



HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

Standing Committee on Access to Information, Privacy and Ethics

ETHI • NUMBER 145 • 1st SESSION • 42nd PARLIAMENT

EVIDENCE

Tuesday, April 30, 2019

—
Chair

Mr. Bob Zimmer

Standing Committee on Access to Information, Privacy and Ethics

Tuesday, April 30, 2019

• (1540)

[English]

The Chair (Mr. Bob Zimmer (Prince George—Peace River—Northern Rockies, CPC)): We will call to order meeting number 145 of the Standing Committee on Access to Information, Privacy and Ethics. Pursuant to Standing Order 108(2), this is the study of the ethical aspects of artificial intelligence and algorithms.

First of all, we'll go to Mr. Angus who has a motion.

Go ahead, Mr. Angus.

Mr. Charlie Angus (Timmins—James Bay, NDP): Thank you. I won't take too much time.

I have a motion, but can I make a commercial announcement first?

I have no pecuniary interest in this, but I implore my colleagues on the committee to watch the BBC television show *Brexit*. Our friend Zack Massingham makes an appearance as one of the central characters, and he certainly does not appear to be as tired and confused and memory-fogged as he did to us.

Mr. Raj Saini (Kitchener Centre, Lib.): Is it a show?

The Chair: Yes. As chair, I'll send out the link.

I uploaded it on iTunes, so get it however you wish.

Mr. Charlie Angus: I think it would be great to have Mr. Zack Massingham back to ask how he considers his portrayal....

The Chair: He was much more involved than he let on.

Mr. Charlie Angus: Much more involved.

Okay, I brought a notice of motion to the committee:

That the Committee begin a study on the ethical aspects of artificial intelligence and algorithms.

This was in response to our clerk, who said that in order to undertake this next round of witnesses, we needed an official motion.

The Chair: I'll speak quickly to that, Mr. Angus, to talk about the aspects of the study, and then I'll pass it on to the clerk. Our calendar is very packed for what we're able to pull off. The analysts are wondering how far we go with this. Do we report back with a report?

Mr. Clerk, go ahead.

The Clerk of the Committee (Mr. Michael MacPherson): I can send around a calendar after tonight's meeting. It will make it clear what we can fit in.

It's basically defining the priorities for the committee. Do you want a report on the privacy of digital government services as well as this new study that you're embarking on?

The Chair: Is there any debate?

Mr. Charlie Angus: I'm very interested in following up on some of the discussions we've had in terms of the effect of how the algorithms in certain platforms are being used to distort public conversation and political discourse by moving people toward more and more extremist and false content, as opposed to being able to find accurate, credible sources.

I think we have not really looked at some of those algorithms, particularly with YouTube—we have put a lot of attention into Facebook—but these algorithms are having a very distinct impact on civil discourse. It's worth knowing how they work. Of course, with the larger issues we talked about of AI and algorithms, I leave it to my colleagues around the table if they feel other witnesses should be drawn, but I think we have a pretty good list of witnesses and not much time.

I'd say let's get down to it.

The Chair: Mr. Kent.

• (1545)

Hon. Peter Kent (Thornhill, CPC): The topic is worthy of a full study, and I think we're going to be talking about elements of that today in testimony from our witnesses. I think we should get to that.

However, given that we have barely six weeks of meaningful committee time left, I'm not sure we could get a formal study itself going. Certainly as we go through discussion of the outline of the draft report on digital government, if there are opportunities, then the more testimony we can hear, the better off the next parliament will be to really take it on seriously and chew on it.

The Chair: Mr. Erskine-Smith.

Mr. Nathaniel Erskine-Smith (Beaches—East York, Lib.): Very briefly, obviously on May 28, we have a much larger committee hearing, with international colleagues coming from other countries. Part of that conversation is going to be about algorithmic accountability and transparency.

Whether or not it leads into a more fulsome report in any way—we might run out of time, and that's fine—I agree with Mr. Kent that regardless, it's worthwhile for us to hear the evidence. I think more than anything for us now, it's worthwhile leading into May 28 to listen to some of the experts. We might have some more pointed questions.

The Chair: That sounds good. Is there any further discussion on that?

Mr. Angus.

Mr. Charlie Angus: Well I'm a firm believer in always having reports so we can show what our committees have done. I recognize that the clock is ticking, so I'm willing to bend on that.

I agree with Mr. Erskine-Smith. This is about setting us up for the international grand committee so that we are fully prepared; we've had a chance to look at some other issues that we may bring to the table. To me, this is a good training session leading up to that.

Then, out of that committee, there may be an international statement, or we may feel the need to follow up with a further report. I will take that after we have the international committee and find out what our colleagues around the world think.

The Chair: Okay. Is there any more debate or any discussion? Seeing none, all in favour of the motion?

(Motion agreed to)

The Chair: That's unanimous.

Thank you, Mr. Angus.

We'll get on to business. We have two witnesses here with us today: Mr. Ben Wagner, assistant professor, Vienna University of Economics, by teleconference; and Yoshua Bengio from Mila-Quebec Artificial Intelligence Institute. Dr. Bengio is the scientific director there and is here by teleconference from Montreal.

We'll start with you, Mr. Wagner. Go ahead for 10 minutes.

Professor Ben Wagner (Assistant Professor, Vienna University of Economics, As an Individual): Thank you very much for the opportunity to speak here. I really appreciate the standing committee dealing with these issues. My name is Ben Wagner. I'm with the Privacy and Sustainable Computing Lab in Vienna.

We've been working closely on these issues for some time, specifically trying to understand how to safeguard human rights in a world where artificial intelligence and algorithms are becoming extremely common. This has included helping prepare Global Affairs Canada for the G7 last year. It was a great pleasure to work with colleagues there like Tara Denham, Jennifer Jeppsson and Marketa Geislerova.

The results that were produced in that, I think, are quite relevant also for this committee. You have the Charlevoix common vision for the future of artificial intelligence. Related to that, last year we were also working on—this is now in a Council of Europe context—a study on the human rights dimensions of algorithms, which I also think would be extremely helpful, especially if you're discussing studies and common challenges faced. Many of the common challenges you're discussing are already mentioned in these G7

documents and also in the statements developed by the Council of Europe.

To come back to a more general understanding of why this is important, artificial intelligence or AI is frequently thought of as some unusual or new thing. I think it's important to acknowledge that this is not a new and unusual technology. Artificial intelligence is here right now and is present in many existing applications that are being used.

It's increasingly permeating life-worlds, and it will soon be difficult to live in the modern world without having AI touch your life on a very daily basis. Its deep embedding in societies of course poses considerable challenges, but also opportunities. I think when we specifically look at the ethical and regulatory dimensions as, I believe, this committee is doing, it's extremely important to remember to try to ensure that all citizens have access to the opportunities of these technologies and that the opportunities provided by these technologies are not limited to just a select few.

With regard to how that can be done, there is a variety of sets of challenges and different issues. One of the most common ones is whether we talk about ethical framework or a more regulatory governance framework. I think it's important that they not be played off against each other. Ethical frameworks have their place. They're extremely important and they're extremely valuable, but of course they can't override or prevent governance frameworks from functioning. Indeed it would be difficult if they could. But if they function in parallel in a useful and sustainable manner, that can be quite effective.

The same is true even if you take a more governance-oriented human rights-based framework. It's very frequent that in these contexts different human rights are played off against each other. The right to freedom of expression is seen as being more important than the right to privacy. The right to privacy is seen as being more important than the right to free assembly, and so on. It's very important that in developing standards and frameworks in this context, we always consider all human rights and that human rights be the basic foundation for how we think about algorithms and artificial intelligence.

If you look at the Charlevoix documents that were developed last summer, you'll also note a considerable focus on human-centric artificial intelligence. While that's an extremely important design component, I think it's also important to acknowledge that human-centric focuses alone are not enough. At the same time, while we're seeing an increasing number of automated systems, lots of actors who are developing automated systems are not willing to admit how they're actually developing them or what exact elements are part of these systems.

It's often joked that some of the most frequently used examples in the start-up business plans of artificial intelligence are closer to Mechanical Turk—that is to say human labour—than to actual advanced artificial intelligence systems. This human labour often gets lost on the way or fails to be acknowledged.

This is also relevant in the context of extra-legal frameworks that are frequently applied when we talk about ethical frameworks, when we talk about frameworks that don't govern in the way that rule of law can. I think we need to be extremely careful there with regard to the extent to which frameworks like this actually come to replace or override the rule of law. That's specifically also the case where we see lots of conversations right now. I'm sure you will have heard about Google's AI board, which was recently created and then shut down within the space of just a week or two.

You'll notice that there's an attempt on the one hand, a great push by some actors, to try to be more ethical, but this ethical framework is not enough and the actors realize this, given the heavy criticism of this that you see, which again isn't to say that ethics isn't important or ethics is necessary but that ethics needs to be done right if it's going to have a meaningful impact on this. That means there's a strong role for the public sector as well. We can't allow ethics squashing. We can't allow ethics shopping. We can't allow for lowering the bar for the standards that we already have.

As I'm sure you are aware, the existing standards in many areas of public governance—when we're talking about existing norms related to how we govern technology and how we govern the activities of corporations, if you look at the business and human rights framework of the United Nations, for example—are already relatively weak. In some areas, there's a danger that these ethical principles will even go below existing business and human rights standards.

● (1550)

At the same time, to take a more positive note as well, there is an extremely important role for the public sector here, and I think it's again possible to commend the work specifically of Michael Karlin, who has done some fantastic work on algorithmic impact assessments for the Government of Canada. There's really an important measure to be seen there in how Canada is also taking a lead and really showing what is possible in the context of these algorithmic impact assessments. I can definitely commend his work there.

At the same time, when you look at the recent accusations now that Facebook has been breaking Canadian privacy laws, we have a serious issue related to implementation. Specifically, these breaches that have been of concern to numerous Canadian privacy regulators do raise a question. Can we just focus on the public sector alone and can the public sector alone lead the way, or do we need to take similar considerations for, at the very least, large, powerful private sector companies? Because in the world we live in right now, whether you're talking about opening a bank account, posting something on Facebook, talking to a friend online or even getting a pizza delivery, algorithms and AI are part of every step that takes place in that context.

Unless we're willing to limit the agency of these algorithms, there are two things when we consider those things democratically relevant. They increasingly begin to dominate us, and this is not a *Terminator*-like scenario where we need to be scared that the robots will come and take over the world.

It's rather that, through these technologies, a lot of power becomes concentrated in the hands of very few human beings, and these are

precisely the types of situations that democratic institutions, such as the parliamentary committee that's hearing about this topic right now, were built to deal with. That is to ensure that the power of the few is spread to the power of the many, and to ensure that having access to AI and to the benefits of AI, but also to the foundational promise of AI that technology can make people's lives better, both inside Canada and beyond, is accessible to every human being, and that basic human rights provide the core foundation of how we develop and how we think about technology in the future.

Thank you very much for listening. I look forward to answering any questions you might have.

● (1555)

The Chair: Thank you.

We'll go next to Mr. Bengio in Montreal.

Go ahead, please, for 10 minutes.

Professor Yoshua Bengio (Scientific Director, Mila - Quebec Artificial Intelligence Institute): Hello. My expertise is in computer science. I've been a pioneer of deep learning, which is the area that has changed AI from something that was happening in universities into something that is now taking a big economic role and where there are billions of investments in industry.

In spite of the progress that's remarkable, it's also important to realize that the current AI systems are very far from human-level AI. In many ways they are weak. They don't understand the human context, of course. They don't understand moral values. They don't understand much but they can be very good at a particular task and that can be very economically useful, but we have to be aware of these limitations.

For example, if we consider the application of these tools in the military, a system is going to take a decision to kill a person and doesn't have the moral context a human can have to maybe not obey the order. There's a red line, which the UN Secretary-General has talked about, that we shouldn't be crossing.

Going back to AI and Canada's role, the thing that is interesting is we've played a very important role in development of the recent science of AI and clearly we are recognized as a scientific leader. We also are playing a growing role on the economic side. Of course, Canada is still dwarfed in comparison to Silicon Valley, but there is a very rapid growth of our tech industry regarding AI and we have a chance, because of our strength scientifically, to become not just a consumer of AI but also a producer, which means Canadian companies are getting involved and that's important to keep in mind as well.

The thing that's important, in addition to the scientific leadership and our growing economic leadership regarding AI, is moral leadership, and Canada has a chance to play a crucial role in the world here. We have already been noticed for this. In particular I want to mention the Montreal declaration for responsible development of AI to which I contributed and which is really about ethical principles.

Ten principles have been articulated with a number of sub-principles for each. This is interesting and different from other efforts in trying to formalize the ethical and social aspects of AI because in addition to bringing in experts in AI, of course there were scholars in the social sciences and humanities, but ordinary people also had a chance to provide feedback. The declaration was modified thanks to that feedback with citizens in libraries, for example, attending workshops where they could discuss the issues that were presented in the declaration.

In general for the future, I think it's a good thing to keep in mind that we have to keep ordinary people in the loop. We have to educate them so they understand issues because we will take decisions collectively, and it's important that ordinary people understand.

When I give talks about AI, often the biggest concerns I hear are about the effect of AI on motivation and jobs. Clearly, governments need to think about that and that thinking must be done quite a bit ahead of the changes that are coming. If you think about, say, changing the education system to adapt to a new wave of people who might lose their jobs in the next decade, those changes can take years, can take a decade to have a real impact. So it's important to start these things early. It's the same thing if we decide to change our social safety net to adapt to these potential rapid changes in the job market. These things should be tackled fairly soon.

I have another example of short-term concerns. I talked about military applications. It could be really good if Canada played more of a leadership role in the discussions that are currently taking place around the UN in the military use of AI and the so-called “killer drones” that can be used, thanks to computer videos, to recognize people and target them.

• (1600)

There's already a large coalition of countries expressing concern and working on drafting an international ban. Even if not all the countries—or even major countries such as the U.S., China or Russia—don't go with such an international treaty, I think Canada can play an important role. A good example is what we did in the nineties with anti-personnel mines and the treaty that was signed in Canada. That really had an impact. Even though countries such as the U.S. didn't sign it, the social stigma of these anti-personnel mines, thanks to the ban, has meant that companies gradually have stopped building them.

Another area of concern from an ethical point of view has to do with bias and discrimination, which is something that is very important to Canadian values. I think it's also an area where governments can step in to make sure there's a level playing field between companies.

Right now, companies can choose to use one approach—or no approach at all—to try to tackle the potential issues of bias and discrimination in the use of AI, which comes mostly from the data that those systems are trained on, but there will be a trade-off between their use of these techniques and, say, the profitability or the predictability of the systems. If there is no regulation, what's going to happen is that the more ethical companies are going to lose market share against the companies that don't have such high standards, and it's important, of course, to make sure that all those companies play on the same level.

Another example that's interesting is the use of AI not necessarily in Canada but in other countries, because these systems can be used to track where people are by, again, using these cameras all over the place. The surveillance systems, for example, are currently being sold by China to some authoritarian countries. We are probably going to see more of that in the future. It's something that is ethnically questionable. We need to decide if we want to just not think about it or have some sort of regulation to make sure that these potentially unethical uses are not something that our companies are going to be doing.

Another area that's interesting for government to think about is advertising. As AI becomes gradually more powerful, it can influence people's minds more efficiently. In using information that a company has on a particular user, a particular person, the advertising can be targeted in a way that can have much more influence on our decisions than older forms of advertising can. If you think about things like political advertising, this could be a real issue, but even in other areas where that type of advertising can influence our behaviour in ways that are not good for us—with respect to our health, for example—we have to be careful.

Finally, related again to targeted advertising is the use of AI in social networks. We've seen the issues with Cambridge Analytica and Facebook, but I think there's a more general issue about how governments should set the rules of the game to minimize this kind of influencing by, again, using targeted messages. It's not necessarily advertising, but equivalently somebody is paying for influencing people's minds in a way that might not agree with what they really think or what's in their best interests.

Related to social networks is the question of data. A lot of the data that is being used by companies like Google and Facebook, of course, comes from users. Right now, users sign a consent to allow those companies to do whatever they want, basically, with that data.

• (1605)

There's no real potential strength for bargaining between a single user and those companies, so various organizations, particularly in the U.K., have been thinking about ways to bring back some sort of balance between the power of these large companies and the users who are providing data. There's a notion of data trust, which I encourage the Canadian government to consider as a legal approach to try to make sure the users can aggregate—you can think of it like a union—where they can negotiate contracts that are aligned with their values and interests.

The Chair: Thank you.

We'll get to questions.

I just want to recognize that we have a special guest and his class with us today. Professor Michael Geist, I thank you for attending. You could probably appear at the same panel with us today, but you're going to take the easy road today and listen in.

Welcome, students.

Mr. Charlie Angus: We also have students from the University of Haifa.

The Chair: They're from Haifa, so we have students from across the water.

Mr. Charlie Angus: That's east of St. John's, I believe.

The Chair: Just a little.

Thank you for coming today.

We'll start off with Mr. Erskine-Smith for seven minutes.

Mr. Nathaniel Erskine-Smith: Thanks very much.

I want to talk more about regulation than ethics, particularly because of the most recent example where Facebook has said to our Privacy Commissioner, "Thanks for your recommendations; we're not going to follow them", so I think we need stronger rules as far as they go.

Mr. Wagner, in a recent article, one of the three examples you use about AI is social media content moderation. At this committee we've talked about algorithmic transparency. In the EU it's algorithmic explainability. In that article you noted that it's unclear what that looks like. It's a new idea, obviously, in the sense that, when we've spoken to the U.K. information commissioner and had recent conversations with the EU data protection supervisor, they are just scaling up their capacity to address this issue and to understand what this looks like.

Having looked at this issue yourself and written about this, when we talk about algorithmic transparency, is there a practical understanding that we ought to have? It's one thing to make an recommendation on algorithmic transparency. What should it specifically look like?

Prof. Ben Wagner: It's an extremely good question. At this point there are quite a lot of proposals out there on what it could be, but I think the first thing, to come straight to the point, is that transparency or explainability itself is insufficient. Just saying we can explain what it does, and therefore it's enough, is not enough. You have to have someone who's in a meaningful way accountable for the actions of these things, and you need a governance framework around it.

When we're talking about, especially in the context of social media, having a framework for how content is moderated, it also means appeal mechanisms, transparency mechanisms and ensuring that there is some kind of external adjudication if there is a disagreement in these contexts, and then adding an extra layer of complexity when we're talking about regulatory responses to this.

There is a challenge that once AI-type systems or automated systems have been embedded within organizations, over time those organizations become dependent on those systems, and it's very difficult to move beyond them or get out of them, so you need to be quite strong on the governance quite early to make sure that you're really having a strong and meaningful effect on how—

Mr. Nathaniel Erskine-Smith: I want to get to what more it could be. With respect to explainability and transparency, you mentioned Karlin here in Canada, and you referenced, too, what the Treasury Board has done with respect to algorithmic impact

assessments on the public sector side. It occurs to me that, if we are serious about that level of transparency and explainability, it could mean a requirement for algorithmic impact assessments in the private sector akin to an SEC filing where non-compliance would come with some sanctions if information is not included. Do you think that is the level we should aim for?

• (1610)

Prof. Ben Wagner: Yes. In principle, I think that's exactly where things should be going. That's exactly the type of proposal that I was trying to suggest as to where things should be moving. What I would add is that, of course by doing so, you don't want to stifle innovation, so you would need some kind of threshold on top of which, let's say for publicly traded companies or for companies of a certain size, there's a certain impact. Now, of course, depending on the amount of data those companies hold, those can also be very small companies, so you would have to have different types of thresholds for different types of organizations. Yes, I think that would be extremely helpful.

Mr. Nathaniel Erskine-Smith: Mr. Bengio, you talked about bias and discrimination. You talked about advertising, the ability to influence more efficiently, and the use of AI in social networks. Each time, I think you were hinting at something. I mean, with respect to bias and discrimination, you explicitly hinted at the need for regulation, or you suggested the need for regulation.

Prof. Yoshua Bengio: Yes.

Mr. Nathaniel Erskine-Smith: If it's not just an ethical framework...and I appreciate the work you've done with the Montreal declaration on ethics, but if we're talking regulation, is there something you would point this committee to in terms of how we ought to regulate algorithmic decision-making to solve some of these problems that you've identified?

Prof. Yoshua Bengio: Yes. I'm not a legal expert, so there might be different ways that one could regulate. In some cases, maybe even current laws are sufficient and they need to pass the test of the courts. Let me give you an example in the case of bias and discrimination. Let's say you consider the insurance industry. You probably would need different regulations for different industries where the way in which issues come up might be different. In the case of insurance, there could be information that is used by the companies that could lead to, say, gender discrimination. Even though the variables used by the insurance company do not explicitly mention gender, or do not explicitly mention race, it might be something that the AI system infers implicitly. For example, if you live in some neighbourhood, maybe it's a good indication of your race in some places.

The good news is that the algorithms that can mitigate this exist, but there will be a trade-off between eliminating the implicit information about gender and the accuracy of the predictions made by those systems. Those predictions turn into dollars. For an insurance company, if I can make a very precise assessment of your risk, of how many dollars you will cost me, that is how I will determine your premium, so that precision is really worth money. There will be pressure from companies to use as much information as they can from their customers, but it might go against our legal principles. We need to make sure we find the right trade-off.

Mr. Nathaniel Erskine-Smith: I want to pick it up there, because I think what it gets at is that we have existing rules from a human rights perspective.

Prof. Yoshua Bengio: Yes.

Mr. Nathaniel Erskine-Smith: In a way, the reason transparency becomes the first step is that it's so hard to enforce any of these rules until a human rights commissioner can adequately assess what is going on. When we were asking questions of the information commissioner in the U.K. in November, her view was that her job was to make it explainable. Other regulators have other rules and perspectives and rights and values that they want to enforce, and it's then their job to take on their roles.

Is that the sense you get? Is that the right approach?

Prof. Yoshua Bengio: That's a first step. We need to have clarity on how these processes are being put in place by companies like insurance companies, which use data to make decisions about people. We need to have some sort of access to that. It's understandable that they might want some secrecy, but government officials should be able to look into how they do it and make sure that it agrees with some of these principles that we put into law or in regulations or whatever. It doesn't mean that the system needs to explain every decision in detail, because that's probably not reasonable, but it's really important that they document, for example, what kind of data was used, where it came from, the way in which the data was used to train the system, and under what objective it was trained so that an expert can look at it and say that, for example, it's fine, or that there is a potential issue of bias and discrimination and maybe you should run such-and-such test to verify that there isn't; if there is an issue, then you should use one of the state-of-the-art techniques that can be used to mitigate the problem.

• (1615)

The Chair: Thank you, Mr. Erskine-Smith.

We'll go next to Mr. Kent for seven minutes.

Hon. Peter Kent: Thank you, Chair.

Thanks to both of you for appearing before the committee today.

Professor Bengio, you're to be congratulated for the work that you did on the Montreal declaration for responsible development of artificial intelligence, but as this committee has learned and as the public is—I hope—increasingly aware, much of the development of artificial intelligence has been funded by the “data-opolies”, by the Facebooks, by the Googles and by the increasing notoriety, as we learn, of their disregard for written and unwritten ethical guidelines and laws.

Just in passing, and you may not be aware of it, when this committee visited Facebook's headquarters in Washington last year, we were told almost in a passing comment that, when we asked if the company would accept increased regulation in Canada, the sort of investment we made in the AI hub in Montreal might not continue to be forthcoming, which hit me like a clunker. It was basically a threat from a “data-opoly” that Canada would be ostracized from AI investment should we increase regulation, even along the lines of the EU's GDPR or elements of it.

The question is to both of you. Large companies are already using and exploiting artificial intelligence in a variety of very commendable, wonderful ways, but also, in any number of ways that disregard ethical and legal guidelines. Should they be responsible for the misuse or the abuse of AI that occurs on their platforms?

Prof. Yoshua Bengio: Let me go to your first question about the Facebook investment in Montreal.

Our AI research centres are mostly funded by the provincial and federal governments right now: Mila in Montreal, Vector in Toronto, Amii in Edmonton. The investment that was made by, say, Facebook to create a lab in Montreal or to be a sponsor of organizations like Mila, is pretty small in comparison to the other investments that are happening.

I'm not really concerned. Facebook and other organizations have opened shop here because they see their interest in it. It makes it easier for them to recruit people they need for their research groups.

Hon. Peter Kent: You don't see as improper a threat or any possibility of partnerships exposing the AI that's being developed across the various partners to the hub by these very large companies?

Prof. Yoshua Bengio: No, they're doing a bit of research here and it's a small part of the bigger pot of research they're doing worldwide, which is somewhat disconnected from their actual business. Unless they want to use threats in an inappropriate way, to pull out of Canada right now would be to their disadvantage.

The other thing is, the investment they made in these other companies is still pretty small compared to the magnitude of the impact that we're talking about for all Canadians. Of course, we would be sad to see them go, and I don't think they would go, but I don't think we should even pay attention to this kind of statement.

Hon. Peter Kent: Okay.

Professor Wagner?

Prof. Ben Wagner: I think threats of this kind are quite indicative of a general regulatory challenge, which is that every country wants to be the leading country on AI right now, and that doesn't always lead to the best regulatory climate for the citizens of those countries.

There seems to have been some kind of agreement between the Government of the Netherlands and the automobile industry which is developing AI into self-driving cars to not look so closely when they build a factory there, in order to ensure that as a result of building that factory, they will bring jobs and investment to these countries.

I think that the impact of AI and these technologies will be sufficiently transformative that while these large U.S. giants seem quite important right now, that may not be the case in a few years' time. A lot of the time I think the danger is actually the other way around. The public sector has historically invested a lot more than many people are aware of, and a lot of the fortunes of these large well-known companies are based on that. Of course, in political terms, it always looks more attractive to have Google, Facebook or Tesla as part of your local industries, because this also sends a political message.

I sense that this is part of the challenge that has led regulators down the path where we have real regulatory gaps. I would also caution from expecting just information commissioners or privacy regulators to be able to respond to this. It's also media regulators, people responsible for elections, people responsible for ensuring that, on a day-to-day basis, competition functions.

All of these regulators are heavily challenged by new digital technologies, and we would be wise as a society to take a step back and make sure they're really able to do their job as regulators, that they have access to all of the relevant data. We may find that there are still regulatory gaps where we possibly need even additional regulatory authorities.

There, I think the danger is to say we just want progress; we just want innovation. If you do that a few times and keep allowing that to be a possibility... It doesn't mean that you have to say no to people like Facebook or Google if they want to invest in your country, but if you start getting threats like this, I would see them as exactly what they are: a futile attempt to resist the change that is already coming.

• (1620)

Hon. Peter Kent: Thank you.

The Chair: Next up, we have Mr. Angus for seven minutes.

Mr. Charlie Angus: Thank you, gentlemen.

You're raising I think some very disturbing, broad questions that are so much beyond the scope of our committee and what we do as politicians. My day job is to get Mrs. O'Grady's hydro turned back on—her electricity. That's what keeps me elected.

However, when we're talking AI with you, we're talking about the potential of mass dislocation of employment. What would that mean for society? We have not even had conversations around this. There's the human rights impact, particularly exporting AI to authoritarian regimes and what that would mean.

For me, trying to understand it, there are the rights of citizens and personal autonomy. The argument we were sold—and I was a digital idealist at one point—was that we'd have self-regulation on the Internet and that would give consumers choice; people would make their decisions and they'd click the apps that they like.

When we're dealing with AI, you have no ability as a citizen to challenge a decision that's been made, because it's been made by the algorithm. Whether or not we need to look at having regulation in place to protect the rights of citizens...

Mr. Wagner, you wrote an article, "Ethics as an Escape from Regulation: From ethics-washing to ethics-shopping?"

How do you see this issue?

Prof. Ben Wagner: I'm sure you've guessed, from the title you mention, that I do see the rise of ethics—and I explicitly wouldn't include the Montreal declaration in this, because I don't think it's a good example of it... There are certainly many cases of ethical frameworks that provide no clear institutional framework beyond them. A lot of my work has been focused, essentially, on getting people to either do human rights and governance, or, if they will do ethics, then to take ethics seriously and really ensure that the ethical frameworks developed as a result of that are rigorous and robust.

At the end of the article you mentioned there is literally a framework of criteria on how to go through this: external participation, external oversight, transparent decision-making and non-arbitrary lists of standards. Ethics don't substitute fundamental rights.

To come back to the example you mentioned on self-regulation on the Internet and how we all assumed that that would be the path that would safeguard citizens' autonomy, I think that's been one of the key challenges. This argument has been misused so much by private companies that then say, for example, "Well, we have a million likes, and you only have 500,000 votes. Surely our likes are worth as much as your votes." I don't even need to explain that in great detail. It's just this logic of lots of clicks and lots of likes surely can be seen as the same thing as votes. This, in a democratic context, is extremely difficult.

Lastly, you specifically mentioned exporting AI to authoritarian regimes. I think there is a strong link between the debates we have about exporting AI to authoritarian regimes and how we trade in and export surveillance technologies. There are a lot of technologies that are extremely powerful that are getting into the wrong hands right now. Limiting that or ensuring, through agreements like the Wassenaar arrangement and others, that there is dual-use control for certain types of technology will become increasingly important.

We have existing mechanisms. We have existing frameworks to do this, but unless we're willing to implement those and sometimes also say that we will do it collectively as a group, even if this means having slightly less—and I emphasize "slightly less"—economic growth as a result of this, we can still also say we're taking more leadership on this issue. It's going to be very difficult to see where these short-term economic gains are going to meaningfully provide for a human rights environment we would want to stand behind in the years and decades to come.

• (1625)

Mr. Charlie Angus: Thank you.

I'm a music buff. Every morning I wake up, and YouTube has selected music for me. Their algorithms are pretty good, and I watch them. I'm also a World War II buff, and YouTube offers me all kinds of documentaries. I see some of these documentaries on the great historian David Irving, who is a notorious Holocaust denier, and they come up in my feed.

Now, I have white hair; I know what David Irving is, but if I'm a high school student, I don't. It has a lot of likes because a lot of extremists are promoting it. The algorithm is pushing us towards seeing content that would otherwise be illegal.

In terms of self-regulation, I look at what we have in Canada. In Canada, we have broadcast standards for media. That doesn't mean we don't have all manner of debate and crazy commentary, and people are free to do it, but if someone was on radio or television promoting a Holocaust denier, there would be consequences. When it's YouTube, we don't even have a proper vehicle to hold them to account.

Again, in terms of the algorithms pushing us towards extremist content, do you believe that we should have some of the same kinds of legal obligations that are for regular broadcast media? You're broadcasting this. You have an obligation. You have to deal with this.

Prof. Ben Wagner: I think there is a distinction to be made between online platforms and media platforms. I think there is a substantive difference. I don't think it's always helpful to just focus on the content. In a lot of these cases, the solutions to this tend to be more procedural and tend to be more, let's say, organizational. If you have ways in which consumers have more control over the algorithms that YouTube is using to present them with music or to present them with information, that can already deal with a large part of the problem.

That's not to say that there isn't a responsibility with these large organizations; for sure there is. It's just also the grave danger that when too much government regulation decides what you can and cannot see on the Internet, that's not always the—

Mr. Charlie Angus: What about illegal content?

Prof. Ben Wagner: If it's illegal in that specific jurisdiction, then steps definitely need to be taken to ensure that.... But a lot of the time, at least in my experience of looking at content moderation, it's not so much about legal or illegal; it's more content that creates a certain atmosphere, and the challenges of that certain atmosphere chill speech and make minorities, different genders or people with different sexual orientations much less comfortable speaking, and that impoverishes the public sphere.

We live in a world, right now, where there is a real challenge that people who are important parts of our communities no longer feel comfortable debating things on the Internet. I don't think just saying that it's identical to media will fix that problem. There is a huge challenge on how to restore a space where people genuinely feel comfortable having a public conversation. I think that's a huge challenge but an extremely important one.

The Chair: Thank you.

Next up for seven minutes is Mr. Saini.

Mr. Raj Saini: Good afternoon to both of you gentlemen.

Mr. Bengio, I'd like to start with you, because I would like to ask a technical question just so I have a better understanding of how algorithms work. I'm sure you're aware of the term "black box problem".

Prof. Yoshua Bengio: Yes.

Mr. Raj Saini: Can you explain that? To me, that sounds like you have an algorithm, the data is not very good, the algorithm produces a result and you just take it for granted that this is the result, without having any human eyes on it or any human interaction. Can you explain that a little more for me?

Prof. Yoshua Bengio: Sure. Actually, we know a lot of things about how that result is obtained. We know that it's obtained as a consequence of optimizing some objectives—for example, minimizing the prediction error on the large dataset—and that tells us a lot about what the system is trying to achieve. When the system is designed, we can also measure how well it achieves that and how many errors it makes on new cases on average. There are many other things you can do to analyze those systems before they are even put in the hands of users.

It's not really a black box. The reason people call it a black box.... In fact, it's very easy to look into it. The problem is that those systems are very complex and they're not completely designed by humans. Humans designed how they learn, but what they learn and detail is something that they come up with by themselves. Those systems learn how to find solutions to problems. We can look at how they learn, but what they learn is something that takes much more effort to figure out. You can look at all of the numbers that are being computed. There is nothing hidden. It's not black; it's just very complex. It's not a black box. It's a complex box.

There are things that we can do very easily. For example, once the system is trained and we look at a particular case where it's taking a decision, it's very easy to find out which of the variables it takes as input that were most relevant and how they influenced the answer. There are things that can be said to highlight it, to give a little bit of explanation about their decisions.

• (1630)

Mr. Raj Saini: Thank you.

Mr. Wagner, I want to ask you a question about a term you used in a recent paper you wrote. You talked about "quasi-automation" and about keeping humans in the loop. Can you explain that to us a little more clearly?

You talked about three places where you felt that human agency, or the involvement of human agency decision-making, was debatable. You talked about self-driving cars. You talked about border searches on passenger name records. You also talked about social content media.

Perhaps you could expand on that term for us so that we have a better understanding of what you meant.

The Chair: Could you hear that question, Mr. Wagner?

I guess not. Are you able to hear me now, either one of you?

Prof. Yoshua Bengio: I'm hearing you fine.

The Chair: Mr. Wagner?

Mr. Nathaniel Erskine-Smith: I'd take that as a no.

The Chair: Yes. I'll take that as a no.

Your time is still ticking, too, Mr. Saini. Hopefully, you'll get it back.

• (1635)

Mr. Raj Saini: Mr. Bengio, I have just one quick question. I came across the term “singularity”. Is it a real thing?

Prof. Yoshua Bengio: No.

Mr. Raj Saini: Could you explain it a little? When I read it, I was alarmed, as you can appreciate.

Prof. Yoshua Bengio: Yes. That is the intention of people who—

Mr. Raj Saini: So, is it a real thing, and if it is—

Prof. Yoshua Bengio: No.

Mr. Raj Saini: It's not a real thing. Then why does it keep being written about?

Prof. Yoshua Bengio: Unfortunately, there is a lot of confusion in many people's understanding of AI. A lot of it comes from the association we make with science fiction.

The real AI on the ground is very different from what you see in movies. The singularity is about the theory. It's just a theory that once the AI becomes as smart as humans, then the intelligence of those machines will just take off and become infinitely smarter than we are.

There is no more reason to believe this theory than there is, say, to believe some opposite theory that once they reach human-level intelligence it would be difficult to go beyond that because of natural barriers that one can think of.

There is not much scientific support to really say whether something like this is an issue, but there are some people who worry about that and worry about what would happen if machines became so intelligent that they could take over humanity at their own will. Because of the way machines are designed today—they learn from us and they are programmed to do the things we ask them to do and that we value—as far as I'm concerned, this is very unlikely.

It's good that there are some researchers who are seriously thinking about how to protect against things like that, but it's a very marginal area of research. What I'm much more concerned with, as are many of my colleagues, is how machines could be used by humans and misused by humans in ways that could be dangerous for society and for the planet. That, to me, is a much bigger concern.

The current level of social wisdom may not grow as quickly as will the power of these technologies as they grow. That's the thing I'm more concerned about.

Mr. Raj Saini: Thank you very much.

The Chair: We have Mr. Wagner back.

Mr. Raj Saini: Mr. Wagner, before we got cut out I was quoting a term you had written in a paper recently called “quasi-automation”. You talked about the lack of human agency in certain decision-making processes that was debatable, for example, self-driving cars, border searches and also social content on the media. Dr. Bengio had also indicated fake news and the misuse of AI.

It seems to me that in some cases human beings are just part of the loop for the minimum amount of contact. How do we make sure there is still the human dimension in making decisions, especially when it comes to certain things like fake news and the use of political advertising?

Prof. Ben Wagner: There's a challenge in that if we assume human intervention alone will fix things, we will also be in a difficult situation because human beings, for all sorts of reasons, often do not make the best decisions. We have many hundreds of years of experience of how to deal with bad human decision-making and not so much experience in how to deal with mainly automated human decision-making, but the best types of decisions tend to come from a good configuration of interactions between humans and machines.

If you look at how decisions are made right now, human beings often rubber-stamp the automated decision made by AIs or algorithms and put a stamp on it and say, “Great, a human decided this”, when actually the reason for that is to evade different legal regulations and different human rights principles, which is why we use the term quasi-automation. It seems like it's an automated process, but then you have three to five seconds where somebody is looking over this.

In the paper I wrote and also in the guidelines of the Article 29 Working Party, guidelines were developed for what is called “meaningful human intervention” and only when this intervention is meaningful. When human beings have enough time to understand the decision they're making, enough training, enough supports in being able to do the only event, then it's considered meaningful decision-making.

It also means that if you're driving in a self-driving car, you need enough time as an operator to be able to stop, to change, to make decisions and a lot of the time we're building technical systems where this isn't possible. If you look at the recent two crashes of Boeing 737 Max aircraft, it's exactly this example where you had an interface between technological systems and human systems, where it became unclear how much control human beings had, and even if they did have control, so they could press the big red button and override this automated system, and whether that control was actually sufficient to allow them to control the aircraft.

As I understand the current debate about this, that's an open question. This is a question that is being faced now. With autopilots and other automated systems of aircraft, this will increasingly lead to questions that we have in everyday life, so not just about an aircraft but also about an insurance system, about how you post online comments, and also how government services are provided. It's extremely important that we get it right.

• (1640)

Mr. Raj Saini: Thank you very much.

The Chair: Mr. Gourde.

[*Translation*]

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

I want to thank the witnesses for joining us.

I'll turn to you, Mr. Bengio. Perhaps you'll answer me in French.

Artificial intelligence finds solutions to our problems and improves the services, knowledge and information that we receive. So far, so good. However, you really worried me when you said that artificial intelligence can find solutions to problems on its own. What would happen if artificial intelligence determined that we were the problem?

You mentioned killer drones, which may be capable of genocide. If artificial intelligence programming includes a list of all Canadian parliamentarians to eliminate within a week, should we be concerned? Is this pure fiction? Could this happen?

Prof. Yoshua Bengio: Some things that you said are pure fiction, but others are cause for concern.

I think that we should be concerned about a system that uses artificial and programmed intelligence to target, for example, all parliamentarians in a certain country. This situation is quite plausible from a scientific point of view, since it involves only technological issues related to the implementation of this type of system. That's why several countries are currently discussing a treaty that would ban these types of systems.

However, we must remember that these systems aren't really autonomous at a high level. The systems will simply follow the instructions that we give them. As a result, a system won't decide on its own to kill someone. The system will need to be programmed for this purpose.

In general, humans will always decide what constitutes good or bad behaviour on the part of the system, much like we do with children. The system will learn to imitate human behaviour. The system will find its own solutions, but according to criteria or an objective chosen by humans.

Mr. Jacques Gourde: We've talked a great deal about job losses, efficiency and the fact that artificial intelligence could eventually replace foremen in factories. I could be informed by means of my smartphone of the work that awaits me today on my production line. Artificial intelligence could arguably do much of the work itself.

Will workers end up going to the factory simply to carry out tasks that are too difficult for robots to perform, such as moving certain items? We're even talking about artificial intelligence controlling transportation. Could a large part of the population be unemployed within 10 to 20 years?

• (1645)

Prof. Yoshua Bengio: Yes, it's quite possible.

Your example of a machine that assigns the work already exists. For example, today, couriers who carry letters from one end of the city to the other are often guided by systems that use artificial intelligence and that decide who will carry a given package. There's no longer any human contact between the dispatcher and the person performing the tasks.

As technology advances, obviously more and more of these jobs, especially the more routine jobs, will be automated. In the courier example that I just provided, the dispatcher's job was the most

routine and easiest to automate. The work of a human who walks the streets of the city is more difficult to automate at the moment. However, it will probably happen eventually.

It's very important for governments to plan, anticipate the future and think about measures that will minimize the human misery that may result from this development if it were left to run its course.

Mr. Jacques Gourde: My last question concerns ethics.

Our government is increasingly using artificial intelligence to provide services to Canadians. All governments in the world are doing the same. If I need services from my government and I see that the responses provided have been generated by artificial intelligence—a reality that's fast approaching—how can I be sure that a human has listened to me? To what extent can I invoke ethical considerations to require that the service be provided by another person?

Prof. Yoshua Bengio: It depends on the type of service. In some cases, all that matters is that the job is done properly.

Personally, I would prefer to receive quick and efficient responses from tax officials, even if the responses must be generated by a machine. I'm using this example because we're in the middle of tax season.

However, if I have questions about my health, if the discussion takes a more personal turn or if I'm ill and in hospital, I want to have a human in front of me. It doesn't bother me that the human uses technology to do a better job. That said, some situations involve human and relational concerns that are better addressed through human interaction.

Mr. Jacques Gourde: Thank you.

[English]

The Chair: Mr. Picard, for five minutes.

[Translation]

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

[English]

Let's try to talk about the positive side of AI, as we have been kind of scared of all the issues.

Mr. Wagner, you talked about the moral leadership. In my view, moral is a bit wider than ethics. Ethics is the values that you decide to promote, and then governance would apply them; but with moral, there's a societal aspect to it where AI is used by individuals, but the system in general is created first by humans. The system must help us govern our values, to the point that you suggest—moral. Who will do that? Who is credible enough for me to say, well, that's a wise person? I should count on this type of person or persons to establish what will be from now on my moral guide in my AI systems.

Prof. Ben Wagner: I think right now we live in a situation where these decisions are overwhelmingly made by private companies, and almost none of these decisions are made by democratically elected governments, and this is the problem for citizens, for rights, for governance. It poses a considerable challenge, but that doesn't mean that it's impossible. Whether it's the trade in the technologies where you choose to export them; the development of the technology, and which ones you focus on developing; the research and the research funding for different technologies, what you focus on, and what you ensure is developed, I do think there is an opportunity for moral leadership—which I think is the right word there.

But also to be perfectly blunt, there aren't that many countries in the world that are seriously trying to develop artificial intelligence in a positive way for their citizens and for its development in the context of human rights. There are many that are discussing it and trying, but a lot of the time they're saying, "Ah, but we're not quite sure. Would it have issues for economic development? Ah, we're not quite sure if some of our companies will have some mild issues here or there."

I think there is a need to be willing and also have the strength to take that stand, but I also do think it's important because if there are no countries left in the world that are willing to do that, then we're in a very difficult spot. I think the European general data protection regulations have a perspective on what things could be done on data. But for artificial intelligence, for algorithms, we have a whole new set of issues, a whole new set of challenges, where I think there will be further leadership required to really get to a human rights basis and a basis that benefits all citizens.

• (1650)

[Translation]

Mr. Michel Picard: Mr. Bengio, I want to hear your comments on this issue.

At this time, the legislation that a country adopts to protect privacy to some extent may become counterproductive. The legislation may prevent that country from fully developing its use of artificial intelligence. The country will then have a leadership issue. However, both of you have stated that Canada wants to be a leader in different areas.

First, I wonder where this leadership should begin, since the concept is so broad. In addition, Canadian leadership, which is probably based on Canadian values, may not be equal to the leadership of another country that uses a different value system.

Prof. Yoshua Bengio: You're asking a good question, but I don't think that there's a general answer.

This requires the use of experts, who will review ethical and moral issues, along with technological and economic concerns in each relevant area. The goal is to establish guidelines to both foster innovation and protect the public. I think that this is generally possible. Of course, several companies have protested that there shouldn't be too many barriers. However, in most cases, I don't believe that the expected results pose an issue.

As we said earlier, there are issues in some situations, but there's no easy solution. We specifically talked about [Technical difficulty—Editor] illegal videos on Facebook. The issue is that we don't yet

have the technology to identify these videos quickly enough, even though Facebook is researching ways to improve this type of automatic identification. However, not enough humans are monitoring everything put on the Internet in order to remove things quickly and prevent things from being posted.

The task is practically impossible, and there are only three possible solutions. We can shut everything down, wait until we've developed better technology, or accept that things aren't perfect and that humans carry out the monitoring. In fact, this is already the case right now, when people have the opportunity to click on a button to report unacceptable content.

[English]

Mr. Michel Picard: Thank you.

The Chair: I just want to bring this before the committee. We actually have three more on the list to ask questions—Mr. Kent, Mr. Baylis and Mr. Angus—but we're at about the five-minute mark. Because we had so many delays, I would suggest we go into committee business slightly later, but it's up to you.

What would you like to see? I think the questions still need to be asked, but I'm looking for direction from the committee just to finish the slate of questions.

Some hon. members: Agreed.

The Chair: It will just push into committee business by about eight minutes.

Okay, we'll continue.

Next up for five minutes is Mr. Kent.

Hon. Peter Kent: Thank you, Chair.

In the interests of time, I have just one last question that I'd like to ask to both of our witnesses, and it comes back to this matter of regulation in the borderless digital world.

Do you see a need for international treaties that would govern the development and the use of artificial intelligence in ways similar to the EU? It has this new GDPR, which is certainly far beyond any regulations we have in Canada. Would either of you see the need for international, meaningful, enforceable—and I think the word "enforceable" is key to this—international treaties to enforce the way artificial intelligence might be used or abused?

• (1655)

Prof. Yoshua Bengio: I think we have to go in that direction. It's not going to be perfect because there will be countries that don't embark, or some countries might be able to water down the strength of these things.

Even some regulation there—and in particular, here we're talking about international regulations—is better than none, by far.

Hon. Peter Kent: Mr. Wagner.

Prof. Ben Wagner: In the experience of what I've seen, when people try to develop general regulations for all of AI or all of algorithms or all of technology, it never ends up being quite appropriate to the task.

I think I agree with Mr. Bengio in the sense that, if we're talking about certain types of international regulation, for example, it would be focused on automated killer systems, let's say, and there is already an extensive process going on in this work in Geneva and in other parts of the world, which I think is extremely important.

There is also the consideration that could be made whether Canada wants to become, itself, a state that has protections equivalent to the GDPR and that, I think, is also a relevant consideration that would considerably improve both flows of data and protection of privacy.

I think all other areas need to be looked at in a sectoral-specific way. If we're talking about elections, for example, often AI and other automated systems will abuse existing weaknesses in regulatory environments. So how can we ensure, for example, that campaign finance laws are improved in specific contexts, but also ensure that those contexts are improved in a way that they consider automation? When we're talking about the media sector and issues related to that, how can we ensure that our existing laws adapt and reflect AI?

I think if we build on what we have already, rather than developing a new cross-sectional rule for all of AI and for all algorithms, we may do a better job.

I think that also goes at the international level where it's very much a case of building on and developing from what we already have, whether it's related to dual-use controls, whether it's related to media or whether it's related to challenges related to elections. I think there are already existing instruments there. I think that's more effective than the one-size-fits-all AI treaty.

Hon. Peter Kent: Thank you.

Thank you, Professor Bengio, for your explanation and your discussion of similarity. I had the occasion to watch an old version of *2001: A Space Odyssey* and the battle between the human and HAL over the control of the spaceship. Your discounting of the reality of similarity was reassuring. Thank you.

The Chair: Thank you, Mr. Kent.

Next up for five minutes is Mr. Baylis.

Mr. Frank Baylis (Pierrefonds—Dollard, Lib.): Thank you.

Is there any country that regulates AI directly?

We'll start with you, Mr. Wagner.

Prof. Ben Wagner: Are there countries that regulate AI directly, as a general thing, that I am aware of right now? No. There are AI-specific regulations in different fields that you will find; for example, the general data protection regulation in Europe is one of those cases.

Mr. Frank Baylis: Mr. Bengio, do you know of any specifically, or are they all subsets of a general regulation?

Prof. Yoshua Bengio: No, I don't think there is. I would agree with Mr. Wagner that we want sector-specific regulations. That's also a protection for innovation—to make sure we find the right compromise that makes sense both ethically and technically.

Mr. Frank Baylis: No one is out there saying we're going to regulate AI in a general sense. They're doing more of what you're suggesting to us, Mr. Wagner, which is to say we already regulate, for example, hate speech. Take that one. How is AI going to be

regulated within the context of hate speech? Is that the approach you would both suggest?

Prof. Yoshua Bengio: Yes.

Mr. Frank Baylis: Are there specific areas in immediate need? As you're at the forefront of the development of AI, do you see specific areas? It's quite a broad thing to start looking at every one of our regulations and say, "Okay, we've got to make every one of them AI-proof." Where would you say we should focus our energies? Where have other jurisdictions focused their energies?

• (1700)

Prof. Yoshua Bengio: I think I already mentioned the security and military applications that deserve more attention. We have to act quickly on this, to avoid the kind of arms race between countries that would lead to the availability of these killer drones. It would mean that it would be very easy for even terrorists to get hold of these things. That's an example where there's no reason to wait. The red line has been defined—

Mr. Frank Baylis: Is that like the anti-personnel mines?

Prof. Yoshua Bengio: Yes. That's right.

Mr. Frank Baylis: Okay. That's the military's one.

Mr. Wagner.

Prof. Ben Wagner: I think the automated drones, or what are termed LAWS—lethal autonomous weapons systems—are definitely areas where further focus is acquired. I would also say that what's been mentioned here about the spread or proliferation of surveillance and AI technologies that can be misused by authoritarian governments is another area where there is an urgent need to look more closely.

Then, of course, you have whole sectors that have been mentioned by this committee already—media, hate-speech-related issues and issues related to elections. I think we have a considerable number of automated technical systems changing the way the battleground works, and how existing debates are taking place.

There's a real need to take a step back, as was mentioned and discussed before, in the context of AI potentially being able to solve or fix hate speech. I don't think we should expect that any automated system will be able to correctly identify content in a way that would prevent hate speech, or that would deal with these issues to create a climate. Instead, I think we need a broad set of tools. It's precisely not relying on just humans or technical solutions that are fully automated, but instead developing a wide tool kit of issues that design and create spaces of debate that we can be proud of, rather than getting stuck in a situation where we say, "Ah, we have this fancy AI system that will fix it for you."

Mr. Frank Baylis: Okay.

One of the areas we hear about, and you mentioned a few times, is transparency. I'm not talking about the transparency of how an algorithm works but the transparency of what you're dealing with. This is a bot, or a design, as opposed to a human being. These are AI-driven bots. What are your views on that?

We'll start again with you, Mr. Bengio.

Prof. Yoshua Bengio: It's usually pretty obvious if you're dealing with a machine or a human, because the machines aren't that good at imitating humans. In the future, we should definitely have regulations to clarify that, so that a user knows whether they are talking to a human or a machine.

Mr. Frank Baylis: Mr. Wagner.

Prof. Ben Wagner: I couldn't agree more. There are also cases, as best as I'm aware, in California, where it's also already being debated—to find mechanisms whereby automated systems like bots would be required to declare themselves as bots. Especially in the context of elections, and also in other cases, that can be quite helpful. Of course, that doesn't mean that all issues are fixed with that, but it's certainly better than what we have right now.

The Chair: Thank you, Mr. Baylis.

Last up for three minutes is Mr. Angus.

Mr. Charlie Angus: Thank you.

Mr. Bengio, my riding is bigger than Great Britain, and I live in my car. My car is very helpful. It tells me when I'm tired, and it tells me when I need to take a break, but it's based on roads that don't look like roads in northern Ontario. I'm always moving into the centre lane to get around potholes, to get around animals and to get away from 18-wheelers. I start watching this monitor, and sometimes I'm five minutes from the house and it's saying I've already exceeded my safety capacity.

I thought, well, it's just bothering me and bugging me. I'll break the glass. Then I read Shoshana Zuboff's book on surveillance capitalism and how all this will be added to my file at some point. This will be what I'm judged on.

To me, it raises the question of the right of the citizen. The right of the citizen has personal autonomy and the right to make decisions. If I, as a citizen, get stopped by the police because I made a mistake, he or she judges me on that and I can still take it to some level of challenge in court if I'm that insistent. That is fair. That's the right of the citizen. Under the systems that are being set up, I have no rights based on what an algorithm designed by someone in California thinks a good roadway is.

The question is, how do we reframe this discussion to talk about the rights of citizens to actually have accountability, so their personal autonomy can be protected and so decisions that are made are not arbitrary? When we are dealing with algorithms, we have yet to find a way to actually have the adjudication of our rights heard.

Is that the role you see legislators taking on? Is it a regulatory body? How would we insist that, in the age of smart cities and surveillance capitalism, the citizen still has the ability to challenge and to be protected?

• (1705)

Prof. Yoshua Bengio: It's interesting. This question is related to the issue of the imbalance of power between the user and large companies in the case of how data is used. You have to sign these consents. Otherwise you can't be part of, say, Facebook.

It's similar in the way the products are defined remotely. As users, we don't have access to the details of how this is done. We may disagree on the decisions that are made, and we don't have any recourse.

You are absolutely right. The balance of power between users and companies that are delivering those products is something that maybe needs rethinking.

As long as the market does its job of providing enough competition between comparable products, then at least there is a chance for things to be okay. Unfortunately, we're moving towards a world where these markets are dominated more and more by just one or a few players, which means that users don't have a choice.

I think we have to rethink things like notions of monopolies and maybe bring them back. We need to make sure one way or another that we re-equilibrate the power differential between ordinary people and those companies that are building those products.

Mr. Charlie Angus: Thank you.

The Chair: I want to thank our witnesses.

I think it's alarming just for you to say that essentially AI is largely unregulated. We're seeing that with data-polies as well, and we're really trying to grasp what we do as regulators to protect our citizens.

The challenge is before us, and it's certainly not easy, but I think we will take your advice. Mr. Wagner, you said to start early. It already feels like we're too late, but we're going to do our best.

I want to thank you for appearing today from Vienna, and from Montreal as well.

We're going to suspend for a few minutes to get our guests out so we can get into committee business.

Thank you.

[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>