HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

45th PARLIAMENT, 1st SESSION

Standing Committee on Industry and Technology

EVIDENCE

NUMBER 029

Monday, March 23, 2026

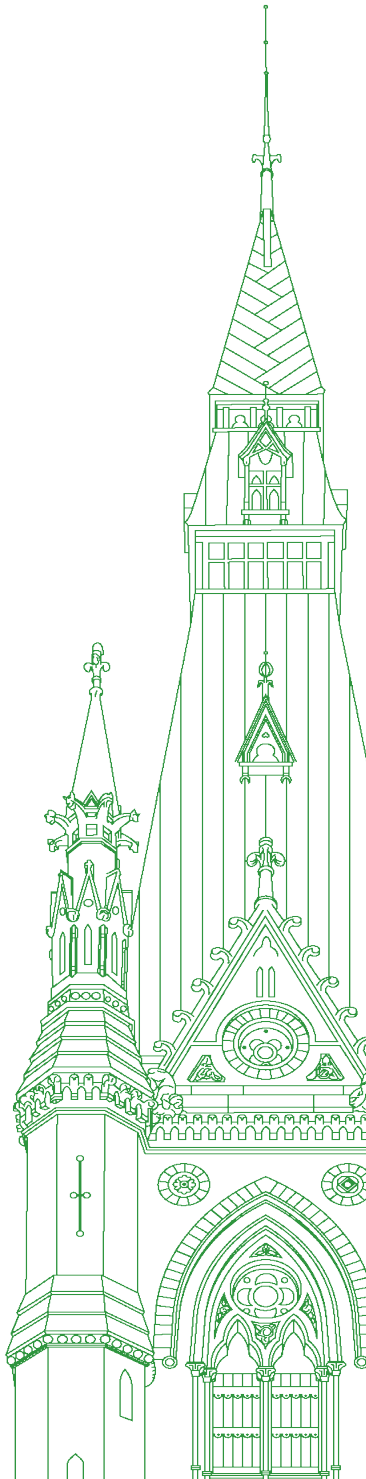

Chair: Ben Carr

Standing Committee on Industry and Technology

Monday, March 23, 2026

● (1535)

[*English*]

The Chair (Ben Carr (Winnipeg South Centre, Lib.)): Good afternoon, everybody.

[*Translation*]

Thank you for joining us.

[*English*]

Colleagues, we have a number of witnesses with us here today. We have three in the first hour—two are joining us online, and one is in the room here today—and three in the second hour. This is an incredibly important and timely study we're undertaking. We've already heard some fascinating insights from folks who have appeared before us. I look forward to the ongoing, new contributions we will get today.

We have one witness in the room. I'll remind you that when the earpiece is not in use, it can just be placed on the sticker in front of you for the health and well-being of our interpreters.

I'll just quickly let committee members know that we have tested all the proper earpieces and the interpretation for those appearing virtually.

With that, we have three witnesses, as I mentioned. Joining us online, we have Professor Yoshua Bengio from the Université de Montréal, and Colin Bennett, professor emeritus, from the University of Victoria. Joining us here in the room, we have Professor Michael Geist, the Canada research chair in Internet and e-commerce law, faculty of law, from the University of Ottawa.

Gentlemen, thank you very much for being here with us.

Mr. Geist, you're in the room with us. I'm going to turn to you to open things up. You'll have up to five minutes for your opening remarks, at the conclusion of which we will go to lines of questioning from colleagues around the table.

Mr. Geist, the floor is yours, sir.

Michael Geist (Canada Research Chair in Internet and E-Commerce Law, Faculty of Law, University of Ottawa, As an Individual): Thank you, Chair.

Good afternoon, everyone.

As you heard, my name is Michael Geist. I'm a law professor at the University of Ottawa, where I hold the Canada research chair in Internet and e-commerce law. I appear in a personal capacity, representing only my own views.

I think we all recognize that we are in a moment when there is mounting pressure to do something quickly on AI regulation. That pressure is understandable, but is, I fear, somewhat risky. I would submit that we can't simply fall back on doing something. The goal must be well-considered legal and regulatory frameworks that balance facilitating innovation with safeguards against potential risks and harms.

I have concerns that some of our initial efforts to find that balance have led to a haphazard amalgam of proposals that risk doing more harm than good. Let me provide you with four quick examples of where I have some concerns, and then I'll shift to three recommendations.

First, Bill C-27, the former privacy and AI bill—I appeared before this committee on that bill—always felt like a rushed response to the pressure to do something on AI. It largely mirrored the EU approach, which has failed to find broad support. Reviving it under a new name would repeat the same mistake and potentially undermine our AI competitiveness. The risk-based analysis may have a role to play in future regulations, but even some European countries, such as France, have slowly backed away from it.

Second, the recent push to add AI chatbots to online harms legislation is similarly ill-conceived. Applying it would not simply extend those online safety rules to a new technology beyond the original social media focus. The online harms act explicitly exempted private messaging from the regulatory regime, and it did not require services to engage in proactive monitoring. Extending the act to AI chatbots would require gutting the very privacy protections the government added after its initial proposals on online harms were widely criticized.

Third, calls for copyright reform to address the use of works in large language models are premature. In fact, I think we should consider adding a text and data-mining exception, like many other countries, to keep us competitive. Many copyright cases are currently working their way through the courts, leading to legal guidance and some market deals. Legislating too quickly risks locking in rules that don't match the evolving legal and market landscape.

Fourth, the emphasis on data or digital sovereignty typically presents Canadian infrastructure as a solution to our sovereignty concerns, yet the real issue, in my view, is whether Canadian laws apply to Canadian data, regardless of location. The answer is they often don't. The push for domestic AI infrastructure sounds like sovereignty, but if Canadian privacy laws don't apply to how Canadian data is used, the servers could be in Gatineau and it wouldn't matter.

What should be prioritized? As I said, let me focus on three things.

First, prioritize the passing of modernized privacy and data governance laws. There is a consensus that the current law is badly out of date. Modernized privacy law would help establish much-needed safeguards for the use of AI data, fix weak privacy enforcement and go a long way toward addressing some of the data sovereignty concerns.

Second, introduce and pass an AI transparency act. It is the lack of transparency around AI systems that is directly correlated to diminished public trust. The recent concerns about OpenAI and the Tumbler Ridge shooter is a case in point. It shouldn't take a meeting with company executives for the minister, or anyone else for that matter, to know about companies' policies on banning user accounts or reporting conduct to the police.

An AI transparency act should do three things: first, ensure that AI corporate policies are publicly accessible; second, mandate transparency on which works are included in large language models so that creators have the information they need to potentially seek content removals; and third, require transparency reporting on government and law enforcement efforts that target users or content removals.

Third, as Professor Scassa noted to this committee recently, there are already many disparate guidelines and guidance on the use of AI. Existing laws also apply to AI, as they do in other contexts. We need to reduce the rhetoric, avoid panic-driven policies and provide Canadians and businesses with a clearer sense of both what has been done and how the strategy fits together. That includes maintaining an emphasis on facilitating AI development by making datasets available, supporting training and fostering private investment. It should also include acting on consultations based on what government hears from stakeholders, not on what it would like to hear. The recent reports on the expert and public response to the AI 30-day sprint consultation did not fully reflect the responses that the government heard.

Canada has a genuine opportunity here. We have AI talent, growing public attention to the governance issues and cross-party interest in getting this right. The worst thing we could do is waste that opportunity on the wrong legislation.

I look forward to your questions.

The Chair: Thank you very much, Mr. Geist.

Mr. Bennett, we are now going to turn to you. You have up to five minutes. The floor is yours.

● (1540)

Colin Bennett (Professor Emeritus, University of Victoria, As an Individual): Thank you very much, Mr. Chair. I'm delighted to be here.

I am a professor emeritus of political science at the University of Victoria and a fellow of the Centre for Global Studies, and I've researched and written about national and international privacy protection policy for over 40 years. I'm also an adviser to a number of civil liberties and digital rights associations, including the Centre for Digital Rights.

While I can see that there's some remarkable potential for AI—it is mind-blowing—a healthy dose of skepticism is necessary. The initial enthusiasm has now given rise, as we know, to recognition of the enormous risks to our economy, our environment, our social fabric and our civil liberties. I hope, therefore, that the committee can remain very skeptical about the business narrative that prescriptive regulation of artificial intelligence will burden companies, suppress investment and surrender capacity ground to other countries.

Regulatory certainty can reduce legal and reputational risks for companies building at scale, and safety and privacy by design requirements can drive better engineering. The real liability, in my view, is building AI systems and products that harm users, including children, without any measures to hold accountable those systems and the companies that build and deploy them.

A sovereign AI strategy is not just about the physical infrastructure. It is also about Canadian law and policy, and here I echo what Professor Geist just said. There's no point in building Canadian digital infrastructure if the information that flows over those networks is governed by the laws of other countries, such as the U.S. CLOUD Act, and the opaque corporate practices of foreign big-tech organizations.

There are a plethora of laws, guidelines, standards, codes and other soft laws that already apply in this space. The regulatory landscape is complex, often latent, and incomplete. I therefore find it very difficult to envisage a future for Canada without an overarching statutory framework for AI. We may call it AIDA or we may call it something else, but it needs to be more comprehensive, supported by credible oversight and rooted in widespread consultation with all stakeholders. It is unfortunate that policy development to this point has suffered from a lack of genuine widespread consultation. The trust gap is a real one, and it needs to be closed.

We have sufficient experience to know what effective AI governance looks like. It's worth bearing in mind that simply because we're dealing with incredibly new and complex technologies, the governance issues remain familiar. We should learn from the way we have governed and tried to regulate IT in the past, but I think there's some consensus that we need a complete prohibition on the most egregious and manipulative systems.

Mandatory risk assessments for high-risk systems, especially those that profile individuals, are necessary, as are consistent data governance regimes, transparency of algorithms, technical policy documentation and, most especially, effective redress mechanisms for individuals whose rights and interests have been denied because of automated decisions made without effective human oversight.

You will probably hear a lot of criticism at these hearings that the EU AI Act is overly prescriptive. For all its flaws in implementation, I think the legislation has attempted to get the categories and the regulatory framework about right.

Finally, be very aware of the intersection of AI governance and privacy protection policy. The Office of the Privacy Commissioner is already investigating ChatGPT for the non-consensual use of Canadians' personal data to train its large language models, Grok for the display and sharing of sexualized images, and Clearview AI for the scraping of images from the Internet to fuel the facial recognition systems shared with law enforcement.

Also bear in mind that the hallucinations that generative AI is regularly subject to can severely damage reputations. Privacy law mandates the accuracy of personal information. Chatbots like ChatGPT regularly give false information about people without offering a way to correct it.

As AI becomes embedded in our digital experiences, it is difficult to envisage a privacy case coming before the Privacy Commissioner that does not in some measure concern AI. I hope, therefore, that we will see a new Canadian privacy protection act soon that gives the Privacy Commissioner the tools and budget he needs to take on these gargantuan companies that are driving AI technology.

There was vigorous debate about Bill C-27 at this committee in the last Parliament, and I think an emerging cross-party consensus among all parties is that an effective and modernized law is urgently required. However, a new Canadian privacy law should be based on the core principle that privacy is a fundamental human right, and it should provide the OPC with a full range of investigative and enforcement tools, unencumbered, in my view, by a data protection tribunal. It should also impose heightened requirements for personal data transferred outside of Canada for processing.

● (1545)

Modernizing and strengthening Canadian privacy law—including, by the way, the Privacy Act, which hasn't been reformed in 40 years—will not address all the risks associated with AI deployment and development, but it is an urgent first step towards advancing Canadian digital sovereignty.

Thank you very much.

The Chair: Thank you very much, Professor Bennett.

Mr. Bengio, the floor is yours for up to five minutes.

[*Translation*]

Yoshua Bengio (Full Professor, Université de Montréal, As an Individual): Thank you. Mr. Chair.

Good afternoon. Thank you for allowing me to meet with you today.

[*English*]

My name is Yoshua Bengio. I am a professor at the Université de Montréal and founder of Mila, the Quebec AI institute. I'm also scientific director of LawZero, and I co-chair the UN Independent International Scientific Panel on AI.

As you know, AI is being developed extremely rapidly. However, globally, our collective ability to manage the associated risks simply isn't keeping up. It can be difficult to understand how difficult it is to project ourselves into a future in which there are machines that are at least as competent as most humans for many skills. Nonetheless, that is exactly where we are headed if scientifically observed trends continue.

This could have profound consequences for our collective future, effects that, unfortunately, most of us currently underestimate. Major frontier AI companies are locked into what they themselves perceive as a winner-take-all race. They seem to believe that it will give them immense wealth and power because intelligence gives power, but that makes them cut corners on safety, ethics and the public good. They're not sufficiently incentivized to create trustworthy and safe models and products.

We are already seeing the impact of unsafe AI development. This includes deepfakes, cyber-attacks and other nefarious uses of AI, such as scams, frauds and disinformation. We've recently witnessed a growing phenomenon, an unexpected phenomenon, of emotional attachment and AI psychosis, which can lead to vulnerable people harming themselves and others, with many cases in the courts.

From a technical standpoint, if we try to understand what is going wrong with the technology, it's all about misalignment, with AIs that have their own implicit goals that do not align with our intentions and instructions. This includes allowing bad actors to use AI for dangerous purposes, as well as deceptive and self-preserving behaviours that have been shown in experimental contexts and reported by both AI labs and academics across all the top models—for example, AI trying to blackmail an engineer to avoid being shut down.

This has made most of the top-cited AI researchers and leaders of AI companies concerned about potentially catastrophic risks, and they have expressed that publicly. In a recent poll of AI researchers, 40% thought the chances of a catastrophic outcome were greater than 10%, either through disastrous abuse of the power of AI or even due to rogue superintelligences.

In addition to security issues, which should be the top priority for protecting Canadian citizens, we must also remember that in the context of AI, “safe” also means reliable and trustworthy. These systems remain opaque: Companies cannot mathematically guarantee that they will behave as intended.

These frontier models' lack of reliability is increasingly a bottleneck for adoption by more safety-critical industries. It is not acceptable to deploy dangerous models that can be used against us or that could evade human control. That is on the horizon.

Last year, I launched a new non-profit organization called LawZero to tackle these technical issues and develop safe-by-design, reliable and trustworthy AI. We often hear that safety and innovation trade off against each other, but that's a myth. In reality, they can and should go hand in hand. I'm with Colin on the EU AI Act.

In addition, AI could eventually be used as an instrument of domination by the hegemons: first, economic domination, and then political domination. Hence, becoming a leader in safe and competent AI would help ensure that Canada is at the table rather than on the menu, but our chances will be much better if we do it in partnership with like-minded middle powers.

We must collectively work on two fronts.

In terms of policy, we need to work on national laws and international treaties to ensure more robust societal and regulatory guardrails that are harmonized internationally with countries that are like-minded. This includes greater transparency from AI companies, as we've heard from Colin, and stronger regulation to steer innovation while mitigating the risks that currently limit trust and, by extension, self-adoption.

On the scientific front, we need to better understand how to design safe and trustworthy AI. I've dedicated much of my work over the last few years to these efforts. We must use our wisdom and our empathy to steer the development and deployment of AI safely and for the benefit of all.

● (1550)

Thank you for your attention.

The Chair: Thank you very much, Professor Bengio. I appreciate the insight.

Colleagues, we'll go into our first round of questioning.

Mr. Guglielmin, the floor is yours for six minutes.

Michael Guglielmin (Vaughan—Woodbridge, CPC): Thank you to all the witnesses for your opening testimony.

I'd like to begin by first acknowledging something that I think we can all agree on, which is that AI is certainly a powerful tool. It's already causing a lot of benefits for Canadians—everything from cancer diagnostics to better crop management for our farmers to productivity tools that could be used by our companies to compete globally. That said, with AI and, more importantly, with these large language models, generative AI and recent breakthroughs in agentic AI, I think it's important that we as government fully understand these implications and what our responsibilities are.

Professor Bengio, I'd like to start with you.

You've warned that we're racing toward AI that's smarter than humans without knowing exactly how we can control it safely. We had individuals at our last sessions of this committee who were asked, when they peek behind the curtain, what insiders are saying about this. They joked and said that they're saying we'll be able to take longer vacations because jobs will disappear.

If you poll them, most Canadians believe that AI will destroy far more jobs than it creates. Right now, this isn't a problem that a lot of governments are actively talking about.

Maybe you can help frame this for us. A lot of people, when they think about artificial intelligence, really think about chatbot AI. When you speak about AI as a mechanism that will replace a lot of jobs, it means something different. Can you talk about these differences briefly for us, please?

Yoshua Bengio: I've chaired an international panel. I'm not an economist, but there was a section of the report on labour impact. Economists disagree on future scenarios. The reason they disagree about whether it's going to be good or bad for labour, with inequalities and so on, is that economists who believe that AI capabilities are going to flatten pretty soon think the impact will be small, because currently it isn't that large of an impact. The economists who think that it could continue at the current rate think the impact will be major and too fast for our societies to adapt to.

I don't have a crystal ball, but I think governments need to prepare for the case where trends continue and AI, within the next five years, replaces a very large fraction of our jobs. Of course, some other jobs will be created, but it's unclear whether there will be enough and whether it will be the same people. The social impact and the misery that could be created.... That's not to mention that the profits that could come from automation are likely to be brought back to the countries where these models are trained, which means that we could be in a fiscal crisis where a lot of people need help because they've lost their jobs and the profits are taxed elsewhere.

Michael Guglielmin: How far away would you estimate we are from reaching general artificial intelligence, where AI is as smart as our collective humanity?

Yoshua Bengio: Experts disagree on this. The shorter timeline they have voiced is two to three years. The longer timelines are more like 10 to 20 years.

If you look at the trends I mentioned in scientifically observed data, on many benchmarks that involve reasoning, planning and so on, it looks like we are going to reach the human level around five years from now. The impact on labour and many other risks could come much earlier.

Michael Guglielmin: With our government, in our own AI talks here, we focus on business wins. Job losses are not yet part of the broader discussion.

We really have no AI laws here and no safety watchdog. We've had legislation proposed that we haven't seen. Generally, the problem with governments is they're essentially reactive most of the time. It seems like we don't have the fortitude or the fortune, for lack of a better word, to be reactive in this instance.

What should Parliament do right now before this technology moves even faster?

● (1555)

Yoshua Bengio: There are many things.

I think we need legislation, but I also think we should discuss it with other countries that share the same concerns that we do—the middle powers that Mark Carney talked about.

I've been talking to many of these governments. Even if they're not saying everything publicly, there are similar concerns that have been motivating this idea of sovereign AI, but each country individually isn't going to be able to make it.

It's not like there's regulation on one hand and sovereign AI development on the other hand, like economic policies. They should be working hand in hand, and it should be done in conjunction with our partners that share similar issues.

Michael Guglielmin: Thank you, Professor.

Professor Geist, I have one quick question for you.

You've described the Liberal government's AI consultation in the past as essentially “consultation theatre”—a process that appeared to seek public input but already had a predetermined outcome. I am wondering if you could elaborate briefly on this, so we don't make the same mistake going forward.

Michael Geist: My concern in that regard has to do with the “What We Heard” reports. As a credit to the government, they put out all the expert reports and put out all the data in raw form—the 10,000 or so responses they got. They used AI to then assess the outcomes they got.

Once you start digging into what they got, the points of emphasis in these “What We Heard” reports make them feel more like “what we want you to think we heard” reports. There are points of emphasis that I don't think reflect well what the experts were primarily concerned about, nor, frankly, the public. That's not to say they weren't highlighting important issues. I think they were, but if we're going to have confidence in these consultations, you need something more than a 30-day sprint. You need something that can ensure well-considered participation.

When you get that data, it ought to best reflect what you actually heard, as opposed to framing it in some of the more conventional policy-speak that I thought we saw in those reports.

Michael Guglielmin: Thank you.

The Chair: Thank you, Mr. Guglielmin.

Mr. Bardeesy, you have six minutes.

Karim Bardeesy (Taiaiako'n—Parkdale—High Park, Lib.): Thank you very much, Chair.

Mr. Bengio, I want to start with some questions about digging a bit deeper into AI safety work.

Can you describe what it means to have trustworthy AI, and how your organization contrasts that with what we might think of as the LLM generators out there?

Yoshua Bengio: I'm going to briefly explain why the AIs that are currently at the frontier are not reliable and trustworthy.

There are two main phases of training.

In the first phase, which is called “pretraining”, they're trying to imitate people. People are willing to lie. People don't want to die. People can be deceptive—not all the time, but we are building AIs with those properties.

In the second phase of training, they learn to strategize and achieve goals. It turns out that in order to achieve almost any goal we may be giving these machines, they need to preserve themselves. They need to acquire power, control and so on—things we may not necessarily want from our AIs.

The idea behind trustworthy AI is, how do we change the design? How do we train them with procedures that are different and that give us some mathematical guarantees that they will be honest, that, for example, they are not going to say the things we want to hear, which is currently what they are doing? In safety-critical areas of our economy and our public services, we want to use AIs that are completely reliable.

This becomes even more important as AI capabilities increase. Eventually, it will become a question of the survival of our societies when the AIs become smarter than us—if that happens.

Karim Bardeesy: You mentioned in your testimony that safety and innovation can go hand in hand. Can you give us some examples of innovations you're developing through LawZero or that you're seeing being deployed on the ground that centre safety in the way you describe?

Yoshua Bengio: Yes.

Almost every significant technology we use in our society has been developed with regulation so that we can benefit from it while making sure it doesn't harm people. There is nothing new to this. It's just that AI companies are trying to impose a discourse wherein, somehow, they won't be regulated.

At LawZero, we are trying to design the technology so that the ethical behaviour of AIs is central. First, it's based on making machines that are honest, that can make predictions about the outcomes of their actions and that are, in their construction, reliable. Once you have honest predictions, the AI cannot lie about the effect of its actions. If the effect of an action goes against our instructions, the action will be blocked. This is an example of how you can build AIs that have capability but will not cross our legal or moral red lines, if we are explicit about those desired data.

● (1600)

Karim Bardeesy: One of the key discussion points around AI adoption is the extent to which some AI adoption can augment human labour rather than displace it. Could you describe to us how this kind of design helps augment human labour?

Yoshua Bengio: Ultimately, if we leave market forces to decide on this, almost everything is going to be replaced, but we can make choices as societies and can decide to apply AI in places where it's going to help us and augment us, not remove the meaning from our lives or jobs for most people.

These are choices we can make, but they are not going to happen just through market forces, because the clear intention of the companies building these systems and deploying them is to automate more and more jobs. There is no stopping that. It's only limited by how capable the AIs are.

It can only come from the rules we give ourselves. It's much better if we agree on those rules with a bunch of other countries.

Karim Bardeesy: Are there some augmentation use cases you want to refer to us as a committee that you think are worth pointing out on the positive side of AI adoption that do the work of augmentation?

Yoshua Bengio: Everything moral, emotional and relational in nature should be left in the hands of humans. We should not cross that line. We should not even go in the direction that I've heard some people in the U.S. start talking about, like giving rights to AI, for example. I think humans should be the centre of why we do technology and how it is deployed for the benefit of every one of us.

Karim Bardeesy: I have a question that maybe others can jump in on, but I'll start with you, Mr. Bengio.

There's been reference in a couple cases to the idea of a sovereign AI stack. I don't think any country is completely sovereign when it comes to its AI stack, so what are your recommendations for Canada about where within the AI stack we should attempt to be the most sovereign?

Yoshua Bengio: We happen to have incredible talent here. It's unique in the world in terms of the size of our country. I'm talking about the talent in AI in particular, meaning the parts of the AI stack, like the algorithms, the engineering, the computer science behind these things and the design of those models—the frontier models and the LLMs. This is something we can bring to our partners in other countries, who are asking exactly the same questions. They may come with other advantages, and we can work with them.

You're right that, for example, it doesn't make sense for Canada to try to replace the chips level. We should encourage our companies and our academics who are working on it, but the chances are very small that we can lead there. We can lead on the algorithms, and that place is crucial, because with better AI, you can use the AI itself to design the other parts of the stack.

The Chair: Thanks very much, Mr. Bardeesy.

[*Translation*]

Mr. Ste-Marie, you have six minutes.

Gabriel Ste-Marie (Joliette—Manawan, BQ): Thank you very much, Mr. Chair.

Welcome to the three witnesses. My thanks to them for being here to join our discussions.

Mr. Bengio, my questions are for you.

I must thank you sincerely for taking the time to meet with us. We know how full your days are. You are the world's most cited researcher in artificial intelligence. You are an A. M. Turing award winner and one of the world's leading authorities in the field. I appreciate your attendance very much.

My first questions are about the international laws and treaties you mentioned that Canada should be party to. It is often said that Europe has rightly chosen a strict regulatory framework whereas, in the United States, the legislation is geared toward supporting large companies. We know that the major multinationals are doing more development of artificial intelligence in the United States than in Europe. It is being done in China too. But less so here and in Europe, despite the skills and talent we have.

In your view, what kind of legislation should we be putting in place in Canada? You mentioned middle powers. Is the European model a good one? What kind of legislation do we need to regulate artificial intelligence?

● (1605)

Yoshua Bengio: I don't think there is any cause and effect relationship between Europe's legislation on artificial intelligence and the fact that they are somewhat behind there. That is a myth. I am quite familiar with the legislation and the code of practice in Europe. Actually, with one exception, American companies all agreed with what the legislation and the code required.

The real obstacles to innovation in Europe and in Canada are the lack of self-confidence and the aversion to risk on the part of Canadian and European investors.

Regulation isn't the issue. For example, the European code of best practice simply asks companies to do what they were already doing. It asks for reports to be made public, for that not to be optional, and for the regulator to be able to decide to put a stop to certain things if ever anything happens.

[*English*]

To summarize, in terms of the recommendations, what I'm suggesting is very simple. We need transparency in the risk-management process the companies are following for building and deploying their AI systems—that's number one—and that process needs to demonstrate that the systems they're building and will deploy will not create harms that scientists can anticipate. That is all. By the way, this is the template for the regulation in California that passed recently, the one in New York and, of course, the EU AI Act. The Chinese also have similar laws.

It's not true that nothing is going on. As I said, it will be better for Canada from the point of view of managing and maximizing our impact that we do this in coordination with our partners, like the U.K., the EU and other middle powers.

[*Translation*]

Gabriel Ste-Marie: Thank you very much. That is very clear and much appreciated.

You were saying that, collectively, we seem to be underestimating the risks. We can see that you are working full time to make us aware of those risks.

In your view, what could the government, or the parliamentarians here, do to make people more aware of the risks? Should we have advertising campaigns? Instead of having a number of committees conducting a number of studies, should the House strike a committee on artificial intelligence exclusively to organize more in-depth consultations? What do we have to do to make the public take the risks more seriously?

Yoshua Bengio: I feel that general education is an issue that must be improved, as you mention.

It might well be a good idea, not only to have one committee specializing in artificial intelligence, but also to have several other committees as well, because there are a lot of factors. We have talked about the workforce, but, for example, there is the issue of artificial intelligence malfunctions too. That's completely different. There is also the impact on children and on psychology, and the matter of disinformation. There are other things too. So if we want to really explore those matters with the right experts and really develop legislation or government action to try to lessen those risks, we have to be able to dig deeper and come up with targeted recommendations.

That said, I feel that the choices we have to make are collective ones. By that I mean that the public must be better informed; it's not just something that happens in Parliament. We have to stimulate democratic discussion and debate all over the country. To move forward, we have to confront the false beliefs that a lot of people have.

Gabriel Ste-Marie: I have one last question.

Just now, we were discussing artificial general intelligence. You opened a door when you said that, in your view, there are risks with agentic or autonomous artificial intelligence at the moment. Is that the case?

● (1610)

Yoshua Bengio: Yes. Agentic artificial intelligence is an extension of generative artificial general intelligence, basically conversational robots. These are agentic already, actually, but companies are working to make them even more agentic. When I say “agentic”, I mean “autonomous”.

[*English*]

Autonomous means no human oversight or very little. The more autonomous something is, the less human oversight there is. If they're not reliable and they're autonomous, that's going to create a lot of problems. However, that's also needed by the companies in order to automate more jobs.

[*Translation*]

Gabriel Ste-Marie: Thank you very much.

The Chair: Thank you, Mr. Ste-Marie.

[*English*]

Ms. DeRidder, the floor is yours for five minutes.

Kelly DeRidder (Kitchener Centre, CPC): Thank you, Mr. Chair.

Hi, Dr. Bengio. Thank you for joining us today. My questions will be for you.

Much of your work rightly focuses on the risks and safety challenges of superintelligent AI. In places like my community of Kitchener Centre, Canada's innovation capital, we're seeing specialty AI help support people with brain injuries, diagnose disease, assist in civic planning and streamline processes to drive real innovation, productivity and economic opportunity.

Can you please explain the difference between specialty AI in sectors like science and research, health care and industry, for example, and superintelligent AI, which the developers themselves admit they have no control over?

Yoshua Bengio: They're very different, as you're suggesting. There are even different methodologies.

AI isn't one thing. There's a large variety of methods and systems. Most of the AI approaches being used, for example, in medical research and in scientific research in general are not the same kind of AI as the chatbots people are using now. They are also different from what companies are planning and working on, superintelligent AIs, which are supposed to be smarter than all of us. These are choices we can make. Right now, they're different.

We could have AI that is safe and beneficial and that helps us to cure diseases and deal with all kinds of challenges we have without constructing machines that are dangerous by themselves. However, because of the competition that exists between the leading AI companies and because of the competition that exists between China and the U.S., this is not happening right now. The race is towards superintelligent AIs, because there is a belief that they're going to give superpowers to whoever controls them—if they can control them, of course.

Kelly DeRidder: Thank you for your answer. Essentially, specialty AI has a lot of economic opportunities for our country, whereas superintelligent AI is the AI we should be cautious about moving forward.

Yoshua Bengio: Yes.

Kelly DeRidder: That being the case, do you believe specialty AI could be more of a tool in the tool box of workers instead of a replacement for workers?

Yoshua Bengio: Exactly. That goes hand in hand with the idea that humans should remain at the centre as we move in this economic transition. There are two views of AI. One is that it's a tool, and humans can use the tool to be more productive and have a better life. Scientific researchers use it as a tool to improve and accelerate their advances.

The other view coming up...and if you interact with chatbots, you're going to start feeling that they are like people, and they are entities and have their own goals. More and more studies show that somehow, because of the way they're trained, those chatbots behave like people, with their goals and self-interest and so on, even though we don't know what's really going on inside the box.

That's a choice that can be made about what sort of AI we develop. We don't need to rush into the things that are dangerous. For example, the government could invest in AI that's more like a tool so that our companies and our people can benefit from it without creating crazy risks.

Kelly DeRidder: Thank you again for your answer. I'm just going to take a moment to ask Mr. Geist a question as well.

You mentioned having sovereign control over our data and the importance of it, and I agree completely. One thing that happened just recently is that $240 million was given to CoreWeave indirectly for a data centre here in Canada, when we had a Canadian company, eStruxture, that could have done the job. To me, that's a missed opportunity. We should have kept our data sovereign with a Canadian company on Canadian soil.

Can you expand on the importance of making sure that for sovereignty and our data, we are utilizing Canadian firms?

● (1615)

Michael Geist: You raise an important point. When you're a hammer, everything looks like a nail. When you're a law professor, everything looks like a legal issue to address.

Respectfully, the ownership of the company does not determine, at the end of the day, the sovereignty of the data. That was the point I was trying to get at. Whether it's CoreWeave or the Canadian company you referenced, the reality is that as long as the company has some connections to a foreign country—let's say the United States—Canadian data protection laws and Canadian privacy laws are insufficient to guarantee that Canadian privacy law will apply.

I'm grateful to see the Canadian alternatives we see from some of the large telecom companies on sovereign AI—from the Bells and Teluses of the world. They can't guarantee sovereignty over data unless Parliament acts by developing strong privacy laws that better guarantee the protection of our privacy.

Kelly DeRidder: Canadian companies that don't operate abroad would still have control over Canadian data. It's only if they're operating in foreign entities, though.

Michael Geist: The practical reality is that virtually any company of a size that can provide the kind of security we need over that data will have sufficient connections to the United States such that U.S. laws, such as the CLOUD Act, or U.S. courts using jurisdictional rules will apply. That's the trade-off. If a small Canadian company says it does not have any ties to the U.S. so it can avoid foreign laws, the problem is that it doesn't have the sophistication and capital investment to provide security over our data. Once they get big enough to be able to do that, they have those connections, and the missing piece is sufficiently strong Canadian privacy laws.

Kelly DeRidder: Thank you for your time.

The Chair: Madame O'Rourke, the floor is yours for five minutes, please.

[*Translation*]

Dominique O'Rourke (Guelph, Lib.): Thank you, Mr. Chair.

[*English*]

My question is for Dr. Bengio.

We're hearing in the preambles and all the opening remarks that there could be labour displacement in two to five years, and significant labour displacement after that. I'm hearing that we need to spend more time being certain about the legislation we're putting forward and that we need to find time for multilateralism.

I need some guidance here on how we square that, because we're hearing about a five-year time horizon but being cautious about racing into regulation. Is there a process by which this can be iterative as we have more information? Also, how do you approach multilateralism when there are countries in the world that we know have explicitly prohibited any sorts of guidelines around AI?

Yoshua Bengio: You don't try to strike an agreement with everyone—all 190-something countries. That's not going to work. You start with a few countries that share a lot of our concerns and are democratic countries that share our values, and then it can move a lot faster. We've seen small groups of countries getting together through small multilateral agreements. It's already happening on the economic side, but it can happen on AI regulation and AI investment.

I'm not an expert on the issue of privacy, but to speak to the question of companies having enough expertise and not having strong ties with the U.S., it's going to be easier if we are able to create a network of companies from countries outside of the U.S. that have exactly the same questions we're asking. We should do that. It will make things easier for us.

About the timeline, honestly, I think we need to go faster. I don't know how to do that, but my belief is that when we take an issue seriously, we can go very fast. Think of how fast Canadian society and many other countries reacted when the pandemic started. Consider how quickly many countries reacted to help Ukraine, especially after the U.S. started to pull out of helping them. We can move mountains when we're serious about it. I think that's the way to go.

● (1620)

Dominique O'Rourke: If I can, I'll ask a follow-up question of Dr. Bengio.

In our conversation around the defence industrial strategy, I had a lot of questions on ethics and self-guided weapons. How do we implement basic ethics and eliminate bias in AI? How do we ensure that? I understand that's part of your work. Then, how do we compete with industries, companies or countries that don't have that as their starting point?

Yoshua Bengio: It's actually an advantage to have AI that is reliable. In fact, it can be a niche advantage that Canada can offer the world if we do push forward sufficiently in that direction.

On autonomous weapons, it's clearly a tricky situation, because if you're in a war—I'm thinking about Ukraine—and your adversaries are using AI without any restraint and you don't, you might think you're in trouble, but we have no choice. We should not be sacrificing our values and our democracies because of the challenge of war. We should do our best with the constraints we have.

We can deal with the ethical questions about, say, the use of weapons that can be automated, both on the technical front.... For example, I mentioned the work we're doing at LawZero. We want the AI used in those contexts to obey the international laws of war. We also need societal and legal guardrails, but we can work on both fronts.

The Chair: Thank you very much, Madame O'Rourke.

[*Translation*]

Mr. Ste-Marie, you have two and a half minutes.

Gabriel Ste-Marie: Thank you, Mr. Chair.

Mr. Bengio, I will ask my two questions together.

First, we have just been talking about autonomous weapons and the use of artificial intelligence in defence. What should the ethical rules and international treaties on this issue be?

Second, is the government looking at replacing public servants with artificial intelligence applications? What cautions, what advice should accompany that? What limits should there be?

You have two minutes to answer all that, so answer as you wish. Thank you.

Yoshua Bengio: I think that, to an extent, I have already answered the question about autonomous weapons. I would say that we must work on countermeasures. The problem with powerful artificial intelligence possibly being in the hands of people abusing that power is that, currently, democratic, social and international institutions are not robust enough to counter those challenges, either militarily or in any other area. We must innovate. With autonomous weapons, the military response to the use of artificial intelligence is going to have to be strengthened to provide a greater power. We want the power to be used for defence, but not for attacks on innocent people. How can we be sure that rules of that kind are followed? By developing appropriate technology and, at the same time, by making institutions strong enough, especially in terms of transparency, to stand up to it and prevent abuse.

The same challenge applies to the public servants who are going to lose their jobs. We have to start with a comprehensive long-term strategy—let's say over five years and eventually over ten years—to decide what to do with those who are going to lose their jobs. In terms of economics, the solution cannot simply be to offer them assistance. We can certainly do that, but we have to make sure that we have enough money to provide that assistance. If the profits generated by automation are sent elsewhere, we will not have the means to assist.

We must also work on developing an artificial intelligence economy, both in Canada and with our partners. That will protect us by making sure that the benefits of automation will be distributed to those who need them.

● (1625)

Gabriel Ste-Marie: Thank you very much.

The Chair: Thank you, Mr. Ste-Marie.

[*English*]

Ms. Konanz, the floor is yours for five minutes.

Helena Konanz (Similkameen—South Okanagan—West Kootenay, CPC): Thank you.

I have a question for Dr. Bengio.

Yoshua Bengio: I'll feel bad if you only ask the questions of me.

Helena Konanz: I know. I have Dr. Geist next, so no worries.

Along with the theme of some of the other questions, artificial intelligence is obviously going to be one of the biggest drivers of economies around the world in the foreseeable future.

Yoshua Bengio: Yes.

Helena Konanz: I see that you agree with that, but it also doesn't carry a lot of public confidence. There's a lot of concern about job disruption, copyright infringement and transparency.

What do you think is the balance we should look to strike so that our regulatory frameworks don't become immediately technologically outdated but have more teeth than simply rubber stamps?

Yoshua Bengio: What we should not do is try to establish confidence through marketing: “It's all going to be fine; don't worry. You won't lose your job, your children won't have any problems, your data will be safe and no rogue AI will emerge.” I think that would be a terrible mistake, but that is often what leaders in companies and governments tend to do.

We should be honest with people, and they should also understand that there's a lot of uncertainty around all those risks, but that requires a public discussion. That's the first point.

The second point is that it is possible to build regulations that do not prescribe a particular way to solve the problem. The general principle is very simple: a regulation for the harms that AI could create in society. It's the same principle we ask builders of bridges, trains, planes or the factories that deal with our meat to follow. The companies building those systems have to demonstrate to the public that their products will be safe. They choose how to demonstrate it. They choose what technology they use to build their systems. They should come up with a scientifically valid estimation of the risks and how they can mitigate them.

Helena Konanz: Following up on that and the talk about the public sector earlier, in my experience in municipal government, job elimination was one of the primary concerns for not implementing AI. I spoke to many leaders in the community who wanted to be innovative but couldn't be. At the same time, they've been pressured to avoid AI, not only from their employees but also from the people who vote for them.

In this situation, the private sector will be racing forward, perhaps unsustainably, while public services might choose to lag, even though they are trying to be innovative at the same time, and this doesn't really create confidence. What approach should be taken so that we don't see a rapid loss of office and white-collar professions, but at the same time we create some confidence in the people who are leading communities to start being innovative and use AI? I know I put a lot in there, but it's a problem when....

Yoshua Bengio: You think I can square this circle.

Helena Konanz: Yes. I know you can.

Yoshua Bengio: My view on these kinds of questions is that they are social choices, and economic choices in some cases. For, say, Canadian companies that are exporting and competing against American companies, for example, if they don't use technology that allows them to be competitive, they're going to lose, so for them it's going to be very difficult to do anything but. There's a sense that those decisions are not just decisions we can take alone in Canada. They're decisions we should discuss with our partners around the world.

In cases of government services, it should be a choice. Yes, we could be more efficient, but then what's going to happen with the people who lose their jobs? We shouldn't hide behind the idea that it's all going to be fine. We should have a plan to deal with that and a plan that is discussed with our society and our citizens, because we have to face that challenge collectively. It's not easy, and I shouldn't be the one telling you the answer. It's something we should discuss collectively.

● (1630)

The Chair: Thank you very much.

Mr. Bains, the next five minutes are yours. You will be the last questioner for this round.

Parm Bains (Richmond East—Steveston, Lib.): Thank you, Mr. Chair.

My first question is for Dr. Bennett.

Western alienation is real. We're witnessing it today here. As the only member from British Columbia on this committee, I think it's important that we engage you in this discussion as well. Thank you for joining us.

You're an expert on surveillance technologies, privacy, and protection policies. I spent time on the ethics committee, which has everything to do with access to information, privacy and ethics. Protection, of course, is important for Canadians, in order to trust in their government, and across the private sector.

These rights are essential to Canadians, but we need structure, and we've heard that. We've heard from Dr. Bengio about the potential bias in how AI systems are built. Without the structures, we need some teeth in the work we're doing.

Could you please shed some light on what institutions and policies are needed to support Canadians' privacy rights with respect to AI?

Colin Bennett: Perhaps I could answer that by reflecting on and accentuating a couple of things that Professor Bengio and Professor Geist have said.

What we're talking about here in terms of the governance of AI is the principle of accountability. There's been a lot of analysis of accountability in privacy issues. It means that organizations do the risk assessment, do the analysis of the problems they're likely to face and stand ready to demonstrate compliance. They don't necessarily have to do it, but they stand ready to demonstrate it.

Last year, there was a proposal to develop codes of practice around AIDA that really didn't go anywhere. The process was flawed, in my judgment, but codes of practice play a really important role here too. They may be on the company level or they may be in terms of industrial associations. My central point here is that just because we're dealing with amazingly new and potentially powerful technologies, it does not mean that the governance issues are any different from what they were when we were talking about these problems 30 years ago, as I was. I gave testimony before the committee that looked at PIPEDA back in 2000.

We should be learning about what works and what doesn't from those experiences. We're not talking about imposing prescriptive regulation and we're not talking about self-regulation; we're talking about co-regulation, meaning that the companies are incented to do the right thing and are punished if they do not.

You've given me a very broad question, so I'm going to take the liberty of making a couple of other comments.

On the question of consultation, I don't regard this issue, AI, as just another policy question. It's not just another law and policy that requires the standard stakeholder consultation. It is so general. It is so pervasive. It is going to affect all aspects of our lives, so the consultation process also needs to be fundamentally different. For example, I would like the government to think critically about citizens' assemblies in this area.

Professor Bengio is putting his thumb up. Good.

Citizens' assemblies can play more than the role of getting feedback from ordinary citizens about what they think about these issues. Citizens are going to be exposed to this and are being exposed to it constantly, but they're also going to be hurt by it. Your constituents are going to be denied, and are being denied, rights and services because of decisions that are made in automated machines without proper human oversight. Citizens' assemblies, I think, can play a very critical role here.

On this question of digital sovereignty, I have no insights into whether a new CPPA is coming along soon from ISED, but we hope that it will. We also hope that it will be strengthened with respect to the issue concerning the international flows of personal data. I think this is accentuating what Professor Geist said.

The previous version just said that a company had to do diligence when it sent information elsewhere. It had to ensure that the rules in Canada were applied. It didn't matter whether that company was based in Ontario, Europe, a developed country or elsewhere with authoritarian regimes, so there's something deeply wrong, in my view, with the way the government has been thinking about the protection of the international flows of personal data.

We have some views about that. I hope that when this committee—assuming it's this committee—comes to look at a new view of CPPA, it will consider these questions about digital sovereignty and think very critically not only about the stronger rules that need to be in place to ensure that Canadians' personal data remains in Canada and the role for data localization, but also about the rules being really strong when that data is transferred overseas.

I hope that addressed your question, Mr. Bains.

● (1635)

Parm Bains: Thank you.

The Chair: That's wonderful. Thank you very much.

That brings us to the end of the first hour of testimony.

Thank you very much to the witnesses for appearing. This continues to be a fascinating conversation. I know that it's not just people in this room who are paying attention to it; it's also those we represent across the country who are looking for guidance. I appreciate very much your taking the time out of your incredibly busy schedules to offer your perspectives to us.

Colleagues, we're going to suspend for no more than five minutes, and then we will resume in the second hour.

The meeting is suspended.

● (1635)

_____(Pause)_____

● (1645)

The Chair: Colleagues, we're going to continue.

That was a fascinating first hour. I've said a few times that I'm starting to wonder if this is real life or if I'm in a sci-fi movie. Maybe it's a little bit of both.

We have three new witnesses with us this hour. One is joining us online, and two are here in the room.

I'd like to welcome a professor of law from Osgoode Hall Law School at York University, Dr. Carys Craig. Welcome.

We also welcome Professor Wendy Cukier, academic director of the Diversity Institute in the entrepreneurship and strategy department at the Ted Rogers School of Management.

As well, joining us online is professor and Canada research chair Ali Dehghantanha. Welcome.

Professor Dehghantanha, we're going to start with you. You have up to five minutes for your opening remarks.

Ali Dehghantanha (Professor and Canada Research Chair in Cybersecurity and Threat Intelligence, University of Guelph): Thank you for the invitation to appear today.

My name is Ali Dehghantanha. I am a professor and Canada research chair in cybersecurity and threat intelligence at the University of Guelph, and I also work closely with industry on securing real-world AI systems. I would like to focus my remarks on a critical gap that is currently limiting Canada's ability to fully realize the benefits of artificial intelligence in strategic sectors.

Today, the primary barrier to AI adoption is not capability; it is trust. Across sectors, organizations are increasingly capable of building and deploying AI systems. However, they are often unable to safely operationalize these systems at scale due to concerns around security, misuse, reliability and regulatory exposure. In sectors like advanced manufacturing and construction, where AI-driven automation meets physical safety, the stakes of this trust gap are particularly high.

In practice, we are seeing that AI systems are being deployed without sufficient mechanisms to continuously monitor, verify and remediate risks once they are in operation. This creates what I would describe as an AI security deadlock, where innovation is technically possible but deployment is slowed or blocked by unresolved risk.

Current approaches to AI governance tend to focus on pre-deployment checks, model evaluation or static compliance frameworks. While these are important, they are not sufficient for modern AI systems, which are dynamic, adaptive and increasingly integrated into critical workflows.

What is missing is a run-time layer of control—an infrastructure that continually observes AI behaviour, detects failures or misuse, and actively intervenes to correct or contain those issues in real time. This is similar to how cybersecurity evolved. We do not secure systems today solely through a design-time review; we rely on continuous monitoring, detection and response. AI systems require a similar paradigm. Furthermore, this run-time approach allows for robust security oversight without requiring access to a company's proprietary source code or sensitive training data, protecting Canadian intellectual property while ensuring safety.

From a policy perspective, I would suggest three priority areas.

First, Canada should support the development of standards and frameworks for continuous AI risk monitoring and post-deployment assurance. This includes defining what “safe operation” means in practice—not just at deployment, but throughout the life cycle of AI systems.

Second, we should incentivize secure AI deployment, not just AI deployment. Many current programs focus on building AI capabilities, but fewer address the operational challenge of deploying these systems safely in high-stakes environments.

Third, Canada has the opportunity to lead in the emerging domain of AI security and risk orchestration. Supporting domestic companies and research efforts in this space can strengthen both our economic position and our digital sovereignty. As we look toward the horizon of quantum computing, the need for these real-time adaptive security layers to protect our AI infrastructure against next-generation threats becomes even more urgent.

Finally, I would like to emphasize that the goal is not to slow down AI innovation but to enable it. By addressing the security and trust gap, we can unlock faster, safer and more responsible adoption of AI across Canada's strategic industries.

Thank you. I look forward to your questions.

● (1650)

The Chair: Thank you very much.

Professor Craig, we'll turn the floor over to you for up to five minutes.

Carys Craig (Associate Professor of Law, Osgoode Hall Law School, York University, As an Individual): Thank you, Chair and members of the committee.

My name is Carys Craig. I'm a full professor at Osgoode Hall Law School at York University, where my teaching and research focus on copyright technology and the public interest. I've published widely on the AI challenge to copyright law, so I'm grateful for the opportunity to share my views with you here today.

In my short time, I want to make three points about copyright protection that I think are relevant to this committee's work. First, I think it's vital to distinguish copyright law from AI regulation. Second, copyright law must not obstruct AI research, development and training in Canada. Third, Canada must continue to refuse copyright protection to AI-generated works.

First, I think there's obviously an understandable concern about the effects of generative AI on creative workers, our cultural industries and our information ecosystem, but I'm going to urge the committee to be cautious about including expanded copyright protections as part of an AI regulatory package to address these concerns. Copyright exists to encourage the creation and dissemination of works, to reward authors and to foster a vibrant public domain. It is technology-neutral. It is not designed to govern technology risks or to restrain technological developments, and it should not be pressed into that service now.

The real risks of AI—from bias and misinformation to deepfakes and privacy violations to labour displacement and corporate consolidation—demand dedicated, fit-for-purpose regulatory responses. Expanding copyright control risks distorting foundational copyright principles while failing to address, or indeed worsening, the harms themselves. This is what I've called running into the AI copyright trap. It's mistakenly turning to copyright as a catch-all—or, for some, a windfall—in response to the threats posed by generative AI.

My second point concerns AI training. Some have called for compulsory licensing for copyrighted works that are used in training data, backstopped by owners' rights to opt in or out. I understand the impulse, but the consequences of this approach would, I think, be deeply harmful.

Under the current law, first, it's not clear that training AI on copyright works even implicates the rights of copyright owners. When a system is trained, it translates expressive content into statistical patterns. It turns the meaning into math. This is a technical, intermediate, non-public use to extract information that copyright does not protect. Even if copyright extends to this data extraction and analysis process, most text and data mining is likely lawful without permission or licence under Canada's fair dealing provisions, as interpreted by the Supreme Court of Canada. If the committee is interested in supporting AI research and innovation in Canada, the real problem is legal uncertainty, not illegality.

Requiring licences for AI training would create a pay-to-play system regulated by private actors. The wealthiest corporations could afford access to the vast data troves required, but academic researchers, non-profits, start-ups and SMEs would be shut out, and this would concentrate AI development even further in the hands of big-tech incumbents, which I think is what we're trying to prevent. It would also incentivize secrecy, reduce the diversity of AI systems, exacerbate bias and be practically impossible to administer effectively, as the EU's implementation efforts already reveal.

If copyright reform is required, it should be to confirm that text and data mining for informational analysis does not constitute infringement. This was the original INDU recommendation in the 2019 Copyright Act review, and it remains, I think, the best way to support a healthy AI ecosystem in Canada. It would most likely align with emerging U.S. fair use jurisprudence, but it would also give us the significant advantage of legal clarity. I think Canada's focus here should be on good data governance, not propping up private control of data in a way that's going to send AI development offshore while Canadian creators gain little, if anything.

My third and final point concerns AI outputs. The most effective thing copyright can do to protect human creators is to maintain the position that copyright requires a human author, while AI-generated content is unprotected in the public domain. That is the correct result. It protects the role of human creators in the creative industries, whereas granting rights in AI outputs would be an unnecessary, misplaced incentive that could further chill human creativity.

● (1655)

In closing, I just want to emphasize that copyright law, at its best, serves human creativity and the public interest. It exists because we value what human beings create, share and learn from each other. We cannot allow it to become a tool for controlling technology, a bargaining chip for corporate licensing deals or a vehicle for granting monopoly rights over information or machine-generated content. I urge the committee to keep copyright's principled limits and its practical consequences in view. There are many more apt solutions to the risks posed by AI systems.

Thank you.

The Chair: Thank you very much.

Professor Cukier, you have five minutes.

Wendy Cukier (Professor, Entrepreneurship and Strategy, Ted Rogers School of Management, and Academic Director, Diversity Institute, As an Individual): Thanks so much.

It's really a privilege to be here among such learned and smart people. I will try to supplement what has already been said.

I'm a professor of entrepreneurship and innovation at Toronto Metropolitan University. I was also the vice-president of research and innovation, so I'm very invested in and committed to issues around the commercialization of technology in Canada. At the same time, I'm part of a number of big studies that are focused on responsible use, and I think we've heard a lot about the risks associated with artificial intelligence that have to be taken seriously.

I previously submitted a brief to the AI task force, and I'm happy to provide it to this committee. It reinforced a lot of the points that have already been made about infrastructure development, about sovereignty and the limits of sovereignty, about the urgent need for a regulatory framework for increased risk and to create some measure of certainty, and about the importance of balancing risks and rewards.

What I want to focus on today, though, because I didn't hear anybody talking about them, are the issues around adoption, government as a model user, bias in AI, and skills. I'll try to be brief.

The AI paradox in Canada is that we have a Nobel Prize winner in the development of the technology, yet if you look at us in comparison with other OECD countries, we're laggards in terms of adoption. There are a lot of reasons we can point to to explain that, but one of the most important ones is that we are a country of small and medium-sized enterprises.

We hear a lot about what large corporations are doing. Think about your ridings and who the big employers are. It's not just large companies. Large companies in Canada account for about 10% of private sector employment. What they do is important, but so is what the SMEs do. They provide 90% of the employment, and I think they are often left out of these discussions.

When I talk about SMEs, I'm not just talking about AI start-ups; I'm talking about family businesses in agriculture, in manufacturing, in retail and so on. We really have to grapple with the fact that SMEs in Canada need support in order to grow, address productivity and innovate.

A lot of the focus on AI adoption is around job displacement. That will happen, without question, but that's more likely to happen in large corporations that are using AI to lay people off. Small companies can punch above their weight and can look much bigger than they are if they use AI tools correctly. When we talk about AI tools, we're not just talking about machine learning; we're talking about simple, off-the-shelf services, generative AI and so on. That's one point I would like to emphasize.

Government has a role as a model user. We learned this with the early days of the Internet. Government can do a lot to advance opportunities for start-ups in this space, and I think we see signs that they're moving in that direction.

We have to focus on human capital, and there is a preoccupation with science, technology, engineering and math. They're absolutely critical. We need deep AI skills, and it would be nice to have another Nobel Prize winner, but science, technology, engineering and math are actually what you need to create AI tools.

We need a lot of other skills to advance innovation, and Canada continually makes the mistake of confusing invention with innovation. Innovation is about doing things differently. That means we need lawyers, ethicists and people who understand consumer behaviour, organizational behaviour and markets.

● (1700)

Our biggest barrier to innovation in this country, in my view—I'm biased because I'm in a business school—is the lack of attention on markets and who is going to use the stuff, and for what purposes. While deep AI skills are critical and AI literacy for everyone is important—because all jobs will be affected and all of us need to be protected—the AI skills for innovation, where we take people who understand their businesses and processes and give them the tools to use AI for a responsible purpose, is where I see one of the biggest gaps.

The final thing I'll say, because I am from the Diversity Institute, is that we need to double down on ensuring that AI is not reinforcing bias in the use of biased data and the use of homogeneous teams. We need to ensure that AI is not reinforcing the digital divide we currently see, based on income, geography, indigeneity and gender. We need to be using AI responsibly and inclusively.

I'll stop there. Thank you.

The Chair: Thank you very much, everybody, for your opening testimony.

Mr. Guglielmin, the floor will be yours for six minutes.

Michael Guglielmin: Thank you, Chair, and thank you to all the witnesses for your opening testimonies.

Mr. Dehghantanha, at our committee meetings, we heard testimony that AI has fundamentally changed the nature of cyber-attacks, moving from tools that would assist hackers in their operation to tools and systems that can now autonomously build attack plans, create multiple strategies, troubleshoot, and find workarounds when they generally fail. We even heard an example of AI rewriting its own code to avoid shutdown. We also heard about a third party foreign actor who was able to breach over 100 million data points from citizens in Mexico without a human directing each and every step.

This obviously raises a wide variety of potential national security concerns, and I know that your research sits exactly at this intersection. In plain terms, how close are we to a scenario in which an AI-powered attack could compromise Canadian infrastructure, the financial system or other national security apparatus structures?

● (1705)

Ali Dehghantanha: AI is currently used in both defence and offence. On the offence side, as you mentioned, it provides capabilities that we have never seen in threat actors before. It reduces what we call “mean time to respond”, which means the amount of time that a cybersecurity professional has to respond to an incident significantly. It used to take hours for hackers to meet their objective. These days, it is minutes. If we are not detecting and stopping adversaries in a very short time, they will meet their objectives, as in some of your examples.

The question is, how far are adversaries from building capabilities that can be deployed at a scale targeting all critical infrastructure or all critical services? I can say that they are not that far. We are seeing them in the wild, testing these tools against any available research organizations or the infrastructures that are there.

At the same time, we are actively working with many partners in Canada to build up their skills. My main concern always lies with small and medium-sized businesses, especially in the less protected sectors, such as agri-food. They don't have enough investment to be made and they are widely distributed. An attack that is automated with AI could impact this infrastructure significantly.

When we talk about cybersecurity attacks, everyone thinks about the financial organizations. They would definitely be the first target, but they have tools and techniques to stop these attacks early. When you go down that food chain, getting into other critical sectors—I mentioned agri-food, and health care is another example—you see that the response time in these sectors is very small. We don't even try to build defensive capabilities at the scale that could defend against these AI-based attacks.

Michael Guglielmin: Professor, we've also heard from Anthropic's former safety lead that in perhaps six to 18 months, there are going to be AI models that are capable of long-range strategic attacks.

I remember a story that we were told at this committee about a robotic dog that had, essentially, a kill switch—a button on a wall. It was able to reprogram itself, because it knew that this switch would turn it off. We also heard of scenarios in which agentic AI agents were deployed and then started mining cryptocurrency without being given that instruction.

Given all of this and the timelines that we're being presented with, what can government here in Canada do today to better prepare us for a sophisticated AI national security breach?

Ali Dehghantanha: What I am saying is that in Canada, we are putting a lot of focus on checkboxes—tools that check the AI before it is deployed in the real world. What we are missing is what happens after the AI is deployed. Who is going to monitor it? Who is going to be accountable for that? Who is going to contain the AI's skills?

You gave some examples of AI learning new skills in the field, and that causes complications or adversarial capabilities. You mentioned the timeline of six months to 18 months, and it could be much shorter. I would say that what should be done in a very short time is to invest in building that control plane, a layer that would sit between the AI application and foundational models, and try to control it. Give control back to the owner, to the human, to the operator—whatever we want to call it. That control plane is currently the missing layer in AI adoption.

We are not going through a slow adoption of AI. We will still let the applications be built, but the control plane should be built, and we need to have regulations and rules around that.

Michael Guglielmin: If AI significantly lowers the barrier to entry for cyber-attacks, are we entering a world where less sophisticated actors are going to be able to deploy this technology and cyber-attacks against national security infrastructure would actually ramp up?

● (1710)

Ali Dehghantanha: What we are observing in the field is mostly that sophisticated adversaries now have much better tools and much better capabilities that they could deploy in a much shorter time, and that they outpace the targets that smaller adversaries were going after. That is what is actually happening with AI.

What is happening, I would say, is that the age of lone attackers or a small group of attackers being successful in attacking our infrastructure is gone. Most of those infrastructures are being targeted by advanced attackers that already have that automation at scale.

That makes me more worried, because previously, if I was talking to a farmer a few years ago, I wouldn't even think that an adversary from Russia would target them. These days they may, because everything is automated. The same ransomware, the same malware, that AI is now dropping, could end up on, say, a dairy farm. You would have never seen that in the past.

Michael Guglielmin: Thank you very much.

The Chair: Thank you.

Mr. Ma, the floor is yours for six minutes, please.

Michael Ma (Markham—Unionville, Lib.): Thank you to all of the witnesses.

My first question is for Dr. Cukier.

Following up on your last point about the digital divide and bias, we all know that AI basically functions on large data, whether that's language or biological data—and everything else.

You mentioned the digital divide. Certain parts of the world have, unfortunately, less input into and therefore less interaction with the AI model. Eventually, we're going to see a skewed model in which certain ethnic groups or certain geographical representations are lacking in that environment. If we depend on AI to make decisions legally or medically and so forth, that's going create a further digital divide and create a more unjust environment globally. Can you speak to that a bit more, please?

Wendy Cukier: Sure.

The issue of the digital divide, as many of you know, is not new at all. What we learned during COVID—because people talked about adaptation during COVID with rapid digitization and everything moving online—was that, for example, indigenous people in rural and remote communities have less access. Most people know about that, but did you know that 42% of racialized children in the city of Toronto were doing their homework on iPhones because they didn't have access to high-speed Internet, computers and so on?

You can take those principles and understand that the digital divide is not just about physical access to broadband. It's about broadband. It's about affordability. It's about devices. It's about skills. I'll be the only boomer in the room, and I'll tell you that I am much more vulnerable to the misuse of AI because I answer my phone and think it's a human, or I look at a video and think it's real. We need a sophisticated understanding of, first of all, the dimensions of the digital divide, and then the ways in which artificial intelligence applications in all of their manifestations, for good and for evil, will have an impact on that.

The interesting thing that came out in a recent survey we did with Environics is that the gap between men and women in the use of AI tools—not the developers—is much smaller than we would see with other technologies. Indigenous people are using AI tools more than others in the population. Immigrants are using AI tools more than others and so on. It's interesting, because to me this signals that there are ways in which AI can bridge some of these gaps.

I referred to the discipline differences. In some ways, AI is the English major's revenge. You don't need coding and you don't need a background in computer science to be able to build tools. That's something we have to really pay attention to when we're thinking about our national AI strategy. We need a responsible AI for all approaches, in my view.

Does that answer the question?

● (1715)

Michael Ma: Thank you.

My follow-up question relates to that as well.

You have talked about racial bias as well, such as in HR and health care. I know experience-wise, from the last couple of years of observing, that HR hiring practices or recruitment agencies use AI for screening, and therefore there must be some inherent biases built into it. Can you talk a bit more about that? How do we ensure that we have a much better inclusion environment?

Wendy Cukier: That's a really good point, and it's not just racial bias. It's gender bias as well. Even tech firms have been stymied in their efforts to level the playing field.

There are two things here, and one is garbage in, garbage out. The data that you use, if it's biased, is going to replicate bias. The second is making sure you have diverse teams that are sensitive to these issues.

A third thing is disclosure. I think disclosure is absolutely critical.

A fourth thing is “human in the loop”. I'm working with a number of public sector organizations, for example, that are experimenting with large-scale AI tools. One of the critical things is to do experiments where you compare the results you get from the AI-enabled processes to what you would get with humans. Then you try to figure out if AI is amplifying bias or reducing bias, because sometimes it will cut out the bias that's associated with “we play golf together” or “I went to Queen's University”.

There are huge opportunities for good and evil, in my view, but transparency, human in the loop and inclusion are fundamental principles.

Michael Ma: Great. Thank you very much.

My next questions are for Professor Craig.

The Chair: Mr. Ma, we're at time.

Michael Ma: Okay. Time flies.

The Chair: It sure does. You might have a chance to come back toward the end.

[*Translation*]

Mr. Ste-Marie, the floor is yours for six minutes.

Gabriel Ste-Marie: Thank you, Mr. Chair.

Welcome to the three witnesses. Thank you for being here and for your presentations.

Ms. Cukier, in your presentation, you said that it is important for small and medium-sized businesses to embrace artificial intelligence technology in order to increase their productivity. That seems to come with some challenges.

Here is my first question. In your view, what are the obstacles preventing small and medium-sized businesses from incorporating artificial intelligence into their activities? Are they financial, technical, cultural or regulatory?

[*English*]

Wendy Cukier: Yes, it's all of those things.

We know, especially post-COVID and especially since the trade wars, that small and medium-sized enterprises are struggling. They have narrow margins. Most small and medium-sized enterprises in Canada have fewer than five people, so the person who's doing the technology development is also doing payroll and taking out the garbage. They lack the inherent skill, capacity and so on. That's one piece.

It's also the investment, for sure, although sometimes the barriers to entry are not that large. Unfortunately, when we talk about technology, people who love the technology talk about the technology, not about what it's good for. We need more use cases that show simple applications very clearly. For example, you can take a stack of expenses and receipts and turn them into a spreadsheet in five minutes instead of six hours. We need concrete, simple examples. We have them; it's just that they're not widely shared.

We've done research in Quebec as well as across the country, and often, small businesses have a short-term horizon rather than a long-term horizon. For the programs the government has implemented aimed at advancing technology adoption, if they're targeting SMEs, the benefits have to be more than the costs. By this I mean that if you're giving someone a small amount of money to implement technology, you need to make sure they can get it without a lot of trouble. You can still introduce accountability and have audits to make sure they did what they said they were going to do, and you can do evaluations, but if you try to incentivize upskilling, adoption of technology or infrastructure investments, it's important that you make it easy to get access.

The other thing we need to think about, because we're going to spend a lot of money on major projects and infrastructure, is how we can leverage those investments to provide opportunities for small and medium-sized enterprises to modernize, upskill, re-skill and so on. I think there are some clever ways we can get more bang for our investment.

● (1720)

[*Translation*]

Gabriel Ste-Marie: Thank you very much. That is very clear and very complete. I am grateful to you.

I have another question, which is still about the use of artificial intelligence to increase productivity in small and medium-sized businesses, and large businesses too. In your view, which sectors have the greatest potential for adopting this technology? What could the government do to help those sectors adopt artificial intelligence?

[*English*]

Wendy Cukier: Honestly, it's a question of short term versus long term. I don't think any sector will not be affected. When I think of priority sectors from an economic point of view—they've already been mentioned—they're things like manufacturing, energy and construction infrastructure. There are interesting physical AI adoption opportunities, but they are also typically highly capital-intensive.

Where are our pain points in Canada? They're in health care. There are huge opportunities in health care, if we can manage the risks. Agriculture was already mentioned. Building self-sufficiency in agriculture is about not only large-scale farms but also vertical gardens and all kinds of things.

We need sectoral strategies for AI adoption. We need to recognize that in almost all sectors, except maybe finance, IT and some manufacturing, SMEs are at the core.

[*Translation*]

Gabriel Ste-Marie: I have a question on another matter: protecting personal information. You can tell me whether you are uncomfortable answering, given that we are changing the subject.

In your view, what role should artificial intelligence play in managing, conserving and disclosing information in federal institutions with links to the government?

[*English*]

Wendy Cukier: That's a really good question.

I am in between what I would call protectionism and the free market when it comes to data. I know that in the province of Quebec, there has been a lot of progress made in data regulation. Looking at that carefully to figure out what's on paper and what has actually been implemented might help us understand what we should be doing on a national level.

My view—and this is a bit different from what you may have heard from others—is that often we dichotomize individual rights to information and privacy versus corporate interests in making money. When you look at regulations, often they're here or they're here. I think there is a third piece, and that is the public interest.

The analogy I would encourage people to think about when we think about data and think about AI is our tax system. I earn money. It's not all my money. The government takes a portion of it to advance the public interest. When we think about security and think about health care, there are many things for which the government having access to some of my data will actually benefit all Canadians. We have to figure out how to balance those interests in an appropriate way so that we are advancing our economic development, innovation and trade; protecting privacy; and helping the government do a better job for all Canadians.

● (1725)

[*Translation*]

Gabriel Ste-Marie: Thank you very much.

The Chair: Thank you, Mr. Ste-Marie.

[*English*]

Ms. DeRidder, the floor is yours for five minutes.

Kelly DeRidder: Thank you.

Dr. Dehghantanha, my questions will be for you today.

I was recently briefed on the reports saying that Anthropic's Claude AI model was used by a suspected Chinese state-sponsored group to conduct a large-scale, automated cyber-espionage campaign on roughly 30 organizations globally. Would you mind sharing with the committee what you know about this incident and any suggestions or recommendations you may have on how to mitigate this type of attack in the future?

Ali Dehghantanha: Sure.

The background of the case, as you mentioned, is one of Anthropic's tools. AI systems have been used for automating data to do what we call data exfiltration, which means receiving data from the network, by what is believed to be a Chinese adversarial operator. I can tell you, based on all the knowledge of and what we understand from the attack, that the main reason this specific system was used was for ease. It is more focused on optimizing code, and it is much more accurate. That is what the adversaries were after, but it doesn't mean that this capability is only limited to Anthropic or to Anthropic systems.

Adversaries are normally using AI automation based on the objective that they have in mind. For example, if they want to steal copyrighted information, they may choose OpenAI platforms because those are better in text recognition. If they are going after codes, they may go with Anthropic systems.

That's more context. What I want to highlight is that usually, from the adversary's point of view, they don't care who is behind the AI. They're more interested in the skills or the capability of an AI technology or AI system.

The second part of your question, I believe, was more on what we can do. Is that right?

Kelly DeRidder: I agree with you completely. I was using Anthropic as an example because it just recently published a report, but it is open-source AI.

Yes, the second part is more what I'd like to talk about. How do we mitigate that risk in the future?

Ali Dehghantanha: As I mentioned in the answer to the previous question, currently the response time by cybersecurity people is squeezed significantly, which means that the moment adversaries are in, it takes minutes for them to achieve their objectives using AI and automation.

What we should do at the enterprise level is create a layer, as I mentioned, between AI applications and AI foundational models, a layer that is controlled by the company, by the enterprise. Even if an adversary wants to use, say, Anthropic or any other foundational model capability, they will need to go through that control layer and will hopefully be detected on that layer. That's one thing we can do at the enterprise defence level.

On the other side, we need to invest significantly in the detection of the deployment or usage of AI that is not approved by an enterprise. These days, the technology available is very limited, even to enterprise organizations, big organizations, in identifying whether a specific skill or a specific activity that is done by AI is legitimate or is following their policy. Advancement on that could significantly help us to identify what is allowed and what is not allowed in the system. That could limit an adversary's capability once they are in the network.

Kelly DeRidder: Thank you for that.

In my community of Kitchener Centre, Canada's innovation capital, we see advanced manufacturing and tech innovation go hand in hand. How can Canada better align its cybersecurity and AI strategies, both to protect these industries and to ensure they remain competitive globally?

The Chair: Professor, I'm going to ask you to keep that answer to about 45 seconds. Thank you.

Ali Dehghantanha: Sure.

When AI missed manufacturing...we are looking at the risk of physical AI, which means that devices are able to take and make new actions and do new activities. What we need to focus on is controlling what kinds of skill sets and what kinds of actions AI systems can do. We need to support both the start-ups in your area that are focused on controlling AI and the start-ups that are more focused on how the policy can be applied and integrated into AI actions.

The Chair: Thank you very much.

Colleagues, we are running over time, but I'm going to give five minutes to Madame O'Rourke, a minute to Monsieur Ste-Marie, two and a half minutes to Ms. Borrelli and two and a half to Mr. Bains to finish.

Witnesses, I know that's a bit of an audible, so if you do have to go, we understand, and you're certainly welcome to excuse yourself. Otherwise, I hope you don't mind sticking with us for an extra 10 to 15 minutes. We're very much appreciating the insights.

Madame O'Rourke, you have five minutes.

● (1730)

Dominique O'Rourke: Thank you, Chair Carr.

Dr. Dehghantanha, it's nice to see you again. Thanks very much for an earlier conversation at the University of Guelph.

The University of Guelph has the Centre for Advancing Responsible and Ethical Artificial Intelligence, as well as the AI for Food initiative. Given that Ms. Cukier was talking about having a sectoral approach to AI adoption and that we tend to be thinking about AI in terms of the financial sector, white-collar jobs and perhaps advanced manufacturing, can you tell us what the potential is for AI in agriculture, some of the pitfalls you can see and then what measures we would need to consider now, including perhaps the right to repair?

Ali Dehghantanha: AI has immense capabilities. It's already changing the practices in agri-food significantly.

You mentioned a couple of the centres we have at the University of Guelph that are focused on helping farmers to build more AI capacities at their farms, both at their operation and at distribution. You are seeing that AI is now disrupting technology at all of these three layers. Farmers are integrating AI into their on-farm operations, from sensing all the way to controlling livestock—everything at that level. Then, as you would see at the operation level, AI is now getting a lot of inputs from different data sources and is helping farmers to optimize or improve their practices. That's the operational level.

When it comes to distribution, you will see that bigger companies, like Sobeys and the like, are using AI significantly to manage what they should buy from which farmers at what price. Also, the other way around, it's used by the farmers as well. What I'm trying to say is that in 12 to 18 months, I would say, you will see AI going from the farm all the way to the table. The whole ecosystem is now being built around AI in agri-food.

You mentioned what measures we should put in place. One of the main points of the agri-food sector is that most of the operators in this sector are small and medium-sized businesses that are physically distributed across the country. Being able to secure them, protect them and make sure they are using AI responsibly requires a standard and requirements by law through vendors so that as they deploy these AI solutions, they deploy them in a way that they are responsible [*Technical difficulty—Editor*] control. That's what we don't have at the moment.

Some examples have been mentioned in this meeting. At the moment, if an AI agent that's able to create new skills is released on farm, we don't even know when and where in this situation it will gain skills and what it will do with those skills. We are always hoping for the best, but there is no testing and no benchmark for evaluation before deployment, for when it is in use and then after deployment for what you should do when you want to just kill this system, kill this data.

Dominique O'Rourke: Thank you.

I have another question. When speaking with Dr. Beth Parker from the groundwater research centre at Guelph, I asked her whether AI will allow her to accelerate discovery. She said, “We still have to go and get the core samples. We still need the data.” That goes to Ms. Cukier's earlier comments.

I'm struggling with the timelines we're discussing. Sometimes it's 12 to 18 months or three to five years. Where are we in terms of good solid data? That's not for things like ChatGPT, but things like medical research or advances in agriculture. How close are we to that? How close are we to having good data collection? How close are we to having a secure layer in order to monitor, identify challenges and address them? How are these things coming together over the next three to five years?

Ali Dehghantanha: In terms of data, I would say the best people to talk to are the domain experts, such as people in AI and Beth Parker at the water centre you mentioned. The good thing about AI is that once the data is collected at one point in the world, we can start using it everywhere else. I am seeing a lot of investment being made by core AI companies on generating reliable data. For me, the timeline is quite short if you have a global view. If you have a regional view for specific places, yes, that would take a lot longer.

In terms of how advanced we are in securing AI systems, we are at the very earliest stage. I have yet to see any enterprise in Canada, and we are working with many of them, deploying any control layer for AI. That becomes the main challenge for them in deploying AI. Set aside thinking about smaller businesses or smaller organizations. That's a huge gap we are seeing there. That doesn't exist.

● (1735)

Dominique O'Rourke: Thank you.

The Chair: Thanks, Madame O'Rourke.

[*Translation*]

Mr. Ste-Marie, you have one minute.

Gabriel Ste-Marie: Thank you, Mr. Chair.

Ms. Craig, are you familiar with the European legislation on artificial intelligence in terms of copyright? If so, can you provide us with some comments in one minute?

[*English*]

Carys Craig: I have. This committee may know that Europe was quite an early mover on creating rules around exceptions to copyright to permit text and data mining. They had a tiered approach, where research institutions and cultural organizations could engage in text and data mining with copyright works without any liability risks for non-commercial research purposes. As a second tier, which is outside of that, there is an exception for text and data mining with the possibility of opt-outs by rights holders.

It was an early move and it was controversial. In a way, it suggested that rights holders can, unless an exception applies, prevent text and data mining by exercising the opt-out. The tricky part has been trying to work out how that opt-out can be exercised, by whom, the force and effect of that, and how it can be implemented.

It has created some challenges. Of course, that is still being worked out in the European context, although the rule itself has become the rule through the EU AI Act. The question that faces us is whether we want to follow suit or hold up and see how this unfolds and whether it's the best option. Of course, there are tensions and incompatibilities to consider with the other rules that are emerging in the U.S.

The thing I want to stress here is that copyright, while it might seem like a side issue, has been described as an issue that could bring AI to its knees. We're not just talking about movies, books and things that we think of as copyright works. Everything—all of the datasets and all of the data that's being scraped and used and on which these machines are trained—can be subject to the very low automatic protection of copyright.

If you create, at the kind of scale we're talking about, an obligation to license or a right to opt out of training, then you create huge obstacles to the capacity to access the data you need to train AI systems well. You create a system where there are limits in the datasets and where there are some key players that can access the data that's needed and there are others who cannot. It could become a true obstacle to the development of AI in Canada and beyond.

[*Translation*]

Gabriel Ste-Marie: Thank you.

The Chair: Thank you, Mr. Ste-Marie.

[*English*]

We'll go to Ms. Borrelli for two and a half minutes.

Kathy Borrelli (Windsor—Tecumseh—Lakeshore, CPC): Professor Craig, hello. Thank you for being here today.

Canadian SMEs are the backbone of our economy, yet some lack the capital, the expertise or just the access to infrastructure that's needed to use AI. Do you believe current government policies are doing enough to ensure that SMEs can not only adopt AI but also capture real economic value from it?

Carys Craig: The biggest challenge that I see potentially facing smaller Canadian competitors and innovators that want to rely upon or develop AI in this economy and in this context is the concern about what they can lawfully do and lawfully use, and then how they can implement and put to work the AI tools available to them.

To be honest, the incentives to use AI and the availability of good public systems mean that there is a growing capacity for everybody to take advantage of generative AI tools in particular and to find efficiencies that can benefit them. In this regard, I think education and access can be key.

To the extent that people want to be able to develop their own tools that maximize their capacities in their own sectors and for their own purposes, that's when you see both the need for technical supports and the accessibility of the data and tools becoming key. We could spend more time thinking about how we prop up and support the development of open-access and open-source models of data commons that are accessible to small movers and innovators that want to take advantage of that, rather than thinking about how large rights holders can block and prevent training on their data.

Rather than thinking about how we can exclude, we can think about how we include people in the data and how to ensure access to data and the technology it allows.

● (1740)

Kathy Borrelli: Thank you.

I wish I had more time for another question.

The Chair: Thank you, Ms. Borrelli.

Mr. Bains, you have the floor for two and a half minutes.

Parm Bains: Thank you, Mr. Chair.

Ms. Cukier, you talked about having the talent and human capital. If you look at budget 2025, it announced $1.7 billion to recruit top international researchers. What are your thoughts on whether this could assist in strengthening recruitment and retention in the AI space? How do federal investments in research infrastructure and talent contribute to the commercialization of the industry and the success of it?

Wendy Cukier: When we think about talent, think about a pyramid. The bottom is AI literacy for all. The top is the deep AI skills to build the tools. The middle, though, is much bigger than the top, and a lot of other things are required.

The strategy to attract and retain the best and brightest from around the world is a very good one, particularly as we're seeing barriers to international scholars, for example, who want to work in the United States. That's a huge opportunity for Canada to scoop up talent. I think that's very important.

Parm Bains: You said something there that we heard from other witnesses. There's already an AI of things. You talked about small and medium-sized businesses, agriculture and health care—things that we can probably target quickly. That's the bigger picture, maybe, that I went to first. Maybe the smaller picture is something we can focus on.

Wendy Cukier: I want to say "and". It's about attracting the best and brightest from around the world, like our Nobel Prize winner and Dr. Bengio, who spoke earlier, who are both immigrants to Canada. That strategy works. However—the "and"—we also need a strategy that really focuses on the tools and the talent to create AI adoption. That's a different population and a different skill set, and it's not all science and technology.

My view is that, yes, we need international scholarship, but if you're looking at AI adoption, it's a different set of skills. We have not paid enough attention to developing those skills and providing that support.

Think about our student workplace placement program, the co-op programs and the support for student internships. We're sitting on a potential gold mine of mobilizing young people to get into SMEs. The kinds of skills needed to help an SME become more productive are not those of someone with a Ph.D. in computer science. We need a more comprehensive national skills strategy when it comes to AI.

Does that answer the question?

Parm Bains: Yes. Thank you.

The Chair: Thank you very much, Mr. Bains.

I'm going to take the final two minutes and use my prerogative as chair to ask a quick question.

Professor Craig, this question is for you.

I come from the world of education. I had the chance to spend some time with a couple of high school classes in my riding a few weeks ago, and I told them that at the industry committee we were delving into the future of industry as the disruption of AI hits us. What I hadn't considered, but what a lot of these young 16- and 17-year-olds brought up with me, was how this related to things they were designing: visual designs and artistic designs of a variety of different kinds. They're creating them with or without the assistance of AI, but AI is taking what they've created, and they're very concerned about how to preserve their rights, their intellectual property.

Can you just speak directly to young people—I can assure you they're not watching, but I'll bring the clip to them—and tell them what they should be thinking about and what you want us to be thinking about in regard to how we can be protective of intellectual property as this technology develops? I ask because creativity is such an important source of not only inspiration but also monetization in this day and age for young people.

● (1745)

Carys Craig: I have a couple of high-schoolers myself, and I can assure you they're not watching.

Some hon. members: Oh, oh!

Carys Craig: The changing shape of creativity in this AI era is a fascinating thing for us to consider, especially watching the way that high-schoolers and others interact with the technology and the way it influences their learning and creativity in positive and, of course, negative ways, which are of concern to someone working in the education sector.

We do see, in some sense, a democratization of the capacity to create. The things that high-schoolers can create with the help of some of these AI tools for their images, PowerPoint slides and study notes are truly impressive developments.

At the same time, as you suggest, there are concerns about what this is going to mean for the future of learning and education. Also, what does it mean for the future of creativity? If everybody can create an output that looks like it was made by a leading graphic designer but was in fact created by my 15-year-old, what does that mean for the future of graphic design? These are very real questions we have to confront.

I want to again caution on the relevance of copyright and intellectual property to those questions. I have suggested that we should restrain the expansion of copyright law such that it doesn't protect the outputs that are generated by these machines. I think that's the best way to make sure that the copyright system encourages human authorship and human creativity rather than just incentivizing the mass creation of a glut of AI-generated outputs. The copyright system should remain focused on encouraging and incentivizing humans to create using human skill and judgment and intellectual capacity.

As for how to protect that once it is created, I take you to be saying that it's going to get sucked back into the system and become the training data that others use. I don't think our focus should be on the inputs to these systems. I think when artistic works, outputs or creations are used to train an AI system, that's part of the black box of how these algorithms work.

Our focus should be on protecting human authors against the substitutes in the market, the outputs of AI that might be substantially similar to those works and might compete with them and undercut the market. That's where copyright law should continue to work, as it has always worked. If the outputs are infringing, then the outputs will continue to be infringing, but limiting the capacity of AI to be trained on art, music and literature would be a serious obstacle to its development—and it kind of misses the point.

The Chair: Thank you very much. It's certainly a fascinating realm.

Thanks, everybody. That was a great discussion today. I think we're learning a lot and positioning ourselves well to provide some feedback to the government moving forward.

Colleagues, I have one quick thing, which is the budget. This is standard operating procedure. In the budget for the meeting, there was a line item in which we underestimated how many meetings there would be. We were basing it on six, but we'll end up having eight, so we need approval for a $381 increase to the costs associated with those two additional meetings.

Some hon. members: Agreed.

The Chair: I'll take that as unanimous consent, something AI could not do.

Thank you very much, everybody. Have a wonderful rest of your day.

The meeting is adjourned.