discriminatory and it seems to be unbiased, but that doesn't mean that problems won't then emerge when the system is essentially used in the wild.

Any sort of impact assessment approach has to be complemented as well by in-process monitoring and post-processing assessment of the decisions that were made, and very clear auditing standards in terms of what information needs to be retained and what sorts of tests need to be carried out after the fact, again, to check for things like bias.

• (1645)

Mr. Nathaniel Erskine-Smith: That's helpful. There's at least a model or a template for algorithmic impact assessments that seems somewhat transferable, at least to the private sector, for bigger companies at a minimum.

We've recommended that there then be a regulatory authority to conduct audits, not only against that original assessment but also potentially ongoing. Is that the kind of thing...? Ought there be some regulator with the power to audit practices of companies that are engaging in the use of algorithms? Is that the idea?

Dr. Brent Mittelstadt: It could be a regulator that does it. I watched Christian Sandvig's testimony to this committee. He pointed out the difference between financial audits and social, scientific and computational audits. I suppose it's more the latter that I'm thinking of here.

You can have a regulator do it, but again, that introduces the problem of whether the regulator actually has the expertise required. Do they understand the system that's being used? Do they have access to actually understand the system, what data it's considering and what its purpose is? There are problems with relying solely on a third party independent regulator.

What I would like to see is more willingness, particularly from private companies, to share a bit more about not only the auditing that they're doing of their systems—in-processing and post-processing auditing—but also just more generally the impact that ethical principles have had on their development and deployment of these systems. In other words, I want to know a lot more about specific cases where they've said no or they've changed the design of the system as the result of an impact assessment or as a result of auditing.

Mr. Nathaniel Erskine-Smith: Isn't, though, the lesson learned to date from big companies that do collect large amounts of data, and then employ algorithms, that they're not implementing ethical principles in the first instance, and that there need to be rules brought to bear?

Dr. Brent Mittelstadt: I'll say that it's not clear the extent to which they are implementing ethical principles. I know that there are some companies that do have feedback mechanisms, but they tend to be more internal. They are very happy to report on positive cases, where ethical considerations have led them to change the system in a positive way or to design a new type of system, but in terms of public-facing sorts of very critical self-assessments, it's not a huge deal.

Mr. Nathaniel Erskine-Smith: What is the economic incentive for them to do that?

Dr. Brent Mittelstadt: In the first instance, yes, you could argue that there is no economic incentive to do that because it can make you look worse than your competitors.

Actually, one of my other slight concerns is that ethics turns purely into something that is marketable in the same way that, say, having an organic label on your product makes it seem more ethical and more valuable. I don't know.... I'm very cynical about that happening.

Mr. Nathaniel Erskine-Smith: Thanks.

Whether you call it algorithmic transparency or algorithmic explainability, as the GDPR does, when some of us were in the U.K. and asking questions of the ICO, Elizabeth Denham, she said her role was to make the algorithm as explainable as possible, and that it was for other regulators—the human rights commissioner, say, or the competition authority—to better assess, with their expertise.

Similarly, we had experts in the technical side of AI before us at the outset of this study who said that transparency rules make sense across the board and that beyond that you need sector-specific regulators and rules to apply, which would simply take AI into account. Do you think that makes sense?

• (1650)

Dr. Brent Mittelstadt: Again, if the sector-specific regulators have the necessary expertise to do so, and if they're sufficiently resourced to do so, it could work. I think it's worth—

Mr. Nathaniel Erskine-Smith: Just to pick up on that, though, isn't that the point? Take regulating in the auto sector, for example. They're employing AI for autonomous vehicles. Do we have this stand-alone regulator, whether it's the ICO model or the privacy data commissioner model, where we roll algorithmic accountability into their function? Or is it simply that the regulatory authorities and the

rules that are brought to bear on the auto sector have to account for and build up capacity with respect to assessing algorithms?

Dr. Brent Mittelstadt: The problem is that I can imagine both models working if there's openness to reforming the sectoral rules that the regulator is enforcing. I don't see any particular reason why it couldn't work—again, assuming that there are sufficient expertise and resources available to it.

At the same time, it does make some sense to have a general regulator, at least for certain types of issues, such as the ICO, for example, the data protection authority. Many of the issues with AI have to do with how data is collected, repurposed and used. For those sorts of issues, yes, it makes sense to have them deal with the challenges of AI, but there will be other sorts of issues that are very specific to specific types of AI where I think having the sectoral regulator deal with them makes the most sense.

Mr. Nathaniel Erskine-Smith: Thanks very much.

The Chair: We'll go to Monsieur Gourde for five minutes.

[*Translation*]

Mr. Jacques Gourde (Lévis—Lotbinière, CPC): Thank you, Mr. Chair.

There is no doubt that artificial intelligence will play a major role in the global economy. However, I have the impression that funding, which comes mainly from the private sector and to a lesser extent from the public sector, is not necessarily intended to ensure the well-being of humanity.

Can we know what proportion of artificial intelligence budgets is allocated to military activities and what proportion goes to health?

I think that, in the health field, this will help everyone. On the other hand, in the military field, we will create super powerful weapons and hope never to use them. These funds might have been more useful to humanity if they had been invested in health. There is no doubt that companies are looking to make a profit. They go where money is available and contracts are easy to obtain. Ethically, there will be a global problem.

What do you think, Mr. Leduc and Mr. Tang?

Mr. André Leduc: With regard to the products and services provided to Canada's Department of National Defence or elsewhere in the world, I don't think it's very different from what we're seeing in traditional sectors. In discussing the ethics of artificial intelligence, we seek to determine in which cases our society will approve the use of artificial intelligence and in which other cases artificial intelligence may be used to develop products for military personnel whose country is in conflict with ours. There are always risks.

As far as funding is concerned, I don't know the answer. So I can't answer you. The Ministry of Defence, when it wants to solve problems it faces, often uses the tools that can be provided to it. More and more, we see in our field that artificial intelligence is integrated into all technologies. Implicitly or explicitly, decisions will therefore be made more and more by artificial intelligence, simply to make products and services more efficient.

• (1655)

Dr. An Tang: Historically, Canada has been a true leader in funding basic research, including through CIFAR, the Canadian Institute for Advanced Research. This has enabled Canadian research laboratories to play a leading international role, particularly in the field of deep learning, which has generated recent interest in artificial intelligence.

In addition, the federal government has funded initiatives such as the Canada First Research Excellence Fund, which is a competition, and has funded things like the Data Serving Canadians initiative. This initiative has transformed fundamental knowledge used in four specific sectors: health, logistics, e-commerce and the financial sector. There are concrete examples of funding useful activities. In addition, supporting such research activities has the effect of attracting a critical mass of industries that will invest in the field.

Mr. Jacques Gourde: Mr. Mittelstadt, did you want to add something?

[English]

Dr. Brent Mittelstadt: I can add a bit, I suppose, not specifically on the budget, although I would note that particularly in research on explainability or interpretability in AI, there's a huge amount of money going into it in the U.S. from their defence department, and what we've seen, at least in the past, is a crossover into the private sector from military developments in technology. Besides that, there is plenty of academic and commercial research outside of the military context that addresses this.

Beyond that, I don't know that I have much of a comment, particularly with not being extremely familiar with the Canadian context.

[Translation]

Mr. Jacques Gourde: Thank you.

[English]

The Chair: Thank you, Mr. Gourde.

Last up for five minutes is Monsieur Picard.

[Translation]

Mr. Michel Picard (Montarville, Lib.): Mr. Tang, you mentioned earlier that the information could be made anonymous, so that no link can be made between the patient and the information.

Doesn't this practice, which is intended to protect the individual, have a perverse effect? Indeed, once the information is anonymous, consent is no longer required to conduct the studies you want with the data you have.

Dr. An Tang: This can be approached from various angles. It is important to know where the information is made anonymous. If the data remains within a hospital, for example, and only the algorithms are removed, there is no breach of confidentiality. In this case, only

the researchers and doctors involved know the identity of the patient. In addition, there is a field of research that allows the teaching of several institutions to be shared, which is extremely advantageous. Data can indeed be kept within institutions, and only the learning process is shared by many hospitals.

[English]

Mr. Michel Picard: Mr. Mittelstadt, can you come back to what you said about protecting data and not protecting people? Did you say that we did protect data but not people, or we should protect data and not necessarily people because that is what AI is all about?

Can you comment on that, please?

Dr. Brent Mittelstadt: Yes. Thank you for the question.

The point of the comment was to say that data protection and privacy law are designed to protect people—or at least that's what inspired them originally—and to protect privacy in all of its various different types.

However, functionally or procedurally, what it ends up doing is protecting the data that are produced by people. This links to the comments I was making around informed consent and the need for identifiability for the law to apply in the first place. As has been described throughout the entire session today, once the data is de-identified or anonymized, you can still do very interesting things with it to create very useful knowledge about groups of people, which can then be applied back to those groups. In the case of medical research, it's very laudable, but in other cases, maybe not so much.

Mr. Michel Picard: That suggests that we then have to find a way legally to.... With not being able to separate data from the people, when there's harm done to data, there is therefore harm done to someone, because somewhere the data concerns someone. You can protect someone but you can't sue the data, but data is the centre of the focus.

Do we lack something, legally speaking ?

• (1700)

Dr. Brent Mittelstadt: I'd say that most places lack something, legally speaking. Dealing with the collective or group aspects of privacy and data protection is extremely difficult. There's not really a satisfactory legal framework for it, outside of specific types of harm such as discrimination.

We could say that we all lack something legally.

[Translation]

Mr. Michel Picard: Mr. Leduc, you surprised me by saying that everyone was caught off guard and no one anticipated how important data would be or how much influence the information would have on our daily lives.

However, we have been talking about information, its added value and its commercialization for some time now, if we remember Alvin Toffler's *Future Shock* and especially his book *The Third Wave*, published in 1980. We have known for a long time that information has an extraordinary and precious value. Could we conclude that we chose to close our eyes rather than say that we hadn't seen anything coming?

Mr. André Leduc: It's not that we closed our eyes. Instead, priority was given to the development of bills that were technology-neutral.

No one could have predicted the quantity of data and transactions an individual can generate in a day because, for example, the smartphone had not yet been invented. As consumers or citizens, we use our phone several times a day, and each time, we exchange data with the service provider we are using.

Every year, we produce more data than we have so far generated in the entire history of humanity. Thus, in 2018, we produced more data than there has been in the entire history of humanity since the very first records until 2017.

We knew that data was valuable, but no one predicted the increase in this data related to interactions between consumers or citizens and businesses, all facilitated by telecommunications technologies.

Mr. Michel Picard: In other words, if I interpret...

[*English*]

The Chair: Okay, Monsieur Picard, we're over time.

Mr. Michel Picard: I had a good one.

Voices: Oh, oh!

The Chair: I'm sorry, everybody. I know we could keep going on this. I think the conversation today was another good one on the really cool aspects of technology and where technology is going, but also the challenges of privacy and security that Monsieur Picard brought up today. That's our challenge.

Thanks for appearing, and hopefully we can have you back again to have a more fulsome discussion.

The meeting is adjourned.

Published under the authority of the Speaker of
the House of Commons

SPEAKER'S PERMISSION

The proceedings of the House of Commons and its Committees are hereby made available to provide greater public access. The parliamentary privilege of the House of Commons to control the publication and broadcast of the proceedings of the House of Commons and its Committees is nonetheless reserved. All copyrights therein are also reserved.

Reproduction of the proceedings of the House of Commons and its Committees, in whole or in part and in any medium, is hereby permitted provided that the reproduction is accurate and is not presented as official. This permission does not extend to reproduction, distribution or use for commercial purpose of financial gain. Reproduction or use outside this permission or without authorization may be treated as copyright infringement in accordance with the *Copyright Act*. Authorization may be obtained on written application to the Office of the Speaker of the House of Commons.

Reproduction in accordance with this permission does not constitute publication under the authority of the House of Commons. The absolute privilege that applies to the proceedings of the House of Commons does not extend to these permitted reproductions. Where a reproduction includes briefs to a Committee of the House of Commons, authorization for reproduction may be required from the authors in accordance with the *Copyright Act*.

Nothing in this permission abrogates or derogates from the privileges, powers, immunities and rights of the House of Commons and its Committees. For greater certainty, this permission does not affect the prohibition against impeaching or questioning the proceedings of the House of Commons in courts or otherwise. The House of Commons retains the right and privilege to find users in contempt of Parliament if a reproduction or use is not in accordance with this permission.

Also available on the House of Commons website at the following address: <http://www.ourcommons.ca>

Publié en conformité de l'autorité
du Président de la Chambre des communes

PERMISSION DU PRÉSIDENT

Les délibérations de la Chambre des communes et de ses comités sont mises à la disposition du public pour mieux le renseigner. La Chambre conserve néanmoins son privilège parlementaire de contrôler la publication et la diffusion des délibérations et elle possède tous les droits d'auteur sur celles-ci.

Il est permis de reproduire les délibérations de la Chambre et de ses comités, en tout ou en partie, sur n'importe quel support, pourvu que la reproduction soit exacte et qu'elle ne soit pas présentée comme version officielle. Il n'est toutefois pas permis de reproduire, de distribuer ou d'utiliser les délibérations à des fins commerciales visant la réalisation d'un profit financier. Toute reproduction ou utilisation non permise ou non formellement autorisée peut être considérée comme une violation du droit d'auteur aux termes de la *Loi sur le droit d'auteur*. Une autorisation formelle peut être obtenue sur présentation d'une demande écrite au Bureau du Président de la Chambre.

La reproduction conforme à la présente permission ne constitue pas une publication sous l'autorité de la Chambre. Le privilège absolu qui s'applique aux délibérations de la Chambre ne s'étend pas aux reproductions permises. Lorsqu'une reproduction comprend des mémoires présentés à un comité de la Chambre, il peut être nécessaire d'obtenir de leurs auteurs l'autorisation de les reproduire, conformément à la *Loi sur le droit d'auteur*.

La présente permission ne porte pas atteinte aux privilèges, pouvoirs, immunités et droits de la Chambre et de ses comités. Il est entendu que cette permission ne touche pas l'interdiction de contester ou de mettre en cause les délibérations de la Chambre devant les tribunaux ou autrement. La Chambre conserve le droit et le privilège de déclarer l'utilisateur coupable d'outrage au Parlement lorsque la reproduction ou l'utilisation n'est pas conforme à la présente permission.

Aussi disponible sur le site Web de la Chambre des communes à l'adresse suivante : <http://www.noscommunes.ca>