



HOUSE OF COMMONS  
CHAMBRE DES COMMUNES  
CANADA

# **Standing Committee on Access to Information, Privacy and Ethics**

---

ETHI • NUMBER 146 • 1st SESSION • 42nd PARLIAMENT

---

TÉMOIGNAGES

**Thursday, May 2, 2019**

—  
**Chair**

**Mr. Bob Zimmer**



## Standing Committee on Access to Information, Privacy and Ethics

Thursday, May 2, 2019

• (1530)

[English]

**The Chair (Mr. Bob Zimmer (Prince George—Peace River—Northern Rockies, CPC)):** We'll call to order meeting 146 of the Standing Committee on Access to Information, Privacy and Ethics. Pursuant to Standing Order 108(2), we're continuing our study on the ethical aspects of artificial intelligence and algorithms.

In the first hour today, we have with us, as individuals, witnesses Marc-Antoine Dilhac, professor of philosophy, Université de Montréal, and Christian Sandvig, director of the Center for Ethics, Society, and Computing, University of Michigan.

Also, as we all know, we're going to discuss in camera, pursuant to Standing Order 108, a briefing by the law clerk of the House on the power of committees to summon witnesses. That will be our discussion following this hour.

We'll start off with Marc-Antoine for 10 minutes.

[Translation]

**Prof. Marc-Antoine Dilhac (Professor, Philosophy, Université de Montréal, As an Individual):** Good afternoon.

Thank you for inviting me to share with you some of the reflections about artificial intelligence ethical issues which we set out in Montreal.

I was asked to speak about the Montreal declaration for the responsible development of artificial intelligence, which was presented in 2018. I will speak about this document.

First I will outline the context in broad strokes. The technological revolution that is taking place is causing a profound change in the structure of society, by automating administrative processes and decisions that impact the life of our citizens. It also changes the architecture of choice by determining our default options, for instance. And it transforms lifestyles and mentalities through the personalization of recommendations, access to online, automated health advice, the planning of activities in real time, forecasting, and so on.

This technological revolution is an unprecedented opportunity, it seems to me, to improve public services, correct injustices and meet the needs of every person and every group. We must seize this opportunity before the digital infrastructure is completely established, leaving us little or no leeway to act.

To do so we must first establish the fundamental ethical principles that will guide the responsible and sustainable development of artificial intelligence and digital technologies. We must then develop standards and appropriate regulations and legislation. In the Montreal Declaration for a Responsible Development of Artificial Intelligence, we proposed an ethical framework for the regulation of the artificial intelligence sector. Although it is not binding, the declaration seeks to guide the standardization, legislation and regulation of AI, or artificial intelligence. In addition, that ethical framework constitutes a basis for human rights in the digital age.

I will quickly explain how we developed that declaration. This may be of interest in the context of discussions about artificial intelligence in our democratic societies. Then I will briefly present its content.

The declaration is first and foremost a document produced via the consultation of various stakeholders. It was an initiative of the University of Montreal, which received support from the Fonds de recherche en santé Québec and from the Canadian Institute for Advanced Research, or CIFAR, in the rest of Canada. Behind this declaration there was a multidisciplinary inter-university working group from the fields of philosophy, ethics, the social sciences, law, medicine, and of course, computer science. Mr. Yoshua Bengio, for instance, was a member of this panel.

This university group then launched, in February 2018, a citizens' consultation process, in order to benefit from the field expertise of citizens and AI stakeholders. It organized over 20 public events and discussion seminars or workshops over eight months, mainly in Quebec, but also in Europe, Paris and Brussels. More than 500 people took part in these workshops in person. The group also organized an online consultation. This consultation process was based on a prospective methodology applied to ethics; our group invited workshop participants to reflect on ethical issues based on prospective scenarios, that is to say scenarios about the near future of the digital society.

We organized a broad citizen consultation with various stakeholders, rather than consulting experts alone, for several reasons. I will mention three, rapidly.

• (1535)

The first reason is that AI is being deployed in all societies and concerns everyone. Everyone must be given an opportunity to speak out about its deployment. That is a democratic requirement.

The second reason is that AI raises some complex ethical dilemmas that touch on many values. In a multicultural and diverse society, experts alone cannot make decisions on the ethical dilemmas posed by the spread of artificial intelligence. Although experts may clarify the ethical issues around AI and establish the conditions for a rational debate, they must design solutions in co-operation with citizens and all parties concerned.

The third reason is that only a participative process can sustain the public's trust, which is necessary to the deployment of AI. If we want to earn the population's trust and give it good reasons to trust the actors involved with AI, we have the duty to involve the public in the conversation about AI. That isn't a sufficient condition, but it is a necessary condition to establish trust.

I should add that although industry actors are very important as stakeholders, they must stop wanting to write the ethical principles instead of citizens and experts, and the legislation that should be drafted by Parliaments. That attitude is very widespread, and it can also undermine the public trust that needs to be fostered.

Let's talk about the content of the declaration. The consultation had a dual objective. First, we wanted to develop the ethical principles and then formulate public policy recommendations.

The result of that participatory process is a very complete declaration that includes 10 fundamental principles, 60 subprinciples or proposals to apply the principles, and 35 public policy recommendations.

The fundamental principles touch on well-being, autonomy, private life and intimacy and solidarity—that principle is not found in other documents—democracy, equity, diversity, responsibility, prudence and sustainable development.

The principles have not been classified according to priority. The last principle is no less important than the first, and according to circumstances, a principle may be considered more relevant than another. For instance, if privacy is in general considered a matter of human dignity, the privacy principle may be considered less important for medical purposes, if two conditions are met: it must contribute to improving the health of patients—under the well-being principle—and the collection and use of private data must be subject to individual consent—the autonomy principle.

The declaration, thus, is not a simple checklist, but it also establishes standards and checklists according to activity sectors. Thus, the privacy regime will not be the same, according to the sector, for instance; it may vary depending on whether we are talking about the medical or banking sector.

The declaration also constitutes a basis for the development of legal norms, such as legislation.

Other similar declarations, such as the Helsinki Declaration on Bioethics, are also non-binding declarations like ours. Our declaration simply lists the principles which the AI development actors should commit to respecting. For us, the task is now to work on transposing those principles into industrial standards that also affect the deployment of artificial intelligence in public administrations.

●(1540)

We are also working on the transposition of those principles into human rights for the digital society. That is what we are going to try to establish through a citizens' consultation which we hope to conduct throughout Canada.

Thank you.

[English]

**The Chair:** Thank you very much, Marc-Antoine.

Next up we'll have Christian Sandvig for 10 minutes.

Go ahead.

**Professor Christian Sandvig (Director, Center for Ethics, Society, and Computing, University of Michigan, As an Individual):** Thank you very much.

I appreciate the opportunity to address the committee. To frame my remarks before I begin with the substance of my comments, I just want to say that my position is that we are at a moment where I'm delighted that the committee is holding these hearings. We're at a moment where there is increasingly a widespread concern about the harms that might be possible as a result of these systems, meaning artificial intelligence and algorithms.

I thought what I could offer to you in my brief opening remarks would be some sort of assessment of what might governments do in this situation. What I'd like to do with my opening statement is to discuss five areas in which I believe there is the most excitement in communities of researchers and practitioners and policy-makers in this area right now. I offer you my assessment of these five areas. Many of them are areas that you at least preliminarily addressed in your earlier reports, but I think that I have something to add.

The five areas I'll address are the following: transparency, structural solutions, technical solutions, auditing and the idea of an independent regulatory agency.

I'll start with transparency. By far the most excitement in practice and policy circles right now has to do with algorithmic automation centres and the idea that we can achieve justice through transparency. I have to tell you, I'm quite skeptical of this area of work. Many of the problems that we worry about in the area of artificial intelligence and transparency are simply not amenable to transparency as a solution. One example is that we're often not sure that the problems are amenable to individual action, so it is not clear that disclosing anything to individuals would help ameliorate any difficulty.

For example, a problem with a social media platform might require expertise to understand the risk. The idea of disclosing something is in some ways regressive because it demands time and expertise to consider the sometimes quite arcane and complicated intricacies of a system. In addition, it might not be possible to perceive the risk at all from the perspective of the individual.

There is a tenet of transparency that we need to be sure that what is revealed has to be matched with the harm that we hope to detect and prevent, and it's just not clear that we know how to match what should be revealed with the harms we hope to prevent.

Sometimes we discuss transparency as a tactic that we use so that we can match what is revealed to an audience that will listen. This is often something that is missing from the debates right now on transparency and artificial intelligence. It's not clear who the audience would be that we need to cultivate to understand disclosures of details of these systems. It seems like they must be experts and it seems like deconstructing these systems would be quite time consuming, but we don't know who exactly they would be.

A key problem that's really specific to this domain that is sometimes elided in other discussions is that algorithms are often not valuable without data and data are often not valuable without algorithms. So if we disclose data we might completely miss an ethically or societally problematic situation that exists in the algorithm and vice versa.

The challenge there is that you also have a scale problem if you need both the data and the algorithm. It's often not clear just in practical terms how you would manage a disclosure of such magnitude or what you would do once you receive the information. Of course, the data on many systems also is continually updated.

Ultimately, I think you have gathered from my remarks I'm pessimistic about many of the proposals about transparency. In fact, it's important to note that when governments pass transparency requirements they can often be counterproductive in this area because it creates the impression that something has happened, but without some effective mechanism of accountability and monitoring matched to the transparency, it may be that nothing has happened. So it may actually harm things to make them transparent.

• (1545)

An example of a transparency proposal that's gotten a lot of excitement recently would be dataset labels that are somehow made equivalent to food labels, such as nutrition facts for datasets or something like that. There are some interesting ideas. There would be a description of biases or ingredients that have an unusual provenance—where did the data come from?—but the metaphor is that tainted ingredients produce tainted food. Unfortunately, with the systems we have in AI, it's not a good metaphor, because it's often not clear, without some indication of the use or context, what the data are meant to do and how they will affect the world.

Another attractive, exciting idea in this space of transparency is the right to explanation, which is often discussed. I agree that it's an attractive idea, but it's often not clear that processes are amenable to explanation. Even a relatively simple process—it doesn't have to be with a computer; it could be the process by which you decided to join the House of Commons—might be a decision that involves many factors, and simply stating a few of them doesn't capture the full complexity of how you made that decision. We find the same things with computer systems.

The second big area I'll talk about is structural solutions. I think this was covered quite well in the committee's previous report, so I'll just say a couple of things about it.

The idea of a structural solution might be that because there are only a few companies operating in some of these areas, particularly in social media, we might use competition or antitrust policy to break

up monopoly power. That, by changing the structure and incentives of the sector, could lead to the amelioration of any harms we foresee with the systems.

I think it is quite promising that if we change the incentives in a sector we could see changes in the harms that we foresee; however, as your report also mentioned, it's often not clear how economies of scale operate in these platforms. Without some quite robust mechanism for interoperability among systems, it's not clear how an alternative that's an upstart in the area of social media or artificial intelligence—or really any area where there is a large repository of data required—would be effective.

I think that one of the most exciting things about this area might be the idea of a public alternative in some sectors. Some people have talked about a public alternative to social media, but it still has this scale problem, this problem of network effects, so I guess we could summarize that area by saying that we are excited about the potential but we don't know exactly how to achieve the structural change.

One example of a structural change that people are excited about and is more modest is the information fiduciary proposal, whereby a government might regulate a different incentive by just requiring it. It's a little challenging to imagine, because it does seem like we are most successful with these proposals when we have a domain with strong professionalization, such as doctors or lawyers.

The third area I will discuss is the idea of a technical solution to problems of AI and algorithms. There's a lot of work currently under way that imagines that we can engineer an unbiased fair or just system and that this is fundamentally a technical problem. While it's true that we can imagine creating these systems that are more effective in some ways than the systems that we have, ultimately it's not a technical problem.

Some examples that have been put forward in this area include the idea of a seal of approval for systems that meet some sort of standard that might be done via testing and certification. This is definitely an exciting area, but only a limited set of the problems we face would fall into the domain that could be tested systematically and technically solved. Really, these are really societal problems, as the previous witness stated.

The fourth area I'll introduce is the idea of auditing, which I saw mentioned only briefly in the committee's last report. The auditing idea is my favourite. It actually comes from work to identify racial discrimination in housing and employment. The idea of an audit is that we send two testers to a landlord at roughly the same time and ask for an apartment. The testers then see if they get different answers, and if they get different answers, something is wrong.

•(1550)

The exciting thing about this area is that we don't need to know the landlord's mind or to explain it. We simply figure out if something is wrong. There's a lot that legislatures can do in the area of testing. They can protect third parties that wish to investigate these systems or they can create processes akin to software's "bug bounties", but the bounties could be for fairness or justice. This is I think the most promising area that governments can use to intervene.

Finally, I'll conclude by just mentioning there is also talk of a new agency, a judicial administrative law or commission agency to handle the areas of AI. I think this is an interesting idea, but the challenge is that it just postpones many of the comments I made in the earlier parts of my remarks. We often would imagine such an agency doing some of the same things that I've already discussed, so the question then becomes, what is different about this area that requires processes that are not the processes of the legislature and the standard law-making apparatus—the courts—that we already have? The argument has been made that expertise makes this different, but it's hard to sustain that argument, because we often do see plain old legislatures making rules about quite complicated areas.

I'll conclude there. I'm happy to take your questions.

**The Chair:** Thank you very much, both of you, for your testimony.

We'll start off with seven-minute questioning rounds.

We'll start with Ms. Vandenberg.

**Ms. Anita Vandenberg (Ottawa West—Nepean, Lib.):** Thank you very much, Chair, and thank you, both of you, for your very informative remarks.

One of the things I have been wondering about—actually, Professor Dilhac, you mentioned this when you talked about involving civil society—is that for most people this is a very misunderstood area, and that's for people who are not technical experts, and even, I would imagine, for some technical experts.

First of all, you have popular culture myths around artificial intelligence that go back decades. Many people aren't aware of how prevalent it already is in our day-to-day lives. If you have systems whereby you involve civil society, legislators or people who are not technical experts to oversee this, how do you ensure that you're not then taking the biases that exist in these systems and in the public and just replicating all of those once again and amplifying the same bias?

[Translation]

**Prof. Marc-Antoine Dilhac:** Thank you for that excellent question.

There is no miracle solution.

The idea is to get all of civil society working together, the informatics, ethics and social science experts, as well as the industry stakeholders. There is more than one useful expertise. We will manage to reduce potential bias, preconceived ideas about the most vulnerable people, among others, by getting the various experts together and getting them to talk to each other. The reason is that the discussion among those experts will rationalize the debate. It might

be a preconceived philosopher's idea to believe in the rationalization of the debate. However, in the context of meetings in Parliament or with citizens in libraries—we have done a lot of meetings in public libraries—dialogue leads to a rationalization of arguments and allows people to identify the prejudice that may be in them. That collaboration is really important.

•(1555)

[English]

**Ms. Anita Vandenberg:** Professor Sandvig, did you want to respond to that?

**Prof. Christian Sandvig:** I'm happy to defer to my colleague.

**Ms. Anita Vandenberg:** Thank you.

Going back to what you said about transparency, Professor Sandvig, because this is something we've heard a lot about in this committee, if people know where the data is coming from and they understand how the algorithms work, this would allow a certain amount of oversight, as it does in many other areas.

You're suggesting that transparency alone would not actually have that effect. In order to have audits and in order to have a regulator, obviously the information needs to be available, even if you were to audit that information. Are you saying that we need transparency but in such a way that we know who the "who" is in terms of who is actually going to be reviewing? Or is it the public in general—civil society—that would have to do that?

**Prof. Christian Sandvig:** Thank you very much for this question, because I think it exposed a weakness in my own explanation.

In the social science literature, they use the term "audit", but they don't use it in the financial sense. The audit simply describes the process I outlined where two testers, say one black and one white or one woman and one man, ask a landlord for a room or an employer for a job. They call that an audit, but it's quite confusing, because obviously the tax authorities also have an audit and it means something else.

I think the reason audits are exciting to me is that you can have an audit without transparency. Remember that I said you don't get to see the inside of the landlord's brain. That's why the audit is exciting. We can audit platforms like Facebook and Google without transparency by simply protecting third parties like researchers, investigative journalists and civil society organizations like NGOs, who wish to see if there are harms produced by these systems. To do that, they would act like the testers in my example. They would act as users of the systems and then aggregate these data to see if there were patterns that were worrying.

Now, this has some shortcomings. For example, you might have to lie. Auditors lie. The people who go in to ask a landlord for a room don't actually want a room; they're testers working for an NGO or a government agency. So you might have to lie; you might have to waste the landlord's time, but not very much.

Usually on systems like large Internet platforms, it's hard to imagine that an audit would be detectable. However, it's possible that you would provide false information that makes it into the system somehow, because you aren't actually looking for a job; you're just testing. There are definitely downsides.

As I mentioned, you also need some sort of system to continue... after your audit finds that there is a problem. For example, if you found that there was something worrying, you would then need some other mechanism like a judicial proceeding, say, involving some disclosure. You could say that transparency comes later through another process, if you needed to really understand how the system works. However, you might never need to understand that. You might just need to detect that there is a harm and tell the company they have to fix it, and they're the ones that have to worry about how.

This is why I'm excited about auditing, because it gets around the problems of transparency.

**Ms. Anita Vandenberg:** Professor Dilhac.

[*Translation*]

**Prof. Marc-Antoine Dilhac:** Yes, I'd like to add something.

I am in complete agreement with what has just been said about transparency. I think we overestimate transparency. The mechanism to test the algorithms is probably the best way to proceed to identify problems.

Nevertheless, I would use the term “audit” in both senses of the word. First in the sense that was just used, and also in the sense that competent authorities must use algorithms to detect the problematic parameters of a decision. We could use both, but first I think that the solution that was proposed before is an excellent process.

• (1600)

[*English*]

**The Chair:** Thank you.

Next up, for seven minutes, is Mr. Kent.

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

Thank you both for your participation today.

Professor Sandvig, with regard to regulation, many lay observers, many with concerns about what they read, what they hear about algorithmic use and AI, may feel—and given some of the testimony we've heard, for example, about Cambridge Analytica, about some of the big data-opoly use of AI to affect and direct consumer retail attitudes, social attitudes and so forth, sometimes I feel..... Is there an element of being a little late to put some of the smoke back into the bottle in terms of regulation? Can some of the inappropriate or unethical applications of AI and algorithms to date be reverse regulated?

I would throw that out to both of you, but Professor Sandvig first.

**Prof. Christian Sandvig:** Well, my background before going to graduate school was as a software engineer, and my memory of that time as a software engineer fits with what many commentators are saying now: that software engineering does not have a safety culture. It does not have a culture that we would analogize to industries like, for example, airline travel.

I guess the question is, can we imagine changing something that's big and that has already happened? You mentioned recent revelations about Facebook and the Cambridge Analytica scandal. Can we imagine changing something that looks extremely bad with something that looks extremely safe?

Again, I used the example of air travel, but I think it's possible.... I mean, I can't imagine that the Wright brothers had a safety culture, right? There was some way in which government regulation started slowly and accreted. We have an industry that's regulated, and we now consider—perhaps the Max 8 is an exception—this industry to be a safe industry. We're not concerned about air travel.

I would say that this is the trajectory we need for these industries. We need a sense where it's the role of the government to make sure that the public is safe, and we can do it with social media if we did it with other dangerous things.

**Hon. Peter Kent:** You gave the example of the 737 Max 8. Is that a failure to audit after the first tragedy?

**Prof. Christian Sandvig:** Well, a hearing on airline safety...so you'll forgive me for saying I'm not sure. I think that in general I would point to the level of comfort that people currently have with air travel as my main point, even despite the Max 8s. I think we could imagine something like that for social media platforms or artificial intelligence.

We're currently very far away from it, so I don't mean to at all minimize your concern. Recent news reports show that many of these companies are at rock bottom in terms of consumer trust in their operations or in customer satisfaction, and I mean really below even any other industry. We're looking at some of the most hated industries.

**Hon. Peter Kent:** Professor Dilhac.

[*Translation*]

**Prof. Marc-Antoine Dilhac:** In the Montreal Declaration for a Responsible Development of Artificial Intelligence, for instance, one of the principles mentioned is prudence. The idea behind that is to state that there are security and reliability criteria for the algorithms, but not only for the algorithms. I would like to expound on the topic because the way in which the algorithm is put in place in a system is important.

There is a whole system around an algorithm, like other algorithms, data bases, and their use in a specific context. In the case of a platform, it is easy since you have an individual user behind his screen. However, when you are talking about aircraft or a complex enterprise, you have to take the entire system into account.

Here, the reliability involved is that of the system and not only that of the algorithm. The algorithm does its work. The issue is to see how the data is being used, what types of decisions are made and what human control there is over those decisions or predictions. From that perspective, it seems extremely important to me that the algorithmic systems—not simply the algorithm—be audited. I'm talking about audits in the sense where people really look into the architecture of the system to find its possible shortcomings.

In the case of aircraft, since you mentioned those two recent tragic air catastrophes, we must, for instance, ensure in advance that human beings keep control, even if they may make mistakes. That is not the issue; human beings make mistakes. That is precisely why we could also put in place algorithmic aids. However, admitting that to err is human and that there is still human control over the machine—that is part of the things we need to discuss. However, this is certainly an essential factor if we are to identify the problems with a given algorithmic system.

• (1605)

[English]

**Hon. Peter Kent:** Thank you.

With regard to the social auditing of AI development and algorithms, we've seen the negative impact on assembly lines, for example, the negative impact on labour as some companies rush to either save their business plans or maximize new profits by relying on AI technology. How do both of you feel in terms of the acceptability or the necessity of a certain amount of social regulation to ensure societal stability in important areas of the labour force?

Professor Dilhac can go first.

[Translation]

**Prof. Marc-Antoine Dilhac:** It's very difficult to direct social change through laws or regulation. When we talk about a technological revolution, we have to take the mode of the revolution seriously. As we were saying, there is a structural change. It does not seem reasonable to me to want to direct a transformation of this nature. What does seem reasonable is to put in place training mechanisms so that the digital society transformation can include everyone who needs to change their skills.

The idea is not to put more pressure on businesses to prevent them from replacing human beings with algorithms. That may be regrettable, and I regret it, but it's not the best approach. The government could, however, put training mechanisms in place to support people in their quest to transform their skills.

[English]

**Hon. Peter Kent:** On technological—

**The Chair:** Sorry, we're way over.

Next up, for seven minutes, is Mr. Julian.

[Translation]

**Mr. Peter Julian (New Westminster—Burnaby, NDP):** Thank you very much, Mr. Chair.

I also thank our guests.

I apologize for arriving a bit late. I didn't get a chance to hear their presentations. I apologize in advance if I ask questions that have already been answered in the presentations.

My first question is about the legislation and the AI regulatory framework. Should Canada develop a regulatory framework to govern the ethical use of artificial intelligence? Can you name some countries where governments, either nationally or regionally, have put in place laws and regulatory frameworks for the use of artificial intelligence?

My question is for both of you, but Mr. Dilhac may answer first.

• (1610)

**Prof. Marc-Antoine Dilhac:** There are many regulatory initiatives throughout the world. The European Union has just produced a normative and ethical framework with recommendations and assessment lists. It is, however, rare that countries resort to legislation.

For certain activities and sectors, the legislative framework can be relevant. Canada can govern certain algorithm-related activities. I am referring to data, of course, since that is what feeds the algorithms. We need to adopt a law to create a regime conducive to governing the use of the data.

We also need to adopt laws with regard to education in order to determine up to what point we want education to be robotized. This is coming gradually, and when we are there, it will be a bit too late to legislate.

I'll give you an example. When a robot follows a student's journey, it makes decisions that will follow the student. I am referring here to "robots" because the interface is robotic, but those decisions are in fact based on algorithms. The question is: to what extent do we want to lock a student into a path where their progress is established or assessed constantly by an algorithm.

**Mr. Peter Julian:** Forgive me for interrupting you, but I don't understand. You are talking about a robot that follows a pupil or a student, but what is the robot doing in the example you gave?

**Prof. Marc-Antoine Dilhac:** You are asking me what the purpose of the robot is?

**Mr. Peter Julian:** Yes.

**Prof. Marc-Antoine Dilhac:** It makes it possible to personalize education. If a student is having difficulties, the algorithm can identify this very quickly and adapt the teaching content to the student. It's a big technological breakthrough. Canada and North America in general are still a bit behind with regard to this technology that is well established in Asia, in China and South Korea. I'm giving you examples of things that will be coming soon. It was on my mind because there was a summit on the use of digital technology in education in Montreal recently, attended by over 1,500 people.

I'll give you one last example. Who should make the decisions? It is up to lawmakers to decide, after a broad consultation, just as doctors are responsible for limiting risk when they make a diagnosis and prescribe treatment.

Those are just a few examples, but there are many others.



**Mr. Peter Julian:** Thank you very much.

Mr. Sandvig, what is your opinion?

[*English*]

**Prof. Christian Sandvig:** Well, I think I'm in sympathy with remarks made by my colleague.

What I can add is that it's hard to foresee specific legislation, in part because we don't have a good definition of what we mean by artificial intelligence. It's really a loose term that covers all kinds of different things. Even ideas within it that we're particularly concerned about, like machine learning...that term is itself a loose term that covers a variety of approaches that are quite different.

One of the challenges for us is the success of computing, because it has meant that things that look like artificial intelligence are all kinds of things, and they are in all kinds of domains. I think it's more likely that we will see legislation that specifically addresses a context and a use of technology, as opposed to an overarching principle.

A colleague of mine said that we are at "peak white paper". We might be near peak principles as well. There are many statements of principles, and these are valuable. However, I think our task is to translate these into specific situations rather than to legislate all of AI, because I just don't know how to do it. There are some exceptions, though. There are a few areas where we might see overarching legislation that's of value.

One example would be that this committee has done some important work on the Cambridge Analytica scandal with its previous report. One of the challenges of that scandal for many countries around the world was that they had taken an approach to communication that said social media platforms essentially do nothing. Many governments, as you know, provide immunity to liability for online platforms or social media companies as conduits.... They did that in a very blanket way. We could say it's a terrible mistake of the United States.

This is an area where you have one legislation that affects a huge swath of activity, because it affects all use of computers to act as intermediaries or conduits between humans. The idea that you would give away freedom from liability seems like a bad one.

There are some areas where there could be broad legislative action, but I think they're rare. It's more likely that we'll see domain-specific approaches.

•(1615)

**Mr. Peter Julian:** Thank you.

**The Chair:** Next up, for seven minutes, Mr. Erskine-Smith.

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

Thank you to you both.

The last answer is a useful segue in terms of regulation that could apply more broadly. We obviously have in the GDPR algorithm explainability, and a right to explainability was referenced in some of the opening comments.

This committee has recommended a level of transparency and the ability of a regulator to look under the hood at times to assess whether that transparency has been sufficient.

Mr. Sandvig, you have expressed some skepticism about transparency, though that does appear, to me at least, to be an initial step that would apply more broadly, in the way that the next step might have to be sector specific.

I want to drill down on some of the skepticism with respect to transparency. You didn't mention algorithmic impact assessments in your opening comments. I wonder if the detailed work that is now being put into formulating AIAs is a better answer to transparency.

**Prof. Christian Sandvig:** I'm going to remain skeptical about transparency because I think that algorithmic impact assessment isn't a transparency proposal. I think that those proposals, as their title implies, owe a debt to environmental impact assessment. There may be elements of transparency required in producing such an assessment, but I think that I didn't mention them in that section in part because I don't see them as predominantly a transparency approach.

I'd be happy to give you additional skepticism about algorithmic impact assessments, though. The challenge of them, for me, is that we might group the negative harms of algorithms and AI into two groups. One group we could say is foreseeable, and one group we could say is not foreseeable. I'm afraid that the second group is quite large. The algorithmic impact assessment stuff that I've seen really takes for granted that it's possible to have some assessment. When we look at many of the scandals involving computer systems, artificial intelligence and algorithmic systems, a number of the scandals—although not all—seem to involve things that no one would have wanted. It could be that an impact assessment process caused or required the developer to think more carefully about the system and to produce a different one, but it might also be that some of the results that we're seeing are hard to imagine as being foreseen at all. I just worry about it.

**Mr. Nathaniel Erskine-Smith:** Let me jump in.

You have to accept that there is a transparency aspect to this. I'll use an example. In the public sector at the moment—and this is very recent for the Government of Canada—there is a questionnaire that any department that is employing automated decision-making needs to fill out. It's 80-some-odd questions. Based upon the answers to those questions, they're assigned basically a level 1, level 2, level 3 or level 4 in terms of the risk.

Then there are measures that need to be taken, some additional notice requirements. They have to obtain experts who peer review the work, but in the initial impact assessment itself, there are questions about the purpose of the automated decision-making that they intend to employ and the impact that it's likely to have on a particular area, such as individual rights, the environment or the economy.

We could argue about the generality of it and whether this could be improved, but it seems, on one hand, to provide a transparency mechanism in that it is requiring a disclosure of the purpose of the algorithm and potentially the inputs to the algorithm, its benefits and costs, and the potential externalities and risks. Then, depending on the outputs to that assessment, there are additional accountability mechanisms that could apply.

If you haven't looked at it yet, my question would be this: If and when you do take a look at the Canadian model for the public sector in more detail, is that something that you could transcribe and treat more like a securities filing—that is, to say “this is going to be required for private sector companies of a certain threshold, and if there is any non-compliance where material terms are excluded purposefully or in a negligent way, then there are penalties for non-compliance”? Would that be sufficient to meet at least the baseline of transparency accountability generally before we get into sector-specific regulations?

• (1620)

**Prof. Christian Sandvig:** I absolutely will agree that there is a role for transparency somewhere. I'm just afraid of it as a proposal because it promises, I think, more than we can expect it to.

**Mr. Nathaniel Erskine-Smith:** That's fair.

**Prof. Christian Sandvig:** So, I agree with that.

Let's look at a particular example. With regard to some of the algorithms that this committee was concerned about when writing its prior report, we know that there are already patents available that give broad outlines as to how the algorithms work.

For example, look at the Facebook newsfeed. Facebook—in public disclosures that have already been made—used to brag that the computation was made based on three factors. As the years went by, it said that it was based on dozens of factors, and then it said that it was based on hundreds of factors. I think that we're now at over 300 factors. There's some value in disclosing these factors, but it's not clear that there's that much because—

**Mr. Nathaniel Erskine-Smith:** But I guess my point is that it wouldn't be just about disclosing the factors. When we had Richard Allan, the VP for global policy, at our international committee in London last fall, he said, you know, if speech crosses the threshold for hate, obviously we should take it down, but if it's right up against that line, maybe we shouldn't encourage and promote it. And I'm sitting there thinking, yes, obviously you shouldn't promote that kind of content, but that's the algorithm. That's the newsfeed algorithm to promote reactions, regardless of what those reactions are. Even if they're negative reactions, they're looking for eyeballs. They're not looking for much beyond that when they want to generate profit.

If there is an algorithmic impact assessment and we are setting the rules of what that assessment should entail, I agree with you that there's an element of transparency and disclosure. It shouldn't just be about the inputs, necessarily. A company should have to come to terms with what the potential adverse affects are as well, I think, and have to put that in such an assessment. They have to turn their minds to that.

Do you think that is a useful and additional layer of accountability and transparency?

**Prof. Christian Sandvig:** Yes. I absolutely do. I just worry—I don't know how large the category is—that they won't be addressed by it. This is why I'm skeptical. It's not that it isn't a valuable proposal in itself.

**Mr. Nathaniel Erskine-Smith:** Thanks very much.

**The Chair:** Normally we have lots of time for these discussions, but I have to give everybody notice that we have time for only two more five-minute questions. Then we have to get into the discussion about the legal advice on the summons. I just want to forewarn you about that.

We have Monsieur Gourde for five minutes, and then Monsieur Picard.

• (1625)

[*Translation*]

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

The ethical difficulty all these new platforms raise is due to the fact that individuals are not all necessarily aware of what artificial intelligence means and the degree of acceptance may vary considerably.

I'll give you an example. Personally, it does not bother me that thanks to algorithms, my favourite colour or the brand of my car is known. However, some people are very reluctant and absolutely do not want anyone to know anything about them on these platforms. Unfortunately for them, I think it is already too late. Private businesses acquire a lot of personal information about us. Their data bases grow annually and they can practically predict the date of our death by using algorithms.

Mr. Dilhac, you said that there were a lot of regulations but very few laws in this area. I can understand that it's quite difficult to adopt laws to manage tools that don't respect borders. Nevertheless, as legislators, we have to protect the population. Where should we start?

**Prof. Marc-Antoine Dilhac:** You have to find a balance between the law and the contract. Today, the contractual mode is predominant when it comes to deploying artificial intelligence in applications. When you click on a button at the end of the contract on Facebook, you accept or not. You don't have time to read it.

If you look at the content of these contracts, you see that they contain totally unacceptable elements that should not be there, and I'll take Facebook as an example. We examined the conditions of use of Facebook a little. That enterprise gives itself the authority to obtain your information through third-party applications.

Whether or not you are online using Facebook, whether or not you have registered with them, the enterprise has given itself the right to go and get information about you from other applications. That type of thing is entirely possible through the use of the contract form. If, as a user, you accept that, well it's too bad for you. That kind of contract should be regulated by law. That is precisely where a balance needs to be found. It isn't easy but it is the government's work to find that balance between what should be in a contract between a service provider and a user, and what should be in the law.

What is the priority? There are a lot of things that need to be done, but I think that in order to protect the public, your main, most serious priority should be the use that is made of the data. It isn't just the fact that you like the colour blue that is important, but if one day you no longer like it, an algorithm may come to the conclusion that you have a mental problem or a disease you don't know you have, for instance, and that will be much more troublesome.

**Mr. Jacques Gourde:** Thank you.

[English]

**The Chair:** Thank you, Mr. Gourde.

Next up is Mr. Picard. That will close this off.

Go ahead, Mr. Picard.

[Translation]

**Mr. Michel Picard (Montarville, Lib.):** Thank you.

My question is for Mr. Dilha, and then, Mr. Sandvig can answer.

The AI revolution comes with its share of unknowns, eliciting a negative reaction from the public. People fear the worst and conjure up all the bad things that could happen. Nevertheless, we lived through the Industrial Revolution at the beginning of the last century. In the 1960s and 1970s, we went through the electronic revolution, which gave us computers.

All things being equal, aren't all three events comparable in terms of their cultural, social, economic and political effects on society? Aren't there lessons we can learn, both positive and negative ones? Conversely, is it not possible to draw lessons because the three revolutions are so different?

**Prof. Marc-Antoine Dilhac:** I'll try to keep my answer brief.

Yes, AI does come with unknowns. A modest stance would be to say that we don't quite know where we are headed. If we look at the past, we can find guideposts. You brought up the Industrial Revolution, which led to major advancements. However, the revolution occurred in the early 19th century—two centuries ago—without any groundwork being laid. It gave rise to more than a century of torment, more than a century of transitions and war, not to mention revolutions and, all told, millions of deaths. Government was completely overhauled.

If the Industrial Revolution taught us anything, it's that we need to address the period of transition that comes with technological advancement and new tools. Economist Joseph Schumpeter, whom you're probably familiar with, coined a relevant expression. He talked about the destructive transition, better known as creative destruction, meaning that something is destroyed in order to create new economic activities. Creative destruction can take a long time, and the destructive aspect is not necessarily appealing.

It's important to focus on the conditions for transition so that there are as few losers as possible. AI and the use of algorithms leads to tremendous progress, not just in medicine, but also with respect to repetitive tasks. That is something we should welcome, but we also need to prepare for the revolution.

• (1630)

**Mr. Michel Picard:** Mr. Sandvig, would you care to comment?

[English]

**Prof. Christian Sandvig:** I will defer to my colleague's assessment. I think it's a big question.

**Mr. Michel Picard:** Okay. For my second question, then, who will decide what will be the norms when you create such an AI software system? If the norms are established by the corporate entities, we're going to end up in a business-type society. I don't think that's the way we want to go, so a wise man or a wise woman somewhere who everyone should recognize.... I haven't found one yet, not in government. Who is it going to be?

Mr. Sandvig.

**Prof. Christian Sandvig:** I don't have a specific answer for you, to be honest. It's exciting to me that you're asking that question, because I worry that many people believe that they don't see an alternative to technologies that somehow come out of nowhere and they are then subject to, so I'm excited by the idea that these are not corporate and that these are things that we all have to decide as a society.

As to how to achieve that practically, I think this is quite challenging. We can endorse the principle of “democratic participation”, for example, as given in the Montreal statement. How do we achieve that? There are some models. There's the Scandinavian model of participatory design. There are ideas, but still, currently, I think we look at a landscape that is dominated by what have been called big social monopoly computing platforms—AI—and it's hard to see exactly how you will have a voice in it.

I'm hopeful that some of the proposals I discussed in my opening statement, and that previously were in your last report about structural changes to the industry, might provide openings for which we see another kind of participation—for example, a public alternative.

[Translation]

**Mr. Michel Picard:** Mr. Dilhac, do you have any final comments?

**Prof. Marc-Antoine Dilhac:** In terms of norms, I tend to come down on the negative side. Making sure the norms don't come solely from the private sector and industry is paramount, but that's increasingly what we are seeing. It's called self-regulation. I'm not saying all the answers have to come from government—that would involve identifying the problem, first and foremost, which is no small feat—but government does need to take back control of the conversation around norms. It must ensure that businesses, professional bodies and civil society groups contribute to the conversation. I think government needs to assume control of the debate, which is why I'm so glad to be participating in this kind of forum.

• (1635)

**Mr. Michel Picard:** Thank you.

[English]

**The Chair:** Thank you, Mr. Picard, and thank you to our witnesses. I know it's something we could talk about for a lot longer than an hour. We appreciate your time today.

We're going to briefly suspend while we go in camera with Mr. Dufresne.

Again, thank you. I appreciate your contributions to our committee. *[Proceedings continue in camera]*

---







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>