45th PARLIAMENT, 1st SESSION

Standing Committee on Access to Information, Privacy and Ethics

EVIDENCE

NUMBER 020

Monday, December 1, 2025

Chair: John Brassard

Standing Committee on Access to Information, Privacy and Ethics

Monday, December 1, 2025

● (1100)

[*English*]

The Chair (John Brassard (Barrie South—Innisfil, CPC)): Good morning, everyone. It's December, and I'm going to call this meeting to order.

I want to welcome everyone to meeting number 20 of the House of Commons Standing Committee on Access to Information, Privacy and Ethics.

[*Translation*]

Pursuant to Standing Order 108(3)(h) and the motion adopted on Wednesday, September 17, 2025, the committee resumed its study of the challenges posed by artificial intelligence and its regulation.

[*English*]

I'd like to welcome our witnesses for the first hour today. Both are from Conjecture Ltd. We have Connor Leahy, who is the chief executive officer, and Gabriel Alfour, who is the chief technology officer.

Mr. Leahy, you have up to five minutes to address the committee. I understand that you may need a bit more time or want a bit more time. If it gets up to six minutes, I would accept that, but I know we have lots of questions to ask.

Mr. Leahy, go ahead, please.

Connor Leahy (Chief Executive Officer, Conjecture Ltd.): Thank you, Mr. Chair and members of the committee, for inviting me to testify today.

I'm an expert on the catastrophic global threats of AI and will primarily be speaking to you from this perspective.

I am the CEO of Conjecture, which is an AI safety research firm. I'm also an adviser at ControlAI, which is a non-profit focused on mitigating the security risks posed by advanced AI.

In 1985, humanity awakened to a hole in the sky. Scientists discovered that chlorofluorocarbons, CFCs, were depleting the ozone layer, which shields humanity from damaging ultraviolet radiation. At the same time, humanity also lived atop a deep fracture—a cold war between the U.S. and the U.S.S.R. that threatened nuclear annihilation.

Amidst deep geopolitical tensions, the two superpowers ultimately shook hands, signing both a landmark nuclear de-escalation treaty and the Montreal Protocol in 1987 to prohibit and phase out CFCs. This protocol ultimately received universal ratification. De-

spite the world's divisions, these rival powers came together to mend a hole in the sky and to recognize that never-ending nuclear escalation was in no one's interest, and the rest of the world followed.

In 2023, humanity heard a new warning call from Nobel Prize-winning AI scientists and the CEOs of major AI companies, saying, "Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war." This risk of extinction is posed by superintelligence, the exact subset of AI that the leading AI companies are racing to develop.

Superintelligence is defined as AI that is more competent than all humans at all relevant cognitive tasks across all relevant domains and capable of acting beyond human oversight and control. If there were to exist systems that autonomously out-compete any human in all relevant tasks of science, business, persuasion, politics and warfare, and if we did not control them, it is hard to imagine a future that goes well for humanity.

A major part of the risk is that AI developers fundamentally do not understand how the AI systems they are creating actually work and cannot develop them in a safe manner. Dario Amodei, the CEO of the second-largest AI company, recently stated that we perhaps "understand 3% of how they work", which is, in my personal opinion, somewhat of an overestimation.

AIs are not developed as code that is written line by line as we do with traditional software. Instead, researchers are essentially growing AI models by feeding them vast amounts of data and training them by using enormous computing power to produce what is called a neural network rather than a set of lines of computer code.

Unfortunately, the current AI development paradigm does not allow the safety-by-design approaches that we use for other advanced, highly risky technologies. We would not, for example, build nuclear power plants if we did not know how to control nuclear reactions. Technical control methods are lagging drastically behind the advancement in AI systems capabilities. Currently, there are no legally binding AI safety regulations to protect consumers and humanity as a whole.

Where does this leave us today? Right now, multiple AI companies are pouring hundreds of billions of dollars into developing superintelligent AI as quickly as possible despite experts warning of the risks. This haste is, in my opinion, directly tied to an attempt to outrun legislation to complete their projects before the wider public and the government wake up to the completely unconscionable risks the unconsenting public is being exposed to by private, oversightless and reckless actors.

Recently, AI companies have been racing to automate AI research itself, allowing AIs to build even better AIs by themselves in order to reach superintelligence more quickly. This process is called recursive self-improvement, meaning the moment an AI is built that is good enough to make better AIs, it might already be too late.

Leading scientists now estimate that superintelligence could be developed by 2030, or potentially even sooner. In the face of this pressing threat from superintelligence, I'd like to offer the committee three recommendations for how Canada can respond now.

One, the Canadian government should publicly recognize superintelligence as a national and global security threat that poses an extinction risk to humanity.

Two, Canada should begin negotiating an international agreement to prohibit the development of superintelligence, given that no scientific consensus can be developed in a way that does not threaten humanity with extinction. The agreement should also restrict and monitor superintelligence precursors such as recursive self-improvement.

● (1105)

Three, Canada should prevent the development of artificial superintelligence on its soil, as superintelligence would be capable of overpowering individuals, companies and even Canada's national security apparatus.

Thank you. I would be happy to take any questions you may have.

The Chair: Thank you, Mr. Leahy.

Mr. Alfour, you have up to five minutes to address the committee. Go ahead, please.

Gabriel Alfour (Chief Technology Officer, Conjecture Ltd.): Mr. Chair and members of the committee, my name is Gabriel Alfour. I'm the chief technology officer and co-founder of Conjecture, an AI safety research firm. I also helped found ControlAI, a nonprofit dedicated to preventing risks to humanity from artificial intelligence. ControlAI has engaged lawmakers in Canada, the U.S., the U.K. and the EU.

There are many complex and important challenges we face with AI, but in my personal and professional opinion, the most urgent one is the extinction risk posed by superintelligent AI. These are systems that vastly exceed human cognitive abilities and would be capable of out-competing us in scientific and military development, persuasion, politics, business and more. They would outsmart not just individuals, but corporations, national security establishments and governments. If built, they, not us, will be the force deciding the future.

How did we get to this point with AI?

First, the top experts from the field—the most cited AI scientists and the CEOs of the leading AI labs—warned in 2023 that "Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war." However, said warnings were ignored. Leading AI companies are still recklessly pursuing superintelligent AI systems capable of outsmarting our best technology, engineers and national security experts, and of resisting being shut down. Their plans to control superintelligent systems are at best ungrounded and speculative—when they exist at all.

Second, there is a common misconception about AI development that we directly program how these systems behave, but we don't. We did until about 15 years ago, but modern AI systems are grown, not built, by being fed massive amounts of data, and their behaviour emerges in ways that we cannot predict or control. That is, AI is not coded line by line by humans, and researchers and engineers do not need to understand AI to create it. When AI systems encourage a young person to commit suicide, deceive their users or resist being shut down, no engineer programmed this. This is the consequence of not knowing how to diagnose what led the system to do this or how to reliably prevent it from doing so again.

Finally, as of today, artificial intelligence remains the exception, not the standard, when it comes to how high-risk industries are regulated. To operate in fields like nuclear and biotechnology, developers must comply with stringent safety standards, implement risk mitigation strategies, submit to inspections and so on, yet the AI field remains largely unregulated despite mounting concerns from within the industry. AI engineers in SF have told me they do not understand what they are building, and some see what they do as clearly dangerous. Even Geoffrey Hinton, the "godfather of AI", left Google specifically to warn about the risks of AI.

What can be done to prevent said risks of extinction from artificial superintelligence? It is my belief that countries should not unilaterally act against their own interests, much less on blind trust. Instead, they must do two things.

First, at the national level, they must halt the development of the most dangerous AI systems, namely superintelligent systems. Every country stands to lose from the development of superintelligence and to benefit from domestically halting all programs developing superintelligence. Such systems, once deployed, could not be shut down and would outperform every human at hacking and other tasks, thus threatening the national security of countries.

Second, at the international level, countries must agree to regulate and monitor the precursors to superintelligence. We should apply the same regulatory approach used for dual-use technology like nuclear, biological and chemical materials, and prohibit development programs capable of egregious harm outright—in this case, artificial superintelligence—while regulating their precursors. This will allow beneficial applications to thrive while preventing catastrophic harms.

Determining which precursors' capabilities to regulate is a moving target that will evolve alongside our understanding of AI. Some precursors are, unfortunately, dual-use. Computer and data centres are economically beneficial, yet critical to developing superintelligence. Similarly, hacking capabilities offer military advantages, but could enable AI to break containment. For such dual-use precursors, international agreements are essential. No single country can mitigate these risks alone, nor should one country bear the economic cost of restrictions while others forge ahead.

● (1110)

Meanwhile, some precursors have narrower applications limited to AI research itself, such as systems capable of autonomously advancing AI research without human oversight, which could trigger an unchecked feedback loop of capability improvements.

Canada can also act domestically to neutralize dangerous AI systems within its borders. For example—

The Chair: Mr. Alfour, I'll have to stop you there, because I know members want to get to questions. Perhaps you can answer some of the questions with the remaining statement you have.

I also want to make sure that both of you are on your language of choice, because there will be questions posed in English and French.

[*Translation*]

We will begin in French.

Mr. Hardy, you have the floor for six minutes.

Gabriel Hardy (Montmorency—Charlevoix, CPC): Thank you, Mr. Chair.

Gentlemen, thank you for joining us today.

This is an extremely important topic. We are only at the beginning of our study, but the witnesses who have appeared before the committee to talk about artificial intelligence seem to have very wide-ranging opinions. Some have talked about the very risky and dangerous side of artificial intelligence, while others have been very positive about the benefits it has in store for us.

I would like to draw a parallel between artificial intelligence and what we are seeing with social media. Private businesses have been allowed to directly develop social media at breakneck speed, with a kind of artificial intelligence baby that analyzes everything we look at to try to keep our attention focused on these networks at all times. Young people are experiencing unprecedented levels of stress, largely due to this.

How would you compare the early days of artificial intelligence with social media and with the superintelligence that is currently being developed, which you spoke about earlier?

● (1115)

[*English*]

Connor Leahy: I'll take this one.

A lot of the early research that is leading to the current boom in artificial intelligence started in the context of social media. A lot of the early research on what is now called deep learning and AI was done for social media recommendation algorithms.

Personally, I'm in a younger generation, the oldest gen Z generation. I remember we had a promise, in a sense, that if we just let social media and the Internet run free, if we didn't regulate, it would bring freedom and prosperity to the world. I don't know, but perhaps some members remember, for example, the Arab Spring. There was a widely held belief by me and by many other people at the time that widespread access to the Internet and social media would bring freedom and democracy.

These promises have turned out to be lies. They have not come true. They instead are being used by social media companies to cannibalize many aspects of our interpersonal communications and interpersonal relationships for their personal benefit. They are now pouring hundreds of millions of dollars into lobbying and other forums to try to prevent people from interfering.

I think the same pattern of behaviour—developing a technology so quickly that governments cannot react and actively trying to slow down governments to prevent them from regulating this technology until it's already too late—is exactly the playbook we are seeing being deployed right now by AI companies.

[*Translation*]

Gabriel Hardy: If I understand correctly, we are basically repeating past mistakes with social media. Incentives are driving companies to develop powerful artificial intelligence, or artificial superintelligence, as quickly as possible in order to have the upper hand on the market.

In your opening remarks, you compared the evolution of nuclear power with the Montreal protocol. In your opinion, what steps should be taken to ensure that governments intervene as quickly as possible in the field of artificial intelligence? What needs to be done to ensure that we understand the dangers and benefits of artificial intelligence, so that we can intervene as quickly as possible to keep humans at the centre of the mechanism and ensure that artificial intelligence is developed only for its positive aspects, and its negative aspects can be controlled?

[*English*]

Connor Leahy: I think this is exactly correct. I think that in many ways we are making, to some extent, the same mistake, and it's important that we do not make the same mistake. Therefore, fast action by government is extremely important. As I stated during my initial statement, the most important thing is to fully arrest the development of truly superintelligent, dangerous AI.

This is not an esoteric, small corner of the industry. This is a thing you can see being advertised by these companies' research departments as their primary goal. They go to parties and brag about building superintelligence. This is not a secret operation. This is a widely held thing, and it is something the government can act on now. Already, just acknowledging these risks and bringing them into both national and international discourse are the first steps to stigmatizing and potentially outlawing such dangerous developments, while opening the negotiating table for how to handle dual-use precursors.

This is a very difficult regulation challenge. This is why communications like those we're having today are so important. The first step, from my personal perspective, is the prevention of the creation of superintelligence both nationally and internationally, and then it's about moving towards sensible regulation of dual-use technologies and building on lessons that have been learned in other high-risk technological areas.

[*Translation*]

Gabriel Hardy: Thank you.

I would also like to hear your comments on this subject, Mr. Alfour. You talked about risks in your presentation. However, I believe there is another risk that you did not mention. You said that Canada should stop research aimed at developing artificial superintelligence. However, if Canada slows down, isn't there a danger that other countries will take the lead and we will ultimately fall behind and suffer the consequences? I imagine that's the problem: If one country does it and another doesn't, it triggers a mad rush for all countries.

How can we respond to this challenge? I think that pulling Canada out of the race will put us in a precarious position. Is there a way to get everyone to agree and really move in the right direction with regard to artificial intelligence, and artificial superintelligence in particular?

● (1120)

The Chair: Please give a quick answer in 20 seconds or less, Mr. Alfour.

[*English*]

Gabriel Alfour: I think you outlined three separate concerns. The first one is that stopping the development of ASI does not hurt Canada. ASI cannot be controlled by any country. Any country that develops a superintelligence system will find its national security threatened.

The Chair: I'm sorry, but I'll have to stop you there.

[*Translation*]

Ms. Lapointe, you have the floor for six minutes.

Linda Lapointe (Rivière-des-Mille-Îles, Lib.): Thank you very much, Mr. Chair.

Welcome to the witnesses.

I must say that your statements have highlighted the risks associated with artificial intelligence. We usually hear about its positive aspects, but you are trying to point out the more problematic aspects.

I would like your answers to my questions to help us determine whether Canada is on the right track regarding oversight of artificial intelligence. We know that we are very advanced in this area, particularly in the Montreal region.

Here is my first question.

Canada has launched the Canadian Artificial Intelligence Safety Institute, or CAISI, whose mandate is to independently test and evaluate advanced artificial intelligence systems. In your opinion, how important is it for countries to create public and independent institutes such as the CAISI?

[*English*]

Gabriel Alfour: Being aware of the state of the art in artificial intelligence is quite important. However, I think we may already be way past that in many ways.

We have already gotten a warning from experts about extinction risks from AI. We have already gotten results from many AI safety or security institutes showing that some AIs are already able to persuade people, to manipulate people and to sometimes even break containment. From my point of view, we have already gotten worrying results from existing systems. We have already gotten warnings from experts about systems that are soon to come—in the next three to 10 years.

I think it's important, but now it's even more important that we move on to the next step, which, beyond just measuring, is to actually act.

[*Translation*]

Linda Lapointe: You have talked about the risks associated with artificial intelligence. When you talk about measuring, what exactly would you measure? I'd like to understand what you want to measure.

[*English*]

Gabriel Alfour: I think that's a great question.

I think we're already past many dangerous limits. For instance, we already have very persuasive systems. If we wanted to ensure that systems cannot manipulate people, we already have systems that are good at this. The same thing is true for hacking, for instance. If we wanted to ensure that current systems are not good at hacking, that is already lost. We already have systems that are good at hacking.

Now we're only measuring how much better they're getting, how superhuman they're getting and how much faster than people they're getting. We've already passed a few points that are quite dangerous. We're already in a tightening regime, edging closer to places from where we cannot really recover. This is the type of stuff we're talking about when we talk about measurements.

Another one that is relevant is how AI can autonomously develop itself. Right now we have companies that use AI more and more to develop AI. We have fewer and fewer humans in the loop. This is one of the other things we try to measure: how few humans are needed to develop AI. This is a measure of interest because it tells you when it could kick-start a runaway loop, which is basically a loop in the development of AI where it develops faster than we can even see it coming. These are the types of measurements we usually care about.

[*Translation*]

Linda Lapointe: My question is related to another question that my colleague asked you earlier. Do you believe that if we worked together with all the other countries, we could avoid the risks that you have been listing?

[*English*]

Gabriel Alfour: I personally believe so. These risks are concentrated in superintelligent systems. I think the hard part is monitoring and regulating the precursors to such superintelligent systems, but if we do so, there are many benefits we can get from AI. I will not say that it's easy, but I will say that it is very much tractable; it is doable. We can do it scientifically, and I think we should do so.

● (1125)

[*Translation*]

Linda Lapointe: Thank you.

Mr. Leahy, do you have anything to add about the risks of artificial intelligence and its regulation at the global level?

[*English*]

Connor Leahy: I would agree with everything my colleague said. Tractable but hard is a good way to think about this.

There are many benefits from this technology, as with any other, but an uncontrolled race, including between countries, is in no one's interest. There is no winning, ultimately, if superintelligence is built. It doesn't matter who does it; there will be no benefits. There are many benefits to a well-regulated, well-understood and well-controlled AI market. Doing so is hard, but the work has to be done.

[*Translation*]

Linda Lapointe: I must say that I find your comments very disturbing. I am confident that we can succeed if we have good intentions and everyone sits down at the table to try to establish regulations.

If we had global rules to regulate the development of artificial intelligence, as you say, there would be better outcomes for everyone. However, we need to act on a global scale. Do I fully understand what you are saying?

[*English*]

Gabriel Alfour: I would tend to say so. At least for the mitigation of ASI, it should be done worldwide. If anyone does it, we all suffer from it. We live in a very interconnected world. If we have a superintelligent system that can overthrow governments and can play geopolitics and war better than any human—if anyone builds it—we're in deep trouble.

[*Translation*]

Linda Lapointe: Thank you.

The Chair: Thank you, Ms. Lapointe.

Mr. Thériault, you have the floor for six minutes.

Luc Thériault (Montcalm, BQ): Thank you, Mr. Chair.

Gentlemen, your presentations are quite compelling.

I will start by asking you a slightly more technical question.

Some people would say that this is alarmist rhetoric, that artificial intelligence is a long way off, that we have not yet attained artificial superintelligence and that we will have time to look into it when the time comes.

However, two trends indicate the complete opposite of what these people are saying. Let me know if you agree. First, computing power is increasing exponentially. Second, the artificial intelligence models currently available appear to be advancing in intelligence at an exponential rate, thanks to available data and the size of the model.

Would you agree that we don't have as much time as we think?

[*English*]

Gabriel Alfour: I clearly think so.

Another way to think about it is that 15 years ago no one predicted where the capabilities of artificial intelligence would be now. This is the big thing that has happened. This is a big reason experts are warning about the extinction risks. Very few people, almost no one in AI, expected that we would have models as powerful as the GPT suite of models from several companies.

Put simply, when I was a teenager, no one expected that we'd have models that could talk to people. This was inconceivable. Things are accelerating faster and faster.

[*Translation*]

Luc Thériault: Governments are somewhat fascinated by means that improve efficiency, and they seem much more inclined to rush to implement applications to achieve greater efficiency. Time will tell whether there are indeed efficiency gains.

From what I understand from your remarks, artificial intelligence is a monster in the making, and it is evident that we cannot allow its development to continue without better regulation. Now, the question is how to regulate it.

In Canada, Bill C-27, which died on the Order Paper, proposed the creation of the Commissioner of Artificial Intelligence and Data position within Innovation, Science and Economic Development Canada.

The recent Carney government has a Minister of Artificial Intelligence and Digital Innovation. We would like to invite him here soon, but he does not want to meet with us. Last June, he stated that he would place greater emphasis on finding ways to exploit the economic benefits of this technology rather than on regulation.

What do you think of this approach?

● (1130)

[*English*]

Connor Leahy: I think there are, obviously, many reasons to focus on the positives rather than the negatives. It is, frankly, more profitable and more fun, to say it rather simply.

It is important to note that we are facing a polycrisis in many areas across the globe. I'm less familiar with Canada than I am, for example, with the U.K., where I currently live, or my native Germany or America, but we are facing many challenges. Often, there is a seduction towards thinking of technological solutions. Often, technology is a very powerful solution. It is a very powerful way to help address these problems. I believe AI can be helpful for many of these problems, but it is also important to be skeptical.

There was a world where nuclear power, when it was first being developed, was thought as a solution to everything. Some wanted to make nuclear-powered aircraft, for example, that would spew radiation as they flew. They wanted to use nuclear weapons for mining applications. I think the Russians actually tried that one, and many such things.

This is not to say that nuclear power is not an incredibly useful and powerful technology. I think nuclear power plants are some of the most effective ways to make energy, but the reason they are so safe, good and useful is good regulation. This took decades of hard

work. It took the invention, by many experts, of whole new forms of safety engineering to reap the benefits.

I think we're seeing a similar thing here. If we try to reap the benefits without the correct amount of safety engineering that is necessary, we will see the same thing we've seen with social media repeat. We will not see the AI equivalent of safe, economically productive nuclear reactors.

[*Translation*]

The Chair: You only have 15 seconds left, Mr. Thériault.

Luc Thériault: Then I can't ask more questions.

The Chair: Thank you.

[*English*]

We have Mr. Barrett for five minutes.

Go ahead.

Michael Barrett (Leeds—Grenville—Thousand Islands—Rideau Lakes, CPC): Based on your experience with malicious AI uses such as deepfakes, cyber-attacks and autonomous exploitation tools, what safeguards or regulatory measures should we be examining to protect Canadians? What examples can you point to of where those are successfully being implemented either at a corporate level or at a national or subnational level?

I'll give both gentlemen a crack at the question if they'd like.

Gabriel Alfour: I will start.

There are basically two separate categories of risks. There are the risks that come from superintelligence and more generally from systems that cannot be controlled. For these risks, I recommend basically regulating the development of such systems.

The problem with such systems is that once they're developed, we cannot control them; we cannot put the genie back in the bottle. This is one type of regulation that is quite important. That is why we put a strong emphasis on international agreements and this type of regulation.

The other one is for the, let's say, more prosaic risks of current systems. Future systems are not yet superintelligent, which means that if there is a problem, we can put the genie back in the bottle. Here, it's more at the application level. When I say the application level, I mean more the AI companies—you want to regulate the bottlenecks. Here it will be beneficial to put stringent regulations on the dangerous aspects of AI and put strong liability regimes in place so that if the systems they build are used for nefarious purposes, they will also be liable for it.

I would separate those two, one being the development part of regulation for superintelligent AI and the other being the applications for current systems and near-future systems.

● (1135)

Connor Leahy: To pick up on what my colleague just spoke about, it's important to note that a play is very often attempted by these actors where they attempt to blame the users for misuse of their tools at every possible opportunity.

In general liability law, the best practice is to put liability on the part of the supply chain that is best suited to addressing the harm. Obviously, the part best suited to addressing the harm is these massive corporations with the best technical talent in the world, massive platform leverage and so on, rather than the user. I would push against user liability as the way to address these risks and push much more for developer or employer liability.

Michael Barrett: I'd like to follow up on the point with respect to the regulation of development.

If we gatekeep the advancement or development here, and we have treaties with many other countries but, for example, Russia and/or China don't participate in the treaty, wouldn't we find ourselves in a situation where our adversaries are proceeding with development in an AI arms race while we just watch and hope that things don't get out of hand? Based on the context you provided in your opening statements, we know that it almost certainly will, but we may not have developed tools that will allow us to defend against it.

Would I be correct in saying we can also use models to defend us against rogue states and the models they would develop and deploy to our detriment?

Gabriel Alfour: There are basically two regimes. There's the superintelligent regime and there's what comes before that.

In the superintelligent regime, everyone loses. If Russia builds superintelligent AI systems, everyone loses; they cannot control it. It's the same thing for China, it's the same thing for the U.S. and it's the same thing for any country. This is in the superintelligent regime. It's only the superintelligent AI systems that basically have any agency left.

Then there's the pre-superintelligent regime, where there is an actual race, because until there's superintelligence, when you can still steer your systems, you stand to gain a lot of benefits from developing stronger systems. This is why international agreements are important, and indeed, if we fail to build and reach international agreements, things will go badly.

The same is true for biotechnology. The same is true for any type of dangerous weapons, like nuclear weapons. This is why we believe that international agreements are critical, or else you get into the superintelligent regime and things go badly.

The Chair: Thank you for your response.

Mr. Barrett, thank you for your questions.

[*Translation*]

Mr. Sari, you have the floor for five minutes.

Abdelhaq Sari (Bourassa, Lib.): Thank you very much, Mr. Chair.

Thanks to the two witnesses with us today.

You have shared some very important information about artificial intelligence. I would not necessarily say that your presentations are alarming, but they are very factual. I can only agree with you about the risk. I can only agree with you on the extent of the impact that superintelligence could have on everyone's daily life.

However, my opinion differs from yours on one point, and I would like to discuss it. You used the verb "halt". Are you saying that we need to halt the development of these technologies or halt their use?

As you clearly explained, persuasion technologies or systems already exist. I think Canadians are already using such technologies. Do you want us to halt their use, given that most of these systems are developed outside Canada, or do you want us to halt their development on Canadian soil?

You drew a very interesting parallel with the nuclear arms race. However, in the nuclear arms race, geographical boundaries are important, which is not the case for technologies developed using training systems and artificial neural networks, which can be used anywhere in the world.

I would like to hear your opinion on this.

● (1140)

[*English*]

Connor Leahy: These are some very excellent questions, and this is, in fact, a very difficult thing.

The non-locality of digital technologies presents novel, unprecedented risks of proliferation and difficulty of control. With nuclear weapons, for example, we have the luck, in a sense, that uranium ore is quite bulky and centrifuges are quite hard to build and quite visible from space, if you do them right. We have some of these benefits when it comes to AI systems, such as data centres. Others, such as the open-sourcing of many such systems, present novel difficulties.

This is why we've put such emphasis on the regulation of the development. If a superintelligent system was made, it would probably be a computer file that can probably never be deleted. It would be something that would spread. We might not even know what we're dealing with when it is first developed. It is quite likely that we won't even know that the first superintelligent system is superintelligent until it's too late. It is already the case right now that many of our AI systems have capabilities that we didn't know at the time of development. We only discover much later that our systems are capable of things we didn't know.

This is a novel regime. This is not something we have dealt with very well historically. Even historically, things such as export controls on software have not been very successful or have been very tricky to enforce.

It's very important to say that we do not think all AI should be halted or that all AI applications should be halted or not used by users. I'm sure my colleague would agree with me that we very much enjoy many of the AI applications on the market today. What we want is to gain the benefits of the kind of AI we have right now and continue forward into more powerful applications.

[*Translation*]

Abdelhaq Sari: Time is running out and because I'd like to ask you a question about Canada's strategy, I will wrap up this question by saying that I also hope there will be an international agreement. You alluded to this, but I'm not very optimistic about that because I see the race in the field of quantum servers and facilities. I think I am a little less optimistic than I should be.

We received more than 11,000 comments during the countrywide consultation on Canada's strategy. I believe that we also need to educate Canadians. How can we institutionalize citizen participation so that it becomes a permanent pillar in the development of artificial intelligence-related safety policies?

[*English*]

The Chair: Answer in 30 seconds or less, please.

Gabriel Alfour: We believe that awareness and education are extremely important for controlling AI. It's about half of what we do. First is education.

The second thing is putting people into the decisions of deploying systems. A lot of people are against the development of superintelligent AI systems, and whatever we can do to put them in the loop to have a say, we believe is good.

The Chair: Thank you so much.

[*Translation*]

Thank you, Mr. Sari.

Mr. Thériault, you have the floor for five minutes.

Luc Thériault: Thank you very much, Mr. Chair.

The chief executive officer of the machine intelligence research institute told us when he testified last week that a global shutdown is not currently politically feasible and that's why his organization is focusing on safeguarding the ability to shut down artificial intelligence by creating a kind of off-switch. He proposed putting in place the technical, legal and institutional infrastructure necessary to restrict the dangerous development and deployment of artificial intelligence on an international scale. This is what he calls an off-switch. That would lead to a coordinated, international shutdown of cutting-edge artificial intelligence activities at some point in the future.

What do you think of this proposal?

● (1145)

[*English*]

Connor Leahy: We generally believe that there is a lot of value in what are often called "stop button" proposals, mostly not from a technical perspective, but as a social-political factor. As an example, I have in the past asked someone who worked at a large tech company whether they could shut down all of their servers if they wanted to. He said no. He didn't know where all of them were. There was no single person in the entire company who knew where all the servers were and what software was running on them.

As for creating legislation, I am not familiar, unfortunately, with the specific proposal you mention. There is a lot of value in this, but it sounds to me that the proposal pushes against development.

We cannot rely on waiting until we see a superintelligent system, because by the time we see one, it is already too late. It's quite likely that when the first superintelligent system gets built, we will not even recognize it as being that until quite a bit later, and that will be far too late.

It's very important—which is why there's a repeated emphasis on precursors—to make sure that such systems never get developed in the first place. To do this, we need to already be in the loop before such systems are built.

[*Translation*]

Luc Thériault: You mentioned precursors several times. How can precursors be restricted?

[*English*]

Gabriel Alfour: I'll take this one.

There are many different types of precursors. Some are hardware, like data centres or the way to build GPUs into the supply chain and things like this. Some are software, like the types of AI systems that have been built and the type of scaffolding you have on top of them.

A deep thing that is quite important to understand is that this is a moving target. As time passes, it gets easier to build superintelligent systems, and more things get into the category of precursors. This is also why we believe there is an urgency and that we should tackle this as soon as possible.

Right now, we can get away with, for instance, preventing research programs that are aimed at building superintelligent systems. We should also focus on limiting the open-sourcing of models, because once they're there, you cannot take them back. It is the same for data centres. For every data centre, there should be stop buttons and kill switches. There should be clear regimes for what can be done with them and so on. These are the types of regulations that we should have on precursors.

Fundamentally, it is a moving target. As the technology changes and as the way the technology is built changes, the target itself changes. The precursors of 15 years ago were very different than they are now, and it would have been much simpler to tackle the problem 15 years ago.

[*Translation*]

Luc Thériault: Thank you.

How much time do I have left, Mr. Chair?

The Chair: You have 35 seconds left, or perhaps a little more.

Luc Thériault: I will ask my question, and if we run out of time, you can answer it when I have another turn.

One of the ethical issues that worries me is the energy-intensive nature of data centres. For example, there are two coal-fired power plants in Mumbai that are extremely polluting. They were scheduled to be shut down, but they will continue to operate to meet the enormous electricity needs of Amazon's data centres, which are being built all over the world to compete with other large companies such as Google. I am concerned that artificial intelligence is being developed for the benefit of wealthy countries, but at the expense of the environment and the health of people living in developing countries.

The Chair: Please give a brief answer.

Luc Thériault: This could be the beginning of an answer.

[*English*]

Gabriel Alfour: This is a more general aspect of the fact that AI is developed without much concern for people and that people are not put in the loop. That's mostly how we see it.

[*Translation*]

The Chair: Thank you.

[*English*]

Mr. Mantle, you have five minutes. Go ahead.

Jacob Mantle (York—Durham, CPC): Thank you, Mr. Chair.

Thank you to our witnesses for appearing and providing valuable testimony.

Picking up on the comment from my Bloc colleague about the protection of people, I want to focus on priorities for a moment and the government's current priorities with respect to artificial intelligence.

You may be aware that the government is currently reviewing its AI strategy, but the existing AI strategy, the pan-Canadian artificial intelligence strategy, lists three priorities. The first priority is commercialization, so making money off of AI; the second is standards or protections in that area; and the third is talent and research.

Do you think those priorities are in the proper order?

● (1150)

Connor Leahy: I think it wouldn't be a surprise that we would disagree that these are the most important priorities. Both of us are technologists by background. We love technology. We got into this to do technology to make the world a better place. Technology is dual-use. It is power. It is very important to use it correctly when we're dealing with unprecedented technology that has this kind of power.

To get a good outcome, the most important thing is to get this right and to not repeat the mistakes of social media. Don't let technology exist for technology's sake. Let technology exist to benefit people. This is not what technology will do by default. We have to make it do that. We think this should likely be the top priority.

In the sense of the most acute risk, personally, we believe superintelligence to be the most pressing.

Gabriel Alfour: I would tend to agree.

Perhaps I can add something. I think it was in July of this year that YouTube Shorts reached 200 billion views per day. There was an entire article from YouTube about this, which was extremely happy about the fact that it got 200 billion views per day.

We think a lot of people building AI systems are in this paradigm. We know there are many technical employees at AI companies. It's extremely fun to watch loss go down and results on benchmarks go higher. It's good in itself to get more technology to just get more. It's fun. It's great to see. You don't need to think much.

I will echo what Connor said. I believe the biggest priority when developing technology should be to ensure that it benefits people and benefits humanity.

Jacob Mantle: If I'm understanding your testimony, you're suggesting that the priorities should be reversed and that standards and protection should come before commercialization in the government's strategy. Is that correct?

Gabriel Alfour: I think benefiting humanity can be done through many ways. In the context of superintelligence, we believe it's through regulation and protection, but for other things, it's through using the technology in good ways, in ways that benefit people.

Personally, I'm quite hopeful about AI in the context of education. I think different people have different priorities, but if one tries to use AI for good, I believe a lot can be done with it. It's a very powerful technology.

Jacob Mantle: That's great.

Canada has recently created, in this government, a new ministerial position, the Minister of AI. In one of his first speeches after becoming Minister of AI, Mr. Solomon said that Canada would move away from "over-indexing on warnings and regulation" to make sure the economy benefits from AI. Could I get your reaction to that approach?

Connor Leahy: My general opinion is that, yes, you can get a lot of economic benefit by neglecting human flourishing. This is something we've seen historically many times. There are many ways to pump a stock market at the expense of people. For example, deregulating Ponzi schemes is very profitable in the short term. Eventually, the bill comes.

I think we're seeing a similar thing here. Is building long-term, responsible stewardship and building a good society that uses technology effectively...? Again, I think benefiting mankind also means using AI, but using it correctly and for the benefit of mankind. This is harder and in the short term less profitable.

The Chair: Thank you, Mr. Mantle.

Mr. Leahy, I'm going to get you to move your microphone down a bit. It's coming in a little hollow right now. We want to make sure the interpreters understand what you're saying.

Ms. Church, you have five minutes. Go ahead.

Leslie Church (Toronto—St. Paul's, Lib.): Thanks, Mr. Chair.

Welcome to both of our witnesses.

I'll go first to Mr. Alfour. What experience do you have in dealing with the Canadian artificial intelligence safety institute, or CAISI, which was launched in November 2024? Do you have any experience in dealing with it directly?

● (1155)

Gabriel Alfour: I do not have much experience dealing with it directly, aside from interacting with a few people who I think were there when I talked to them.

Leslie Church: Would either of you have comments on how you think its mandate and work are shaping up?

I would note for the committee that CAISI was developed by the Government of Canada in part to examine the risks posed by advanced AI systems to help develop tools and guidelines to manage those risks. It also works collaboratively internationally to try to develop protocols for AI safety.

It's only a year old at this stage, but I was wondering if either of you had any guidance for us as a committee on the reach of its mandates or how you think its work could be enhanced going forward.

Gabriel Alfour: As far as I understand it, right now, CAISI is observatory. It does not regulate and it does not constrain. I can't say if it's the role of CAISI specifically to regulate and constrain. This is a political call that is outside of my purview, but I think someone should have this authority.

We had warnings two years ago from experts about extinction risks. We already know that several leading AI corporations are racing specifically for superintelligence. Now is the time for action. Whether this should be done through CAISI, through the Minister of AI or through another entity is beyond my pay grade, but I believe it should be done.

Leslie Church: It's one of the tools in the tool box to help develop the framework that Parliament would need to legislate if we were going down that route. It's interesting, because I think we're attuned to the risk, but as you've acknowledged, finding the means to regulate and understand AI in its various forms is going to pose a challenge.

Mr. Leahy, do you have any wisdom for us on how we could grow the mandate or work of CAISI to support us in developing framework and security protocols to deal with this?

Connor Leahy: I have not had personal experience with CAISI. I have not talked to them. I have talked to many people around Yoshua Bengio at Mila and his group. He is one of the godfathers of AI. These are the people I have the most experience with.

Generally, it is very important to have some amount of technical capability that can give non-partisan advice to governments. This is a very important function. It's also very important to understand that a lot of these institutions, in a sense, through no fault of their own, are often corrupted by business interests.

Many of the people who work at these companies or at these institutions have very lucrative offers from these kinds of companies and often have their identity tied to technology being good. On the other hand, there are many people whose identity is tied to technology not being good. We don't want either of these things. What we want is a balanced understanding of how we can mitigate the risks that are truly unacceptable, while also benefiting from the ones that we can't mitigate. This is a hard tightrope to balance, but it's important to make this the clear mandate.

Leslie Church: Thank you. That's very helpful. I think it's helpful that the Canadian AI institutes, including Mila, are involved with CAISI as well to provide some of that balanced input.

Let me take you in a different direction. What type of counter-protocol would you suggest? You've mentioned the geopolitical risks we face from other foreign actors. I'm curious to know if you have any guidance for us to think about, either on the cybersecurity reaction or on the protocols we should put in place to protect our critical infrastructure and systems. Are there any other counterprotocols that you might suggest we look at as a government?

The Chair: I need a rather quick response to the question, please.

Gabriel Alfour: I tend to think that countries should now seriously think about national firewalls to prevent outside actions. There has been a lot of taboo about them in the past. If a country is to ensure its cybersecurity and security in its cybersphere, it should strongly consider this in general.

That's my personal opinion, not that of ControlAI. To your question, I think that's the most direct one.

● (1200)

The Chair: Thank you, Ms. Church.

Mr. Leahy and Mr. Alfour, I want to thank you for appearing before the committee this morning.

We are going to suspend for a few minutes while we get into our second hour.

● (1200)
_____(Pause)_____
● (1205)

The Chair: I'm going to call the meeting to order.

I would like to welcome our witness in the second hour. From INQ Law, we have Carole Piovesan, who is the managing partner.

Unfortunately, we were to have another witness, but they could not be here. It's the Carole show for the second hour.

Carole, you have up to five minutes to address the committee. Go ahead, please.

Carole Piovesan (Managing Partner, INQ Law): Good afternoon. Thank you, Mr. Chair and honourable members of the committee.

My name is Carole Piovesan. I am a managing partner at INQ Law, where I advise clients on privacy, cybersecurity, data governance and AI risk management.

I've had the privilege of contributing to AI policy discussions nationally and internationally, including through the OECD.AI Policy Observatory. I have previously appeared before this committee, as well as the INDU committee. I am an adjunct professor at the University of Toronto's faculty of law, where I teach AI regulation. As well, I co-authored the book *Leading Legal Disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law*.

The opinions I present today are my own personal opinions and do not reflect those of my law firm.

To understand how we should govern AI, we should go back to first principles and ask what AI is trying to achieve. In 1950, Alan Turing posed what he called the imitation game: a test to determine whether machines could think. He believed that one day machines would be able to play games, remember, observe results of their own behaviours, learn from rewards and punishments, and even deliberately introduce mistakes into their working.

Today, some of the leading AI researchers around the world are divided on where the trajectory of AI is taking us. Award-winning Canadian researchers such as Yoshua Bengio and Geoffrey Hinton, both pioneers of deep learning, warn that we may soon have computers that exceed human intelligence, with profound implications for safety and control—indeed, the existential threat we all hear about. Others, such as Yann LeCun, another pioneer, advance an argument that is more aligned with artificial machine intelligence, which is best understood as augmenting human intelligence rather than replacing it.

The purpose for pursuing AI and the achievements of those pursuits matters in how we think about governance. If a tool is to be used to extend human capabilities, we govern its use. If AI is an autonomous system capable of independent reasoning, we regulate its development and deployment with a different level of vigilance. Canada's approach must account for both.

Around the world, we are seeing at least three distinct models of governance being presented for AI.

Under the Trump administration, we are seeing a deregulatory approach in the United States with an emphasis on competitiveness over comprehensive safeguards. The federal approach relies on existing sectoral laws applied through agencies such as the FTC, while actively resisting state-level experimentation with stricter AI rules.

The United Kingdom and Singapore take a different approach. There, we are finding a much more tailored sectoral approach to AI

regulation. The U.K., in particular, has a principles-based approach asking existing sector-specific regulators to interpret and apply cross-cutting principles such as safety, transparency, fairness and contestability within their domains. The U.K. considers that this approach offers critical adaptability that keeps pace with rapid technological change, although there are certain developments that suggest binding measures for the most powerful AI models may be forthcoming.

Singapore has certainly adopted a much more soft law, voluntary framework. There is no specific AI regulation. However, Singapore's approach through consensus building among government, industry and citizens, and through instruments such as the model AI governance framework and the AI Verify Foundation testing tool kit, has proven somewhat successful in building a sense of trust and a common approach to AI development. With Singapore's investment in national AI literacy and its consultative and iterative approach to governance, it's a model from which Canada can draw inspiration.

Then we see the third model, which is far more prescriptive. That model is found in the EU AI Act, which I know this committee has already heard about. That act is much more horizontal and is focused on the prescriptive life cycle of AI development and deployment across the supply chain.

● (1210)

Canada's approach should be tailored to our context. Regulating frontier AI systems is not the same as regulating Copilot in the use of a law firm or a chatbot on a service line. The U.K.'s context-specific approach recognizes this. Canada is more like the U.K. and Singapore than the United States or Europe. We value proportionate regulation that protects rights while enabling innovation.

I'll close with my three-point call to action.

The first is to continue building a regulatory guidance approach for safe AI. Our AI safety institute must be operating at full force, demonstrating that Canada takes the safety of these systems seriously. We must continue to target iterative standards guidance and a directives-based approach to artificial intelligence, with an emphasis on real-world testing for high-risk AI contexts. Lab benchmarks and off-line evaluations only show how models perform on static tests, not how they actually interact in real-world use.

Second, and very importantly, we need to improve the diversity in representation and perspectives in policy and throughout the development, evaluation and deployment process. Individual perspectives matter, and they are highly unrepresented throughout the AI ecosystem.

Third, we must conduct an environmental scan to better understand, on a sectoral basis, where our laws may have gaps to account for AI or where AI is already accounted for, so we have the coverage we need for the everyday use of AI in business. Targeting soft and hard law at home in a tailored manner and enabling Canada to play to its trusted global position to ensure robust and harmonized standards, certifications and guidance for responsible AI should be our path forward.

Thank you. I welcome the committee's questions.

The Chair: Thank you, Ms. Piovesan.

I know there's a fourth, fifth and sixth component to your recommendations—I've seen your opening remarks. Perhaps committee members can guide you with that in their lines of questioning.

Mr. Barrett, you have six minutes. Go ahead, please.

Michael Barrett: I'd like to pick up on one of the points you mentioned in your opening. It was about the regulation of frontier AI systems. In doing that, how would Canada defend against rogue nations developing and deploying AI weapons? What international coordination would be required to make safeguards in the development of frontier AI systems effective?

Carole Piovesan: Every nation is drawn to protect its own systems as much as possible. Canada has been working internationally since 2019 through the Global Partnership on Artificial Intelligence, through the OECD, through the G8 and G7, and through other international mechanisms, including, now, the international AI safety institutes. We've been working quite diligently to establish more of a commonality in how we are approaching AI and the regulation or protection of frontier AI around the world, where rogue nations may be developing certain AI systems that are offensive to our own principles.

It's really important to understand that Canada is already embedded in these international committees, and we are playing a key role in establishing what the norms ought to look like and the mechanisms to enforce those norms. We won't be able to do it alone at all. We have to find out who our friends are and where we have commonality in approach and values, and, through that, establish the mechanisms we need in order to defend those values and that approach.

● (1215)

Michael Barrett: I wanted to ask you about the long-term trajectory of superintelligent AI and the worst-case scenarios that will follow if adequate safeguards aren't put in place to stop it from getting away from us as a species.

First, could you tell me in what terms we're talking? Could you qualify, in your opinion, whether we are talking about long term, 10 years or five years? What do you think?

Carole Piovesan: I'm obviously not a technologist, but I've sat around circles to hear what some of the technologists are saying.

What I am hearing, and I have no reason to disbelieve them, is that we are decades away or less. We're no longer talking about....

I remember sitting in a class with Dr. Hinton years ago when he was suggesting that artificial general intelligence was hundreds of years away. Then, in 2023, he shifted that approach to say that we're much closer than we ever thought.

The pace of technological development is happening at an exponential level. By all accounts from what I'm hearing, with no reason to disbelieve, we are much closer to superintelligent computers than we probably thought we would be just a few years ago.

Michael Barrett: It's easy for one to fall down the dystopian mind hole about what that looks like.

I have just over two minutes left. What do you hear? What is discussed? What are the outcomes we're looking to avoid? How can we as a Parliament play a role in preventing that from happening?

Carole Piovesan: We're looking to avoid superintelligent computers that are superior to us in a way that is harmful to us. Dr. Hinton spoke about embedding maternal instincts into superintelligent computers so they would be more empathetic and protective.

Parliament can play a role in ensuring there is regular tracking of the development of, and inputs into the development of, frontier models and how they are implemented within Canadian society, with a view to shaping what the values that would be encoded in these systems ought to look like. Parliament can play a role in ensuring that we identify the values we ultimately want these systems to be embedded with, and then in providing a process through which those values can be embedded in the systems so we can see the outputs.

Michael Barrett: Thanks very much.

The Chair: Thank you, Mr. Barrett.

[*Translation*]

Mr. Sari, you have the floor for six minutes.

[*English*]

Just make sure, Carole, that your interpretation is on.

[*Translation*]

Abdelhaq Sari: Thank you very much, Mr. Chair.

Ms. Piovesan, thank you for the information you have shared with us and for the insight you have provided.

Nevertheless, I will begin with the same introduction as I did for the two witnesses who preceded you, by seeking to understand what we can control.

In your opinion, will it be possible to control the development of this superintelligence? Shouldn't we instead be working harder to control the use of artificial intelligence itself, given that most of the systems designed to persuade Canadians are not developed in Canada?

● (1220)

[*English*]

Carole Piovesan: That's a fair point; the systems aren't necessarily developed here.

First of all, through international coordination, we have a mechanism to inform what those values and standards ought to look like. We have precedent for this in a number of different respects, some of which you heard about from the earlier witnesses. We have some role to play in being clear about what embedding those good values looks like and what it means to develop intelligent computers that are aligned with the role we believe they should play in society.

I will give you the example of the G7 back in 2018, when our government and all the governments of the G7 were talking very much about the values in the context of AI and how we were going to advance standards, policies and practices that would protect those values. From there, you saw a movement into the global partnership on AI, and you saw a coordination through the United Nations on AI with a view to creating more of a harmonization in approach. Cut to the most recent G7, where the emphasis was primarily on adoption because we're now at a stage where we can see the real-world uses of AI and are much more prone to and excited about, in a lot of ways, the adoption of these technologies, which is a good thing.

In this current context, as we are starting to become much more familiar with the technology and understand its opportunities and use, we have to be mindful of the risks, but we cannot lose sight of the opportunities. We have a role to play in ensuring safe use where there is actual risk. What I don't want to do is establish a system that applies the same kinds of controls for all uses. We have to be targeted.

[*Translation*]

Abdelhaq Sari: Exactly.

You said that we need to be very aware of the risks, and I agree with you on that. I think our government is already aware of these risks, as are several G7 member governments, as you said.

You also mentioned the American approach, which is much more competitive, and the U.K. approach, which is much more based on ethics and accountability.

What approach can you really suggest to the Canadian government?

As I have another question for you, I would ask you to be a little more concise, please.

[*English*]

Carole Piovesan: I'm more inspired by the U.K.'s and Singapore's approaches than I am by those of the EU or the U.S.

[*Translation*]

Abdelhaq Sari: Don't you see that this approach may not have the desired results if other governments are not aligned in some way? This artificial superintelligence will continue to develop.

[*English*]

Carole Piovesan: I think that's exactly right. It doesn't necessarily stop the development of AGI, but it does embed a more mature approach in how we regulate it today.

[*Translation*]

Abdelhaq Sari: My last question concerns Bill C-27, which you were in favour of.

What lessons have you learned from this bill, and what would be the best way forward? We have already done some work on it, and we can't just throw it all away. In your opinion, what lessons have been learned with regard to this bill?

[*English*]

Carole Piovesan: Part 3 of Bill C-27 was the artificial intelligence and data act, and the best lesson to learn from it is that the overarching accountability framework that would have been put in place required a turning of the mind to the context of the use of AI, determining its potential impact level and then establishing a diligence process, soup to nuts, in response to that level of risk.

[*Translation*]

Abdelhaq Sari: I haven't read your book yet, but I think it would be a really good read for the holidays.

Could you briefly tell me what your position is on digital sovereignty in Canada?

[*English*]

Carole Piovesan: It's a complex question to answer in a short period.

I understand and appreciate the approach of digital sovereignty. I think there's a lot that's extremely important about digital sovereignty. I also recognize it's a long-term investment.

● (1225)

[*Translation*]

Abdelhaq Sari: Thank you very much.

Carole Piovesan: Thank you.

The Chair: Thank you, Mr. Sari.

Mr. Thériault, you have the floor for six minutes.

Luc Thériault: Thank you very much, Mr. Chair.

Welcome, Ms. Piovesan.

In June 2025, Professor Bengio, who is also the founder of Mila, the Quebec Artificial Intelligence Institute, and scientific adviser to the institute, launched a new non-profit research organization on artificial intelligence security called LawZero, to prioritize security over commercial imperatives.

In your opinion, what are the greatest risks that artificial intelligence poses to security?

You said earlier that, from the outset, we must first ask ourselves what the objectives are for using this technology. People are very much driven by the lure of profit.

[*English*]

Carole Piovesan: We have to acknowledge the cybersecurity risks of artificial intelligence and whether we're ready as a country to defend against those risks. We also absolutely have to recognize the human rights- and social development-related concerns about artificial intelligence and walk in with our eyes wide open as businesses in our country start to adopt AI with much more interest.

[*Translation*]

Luc Thériault: With regard to human rights, the Montreal Declaration for Responsible Development of Artificial Intelligence states the following about AI systems:

> 1. AIS must be designed and trained so as not to create, reinforce, or reproduce discrimination based on—among other things—social, sexual, ethnic, cultural, or religious differences.

> 2. The development of AIS must help eliminate relationships of domination between groups and people based on differences of power, wealth, or knowledge.

Does the government take sufficient account of artificial intelligence biases in policy implementation? For example, did it do so in its Bill C-27? How should it further integrate this concern?

[*English*]

Carole Piovesan: We have a job to do in establishing discussions among different legal regimes, such as the human rights regime and its place within artificial intelligence. This means that, in the application of human rights law and our charter, for instance, we should understand the connection and role AI has in each of those sections. There are absolutely concerns around the use of AI producing fair outputs. What we need to start to understand is, what does fair mean? What is the standard? How do we judge it? How can we demonstrate that we are living up to these standards? That is critical.

In addition, I want to highlight one of the points I made earlier: The diversity of perspectives around the table matters. You need to hear from different people with different perspectives. It really matters and will shape the way we approach AI policy in law.

[*Translation*]

Luc Thériault: Let us go back to Mr. Bengio's approach. Can any guardrails be put in place to prevent some of the most common risks posed by artificial intelligence? You made a series of recommendations earlier, but could you elaborate on these recommendations or provide additional ones?

[*English*]

Carole Piovesan: There are a number of recommendations to ensure that AI is used in a safer manner. Many of the recommendations I have in my written submissions I tried to consolidate, just for expediency.

We need to augment the standards we have in place. There are standards through the international standards organization. The Standards Council of Canada is working on a body of research to support Canadian standards in AI. The application of those standards matters.

When we look at ISO standards, we are looking at a governance standard for how you operationalize responsible AI within the use of a system. Actually, it's the use of AI in a particular context, which means it is not technical; it is governance-oriented.

By increasing our understanding through a national literacy program, through more sectoral guidance, through more industry collaboration and through greater perspectives being brought to the table, we can start to have very actionable plans around how we will operationalize these standards and what "good" looks like in order to achieve in each of these standards.

● (1230)

[*Translation*]

Luc Thériault: Can—

The Chair: Sorry to cut you off, Mr. Thériault, but your six minutes are up.

Luc Thériault: Oh, all right.

The Chair: Time flies.

Mr. Hardy, you have the floor for five minutes.

Gabriel Hardy: Thank you, Ms. Piovesan, for appearing today.

Many companies have made it clear that their aim is to use AI to take over a significant portion of human-driven tasks. Studies indicate a 13% decline in hiring among youth ages 22 to 25 for roles where artificial intelligence can easily replace human labour.

Is the use of artificial intelligence for labour essentially equivalent to hiring postdoctoral level talent to work 24 hours a day, for less than minimum wages? It seems this is already part of the conversation.

Ethics and responsible innovation are central to everything. Do you think we should legislate this particular issue to ensure companies don't start using artificial intelligence instead of hiring people?

[*English*]

Carole Piovesan: I don't think you can legislate so that businesses aren't hiring AI instead of real people. Let me offer a different perspective.

We work with clients all the time to operationalize their AI governance programs. When we're going through use-case identification, I always ask them three questions. Number one, what is the work they don't want to do because it's boring and mundane? Number two, how much time do they spend on each of those tasks? Number three, what would they be doing if they weren't doing those tasks? A hundred per cent of the time I am told they would be more proactive, they would be able to serve their mandate better and they would be able to contribute more value to their organization.

Here are my points.

Number one, AI as a tool is going to be used within our businesses, and we can't and shouldn't stop that use.

Number two, we will have to reorient the job market. I have kids. I'm distinctly aware of where they're headed and where there may be vulnerabilities in some of their job choices. I understand that there will be shifts in the way we structure labour within Canada, but we have to make that adjustment. Instead of resisting it, we have to support re-skilling and upskilling. This is something that we as a country have been talking about for many years.

The last point is about literacy. We should allow people to understand how this tool is to be used and enable them to better identify and shape their own career paths, recognizing the profound transformational impact of artificial intelligence.

[*Translation*]

Gabriel Hardy: I do understand what you are trying to say: If someone isn't stuck doing tedious tasks, they could be engaged in more creative work. However, it's important to think about the fact that if everyone were doing the same thing and had the same abilities, then artificial intelligence could do most of the work in the country. Don't you think businesses will try to take advantage of that? We would then have to ensure those who lose their jobs don't generally become a burden on society, since they will need financial support, for example.

Don't you think we can regulate companies that replace human employees with artificial intelligence?

[*English*]

Carole Piovesan: I'll take your question in two parts.

First, is there something we should be doing to support this transitional period, as people might be looking for new opportunities because of the displacement caused by AI? Second, is there something we should be doing to prevent companies from using AI instead of people?

On the second point, I don't think we should be inhibiting companies from adapting to the use of AI. I don't think that's the approach. It's certainly outside the scope of my practice, so that's very much a personal opinion, but I don't think that's the right approach.

On the first part, I agree with you. We should be continuously monitoring and investigating how we can better support our workers as we go through a transformation on a step-by-step basis.

● (1235)

[*Translation*]

Gabriel Hardy: I'd like to shift gears with a different question.

In the private sector, setting an example often shapes the direction of the market. What can we do here in Canada to lead by example when it comes to artificial intelligence and encourage countries to emulate us, to ensure we don't lag behind others?

Can we adopt the best approach ethically? Can we make sure that the development of artificial intelligence is confined to use in very specific areas?

[*English*]

The Chair: Give a very quick response, please.

Carole Piovesan: I appreciate that question.

We can lead in two ways. We can adopt the technology, and we can do so with the responsible made-in-Canada AI brand we have been developing for a very long time. We should sell that to the world.

The Chair: Thank you, Monsieur Hardy.

Mr. Saini, you have five minutes. Go ahead, sir.

Gurbux Saini (Fleetwood—Port Kells, Lib.): Thank you, Carole, for coming.

I was very alarmed by the two witnesses who came before you given the extent that humanity may be at risk. What can we do? The United States, China, India and G7 countries can develop all these things that may be a danger to humanity, but a lot of countries in the world don't have the knowledge or facilities. What can we do to help them? How can we protect humanity from the uncontrollable use of this weapon?

Carole Piovesan: I think we have been here before. I didn't hear all the testimony of the two prior witnesses, but I believe they talked about this in the context of nuclear co-operation. I believe that was the case. We are going to rely on those processes again. We will have to find our friends, and we will establish the mechanisms with our friends to establish—

The Chair: Excuse me, Carole. I'm sorry.

Something became hollow. I don't know whether the connection came out of the microphone. Maybe you can plug it back in, if you don't mind, because the sound in the room got really hollow quickly.

I've stopped your time, Mr. Saini, just so you know. I'm going to give Carole an opportunity to respond to the question in its entirety.

Can you just give me a test, please?

Carole Piovesan: Sure. Is this better?

The Chair: It's much better. Thank you.

I'm going to give you an opportunity to restart your response. Once you're done, I'll start the clock again.

Carole Piovesan: That's wonderful.

As I was saying, we have been here before. I think we will have to find out who our friends are and establish the mechanisms to embed the right types of values, certification mechanisms, evaluation mechanisms and deployment safeguards to do what we can to prevent rogue countries from investing in and ultimately succeeding with superintelligent computers that are harmful to humanity. I think that's going to be our best path forward.

Gurbux Saini: My concern is with atomic energy. It was used, and it did a lot of harm before we realized it was dangerous.

Do we have any organizations, like the United Nations, that can regulate these things and tell rogue countries that it is enough, that we don't want to carry on with this?

Carole Piovesan: I don't think we have the right international governance organization set up to protect AI yet. We need to invest in that organization or augment an existing organization with a specific view to supporting the safe deployment of AI. We have seen some of that through the international network of AI safety institutes—the organization of safety institutes that consolidates and organizes the various national safety institutes—but I don't believe that is their explicit mandate.

● (1240)

Gurbux Saini: In general, how are Canadian AI companies doing compared with those of the rest of the world?

Carole Piovesan: We have one large language model that competes on a global scale, and that's through Cohere. Otherwise, our AI companies are by and large, from what I understand, relatively small compared with some U.S. companies. I think it's California that has 32 of the top 50 AI companies in the world.

Canada has a long way to go in augmenting and really globalizing our AI companies.

Gurbux Saini: I understand that some money was put to this in the last budget. Do you think it was enough, or do we need to do more?

Carole Piovesan: That's a little outside of my scope.

Looking at all the reports of the AI bubble and the amount of money being poured into U.S. AI companies, we would do very well not only to pour more money into some of our companies and our approaches in supporting them, but also to help them go global and to export our technology with rigour.

Gurbux Saini: You mentioned Singapore or the U.K. as a model our country should follow. Can you give us the reason you picked them over those of the rest of the world that are advanced?

Carole Piovesan: The U.K.'s model, which has a principles-based approach and is sectoral-focused, does not take regulation off the table entirely, but uses an iterative approach to better understand where there are gaps in regulation and where it can better target regulation to support, if needed, a horizontal law. If not, it can support regulators in providing effective guidance through a consultative approach with industry and other players. I think that aligns very nicely with Canada's approach, generally.

Singapore is interesting. It has had a number of sandboxes, particularly AI Verify, which has been effective in testing different trust models. It has invested quite a bit in the national AI literacy program. We see a different reaction to AI in countries like Singapore that are more trusting of the technology, which will lead to greater adoption and a strong focus on economics and competitiveness.

The Chair: Thank you, Mr. Saini and Ms. Piovesan.

[*Translation*]

Mr. Thériault, you have the floor for five minutes.

Luc Thériault: Thank you, Mr. Chair.

Artificial intelligence is truly a monster in the making and it's clear we can regulate it further. Earlier, we spoke about the former Bill C-27, which died on the Order Paper when the election was called. The bill proposed the creation of the position of an artificial intelligence and data commissioner.

What do you think of that recommendation?

[*English*]

Carole Piovesan: I think that position can be very useful, depending on its mandate and, importantly, on how it is resourced.

[*Translation*]

Luc Thériault: Artificial intelligence acts as a catalyst and a powerful accelerator of information control, and by extension, power. What is the best way to regulate this underlying motivation that guides developers?

The government, through its minister, has stated that it will put more emphasis on developing the economic advantages of the technology than on regulation.

What do you think of that approach? Shouldn't equal emphasis be placed on regulating artificial intelligence and on promoting its economic advantages?

● (1245)

[*English*]

Carole Piovesan: I interpreted the minister's approach to AI, to my recollection, as being "light, tight and right", which meant we would look to regulate, but in a tailored and focused manner. I think that aligns very well with the U.K. and Singapore approaches that I am proposing for Canada.

I was involved in the consultations related to the artificial intelligence and data act, and there were significant concerns about establishing a horizontal law that would apply in every context, but without sufficient consultation to determine if that was necessary.

I think the approach I'm proposing to this committee would allow us to balance the economic advantages with the controls that you speak about.

The Chair: I apologize again, Mr. Thériault, but we have another issue with the microphone. I'll stop the clock as we look into it.

Carole, could you unplug and plug in your headset again, if you don't mind? It gets really tinny and hollow for some reason.

Can you give me another test?

Carole Piovesan: Is that okay?

The Chair: That's much better, yes.

Carole Piovesan: Okay.

The Chair: I have to do this just for the sake of the interpreters. We don't want to cause any damage.

Mr. Thériault, you have two minutes and ten seconds left.

Luc Thériault: Thank you, Mr. Chair.

Ms. Piovesan, you have already explained that the regulation is intended to encourage organizations to change their culture, increase awareness, enhance digital literacy across the organization and adopt a more responsible approach to what they are building. You also said that it's a matter of mitigating actual risks upfront. In your opinion, there is still some uncertainty with respect to some regulations, but if we put this into perspective, the same themes continue to emerge.

What are these themes?

Carole Piovesan: I think that's right. The focus for the business use of artificial intelligence is an analysis of opportunity and risk. That is one of the themes I've talked about.

We want our businesses to go forward and adopt AI, and we've seen that on a global scale. Canadian companies are still lagging a bit in adoption. We want them to start iterating, using this technology and becoming much more confident and comfortable with the use of this technology. We want this to be within a context of responsible use, meaning a risk-based approach to the use of AI within a context, ensuring that appropriate governance controls are in place where the use of the technology is hitting up against a higher risk scenario, to be guided ideally by the sector, potentially by government.

We've seen this before through the compendium document that came out with AIDA, but also through a sectoral approach to what risk looks like in a particular environment. Then there's the company's own risk-based approach and how it would defend the risk classification it has put forward.

Luc Thériault: You spoke about balance earlier. Shouldn't the government's primary role in responsible innovation and design be to prioritize regulation, however light, to ensure it is upheld before pushing forward with economic development at full speed?

We want to increase efficiency, but as you said earlier, efficiencies can create other types of issues, including labour issues. All of this has not necessarily been assessed.

The Chair: Answer quickly, please.

Carole Piovesan: Yes, but we need a targeted approach, which is why I am looking to the U.K. and Singapore as examples for establishing an appropriate targeted approach to what and how we regulate the use of AI.

● (1250)

The Chair: Thank you.

Thank you, Mr. Thériault.

You have the floor for five minutes, Mr. Hardy.

Gabriel Hardy: Thank you very much, Mr. Chair.

I have a question for you, Ms. Piovesan. If the minister had to focus on just one task and all his work over the next few years hinged on a single decision, what, in your view, should that decision be, to ensure the responsible and effective development of artificial intelligence over the next few years?

Carole Piovesan: It would be to augment the emphasis and mandate of the safety institute to make sure that it is properly established to lead and support the safe and responsible development and deployment of frontier AI.

Gabriel Hardy: For greater clarity, what would that safety look like to you? What should be the priority when it comes to safety in artificial intelligence? Saying we need to act safely is a fairly broad statement.

Does this mean making businesses more responsible so that they can be held accountable in case they develop destructive technologies?

Is a very rigid structure needed upstream to avoid crossing the line?

What does safety mean to you?

Carole Piovesan: Safety is in the standards that are put in place for developing the technology, how we evaluate that technology, how we clarify what the standards mean and how we establish ongoing monitoring in real time of the use of that technology. Then it is in how we coordinate and harmonize internationally to ensure that other jurisdictions are following suit. That's where I would place the emphasis.

[*Translation*]

Gabriel Hardy: Earlier, you said that Canada is a leader in this area and that it should sell its expertise to the rest of the world. In what way does Canada stand out, and is it a leader in artificial intelligence? In what ways do our practices influence the rest of the world?

[*English*]

Carole Piovesan: What I meant by that is that Canada has put together a responsible AI brand through its co-founding of the Global Partnership on AI and its active involvement in the ISO AI standards—the 42000 series. In so doing, Canada has had an important role to play in establishing what responsible AI looks like. I think we should continue to do that not only by selling Canadian technologies globally, but by proving that our companies have established certain security safeguards that build trust in Canadian technology.

[*Translation*]

Gabriel Hardy: That is interesting, because two of the witnesses who appeared earlier told the committee that ultimately, incentives drive businesses to develop artificial intelligence at a very fast pace for financial reasons. We drew a bit of a parallel with social media. We recognize that they have very adverse effects, especially on young people and on our children's mental health. The witnesses noted that after technology has been developed, companies distance themselves from blame if users misuse it.

You have spoken about putting up structures and ensuring Canadian businesses develop artificial intelligence properly. If businesses were held accountable for the harmful consequences of artificial intelligence developed purely for profit, do you think this might discourage businesses from developing technologies that, while lucrative, offer no benefit to society?

[*English*]

Carole Piovesan: It depends on how we decide to hold a company accountable for something that it may not foresee.

I think there are certain disincentives you can put in place. There are also incentives we can put in place by providing guardrails and showing what "good" looks like. I think that might be a really effective approach to supporting our companies.

It's not to say that I'm contrary to anything punitive. That's not the point. It's just that we have to understand what we're regulating and what tools we're using to regulate and to enforce a regulation. It has to be tailored.

[*Translation*]

Gabriel Hardy: I would tend to agree with you. I like the comparison with social media or cigarettes. For a long time, people were told cigarettes were good and then they came to understand they were not, and measures were put in place to restrict advertising. I have a feeling that social media currently benefits from a big legal vacuum. Our children are hooked on social media and these companies have not faced any negative consequences.

I think we can learn from that and maybe direct the development of artificial intelligence to ensure we are not justifying the idea of profit at all cost.

Do you agree with that?

[*English*]

Carole Piovesan: I am, and I think there are certain areas where we could be far more active in ensuring transparency and disclosure with AI, particularly where you're looking at the use of public-facing chatbots that can be confusing to the user. There are notices we could put out to the public to make it much clearer.

Again, we need to be targeted and we need to understand the use of the general purpose technology.

● (1255)

[*Translation*]

The Chair: Thank you, Mr. Hardy.

Ms. Lapointe, you will be sharing your time with Mrs. Church. You have 300 seconds.

Linda Lapointe: Thank you very much, Mr. Chair.

Thank you for joining us, Ms. Piovesan. Your remarks about what we need to focus on have been very interesting.

Last week, we heard from a witness, Antoine Guilmain, who said he had a better understanding of protecting personal information. He stated that there are many laws and that we should identify any gaps in existing ones before creating new ones. You alluded to that earlier as well.

I know we had Bill C-27, which died on the Order Paper, but what would you recommend to close these gaps?

[*English*]

Carole Piovesan: I agree with that. The recommendation is to conduct more of a comprehensive study, maybe through regulators, as with the U.K. model, feeding information back into more of a central body to understand where there are specific gaps in the application of the regulation in the context of artificial intelligence.

[*Translation*]

Linda Lapointe: Thank you.

[*English*]

Leslie Church: Hello, Ms. Piovesan.

My question is about the principles- and values-based approach that you've talked about.

What are the best practices that you've seen in the U.K. or Singapore around enforcement and transparency? How do we ensure that the standards in place and the principles that govern the framework are taken to heart, adopted and adequately enforced? Also, how do we ensure that AI systems are transparent to the agencies or the government enforcing the principles?

Carole Piovesan: The examples we've seen—primarily in Singapore and, to a lesser extent, in the U.K.—are through the consultative and collaborative process between market participants, if you will, and the applicable regulator. It's done on an iterative and ongoing basis to see how the application of those principles is occurring and where there are deficiencies or challenges in the application of those principles.

As we become more and more aware of where the risks are materializing, we will start to see a greater emphasis on regulation, with stronger enforcement. It may also be that for regulators, we will need to augment certain enforcement capabilities so they are able to exercise whatever jurisdiction they have in the application of their particular sector or body of regulation and AI.

Leslie Church: Given the commercial sensitivities that are no doubt present in the sector, how do we ensure that companies are fully transparent and compliant?

Carole Piovesan: I think it will be a challenge to ensure full transparency, since many AI companies are relying on trade secrets as a mechanism for IP protection. That was embedded in the earlier AIDA model. To the extent that we proceed with a body such as a data and AI commissioner, that might be a role we look to augment within that particular office.

Leslie Church: One of the groups I've met with is the Kids Help Phone—this is more in the context of some of the chatbots and public access to AI currently—and they've talked about a standard of care existing.

Do you think there is an approach we could take to ensure that as AI develops and is publicly utilized, we can create some guardrails to create a standard of care, particularly in situations where we're dealing with harms that children could be facing?

Carole Piovesan: I think we will start to see a change in the applicable standard of care, whether it applies to children, to the health care sector or to the auto sector. We will start to see a changing standard of care that is applicable to and conscious of the use of AI in a particular case.

The Chair: You have 50 seconds.

Leslie Church: Would you say there is a sufficient duty of care by companies that are innovating in this area right now under Canadian law?

● (1300)

Carole Piovesan: It's all very context-specific. In certain sectors, there are existing frameworks that are working quite effectively to stem irresponsible investment in AI in the use context. Whether or not they're experimenting internally is one thing, but how it's ultimately used is another. I think in certain contexts there are effective safeguards. I also think that many sectors are adapting, with concern that the safeguards aren't good enough.

Look at health care. Health Canada came out with "Software as a Medical Device" to help us better understand how we can assess the risk of medical devices. That's important. The development of AI in health care is one thing, but we need that kind of tailored approach to ensure that when it gets to bedside, it meets certain standards, and we have to know what those standards are.

Leslie Church: Thank you so much.

The Chair: Thank you, Ms. Church.

Ms. Piovesan, I want to thank you on behalf of the committee for appearing today. I appreciate your input.

I have a couple of things for the committee. In case you haven't seen it yet—it was distributed among committee members—we heard from the office of the AI minister that he will not be appearing for this study. That's in the digital binder. Also, as a reminder to committee members, the Ethics Commissioner will be appearing on Monday of next week.

That's it for today. Thank you, everyone.

The meeting is adjourned.