



HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

45th PARLIAMENT, 1st SESSION

Standing Committee on Access to Information, Privacy and Ethics

EVIDENCE

NUMBER 024

Monday, January 26, 2026

Chair: John Brassard



Standing Committee on Access to Information, Privacy and Ethics

Monday, January 26, 2026

• (1530)

[English]

The Chair (John Brassard (Barrie South—Innisfil, CPC)): I'm going to call the meeting to order.

I want to welcome everybody and wish everybody a very happy new year.

[Translation]

Welcome to meeting number 24 of the House of Commons Standing Committee on Access to Information, Privacy and Ethics.

[English]

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

I would like to welcome our witness for the first hour today.

We have Wyatt Tessari L'Allié, who is the founder and executive director of AI Governance and Safety Canada.

Now, we were to have a second witness, Etienne Brisson. My understanding is that he is facing some serious road conditions.

Mr. Hardy, you took the same route this morning.

[Translation]

Roads conditions were bad, weren't they?

[English]

Okay, so if he does get here for the second hour, I'm going to be glad to include him as a witness in that second hour.

Mr... I'm going to call you Wyatt, okay? You have up to five minutes to address the committee. Go ahead, sir.

[Translation]

Wyatt Tessari L'Allié (Founder and Executive Director, AI Governance and Safety Canada): Mr. Chair, committee members, thank you for honouring me with the invitation to address you today.

AI Governance and Safety Canada is a non-profit, non-partisan organization, as well as a community of people from across the country. We start by asking the following question: What can we do in and from Canada to ensure advanced AI is safe and benefits everyone?

Since 2022, we've been making public policy recommendations to the federal government, such as our presentations on the AI bill and related data, as well as addressing parliamentary committees on the matter.

[English]

So far in this study, you've heard about the impacts that Canadians are already dealing with. Even with current systems, chatbots have talked teenagers into suicide, and developers can't reliably predict what the models will do.

You've heard that, with capabilities continuing to accelerate, there are much bigger risks fast approaching and that global companies like OpenAI and Google are competing to build smarter-than-human AI systems in the near term, systems that they themselves admit they won't know how to control.

If a nuclear power plant melts down, it's a tragedy, but the rest of the world moves on and eventually recovers. With smarter-than-human AI, we may not get a second chance. If, through accident or poor design, a system interpreted human beings as an obstacle to achieving the goal it was given and started taking actions against us, there is no guarantee that technologists or governments would ever be able to regain control. It would be a global crisis the world might never recover from.

If you find the situation downright scary, you are not alone. The question is, what do we do? As Canadians sitting around this table in 2026 looking at the exponential advance of AI, mostly driven by entities outside of our borders, what can we do?

If we want, we can try to play whack-a-mole with current AI impacts and ignore the bigger picture within which they fit. We can try to deny or dismiss what the leading labs are building, wasting the limited time we have to operate, or we can take a hard look at where things are heading and start preparing now in a manner that also addresses current risks, because if we're not ready to give up, Canada has a number of options at its disposal.

In October, we published our white paper, “Preparing for the AI Crisis: A Plan for Canada”. In it, there are four key recommendations.

First, pivot to meet the AI crisis. The development of smarter-than-human AI is the biggest threat to Canadians’ safety. For that reason alone, it deserves to be a top priority. AI will disrupt almost every other file you’re working on, from national defence to jobs to health care to education to energy and the environment. Much like with COVID in 2020, there are times when the responsible thing for government to do is pivot to address the developing crisis and reassess the priority of other files accordingly. Given its wide scope and long-term implications, AI needs to be a cabinet-level priority, and action needs to be coordinated with opposition parties and the provinces.

Second, spearhead global talks. The race to smarter-than-human AI is a global phenomenon that no country can manage on its own. At this time more than ever, the world needs leadership, and Canada is well placed to deliver it. The strongest card we can play is to advance global talks and solutions and lay the groundwork for an AI treaty that the U.S. and China might sign when the crisis hits and they realize they have no alternative.

Third, build Canada’s resilience. While domestic action alone cannot protect Canadians, plenty can be done to mitigate the secondary impacts, such as putting in place supports for displaced workers, banning deepfakes and strengthening critical infrastructure against cyber-attacks. By taking the initiative at home, Canada will be in a stronger position to navigate the AI crisis and negotiate from a position of strength.

Fourth, launch a national conversation on AI. Canadians deserve to be informed and consulted on a technology that will fundamentally reshape their lives. We need nationwide public hearings to educate and consult on core societal decisions pertaining to our future with AI.

Last week, Prime Minister Carney put Canada in a leadership role on the world stage. This is an unprecedented opportunity to push for global AI safety measures while building resilience at home and to be the adult in the room when it matters most. The stakes couldn’t be higher. The clock is ticking. Let’s get to work.

• (1535)

The Chair: Thank you for your opening statement.

We’re going to have six-minute rounds, starting with Mr. Barrett. These are questions, and it’s going to be back and forth.

For the sake of the interpreters, I would ask you to speak a little bit more slowly in your responses, if you don’t mind. I know they had your opening statement, which was fine, but we want to make sure we have proper interpretation.

Go ahead, Mr. Barrett, for six minutes.

Michael Barrett (Leeds—Grenville—Thousand Islands—Rideau Lakes, CPC): What specific evidence-based risks justify urgent federal intervention of the kind you’ve suggested? How do we avoid the situation where the policies that are made are driven by fear instead of the actual situation on the ground? Are there ex-

amples or evidence, perhaps, that you would be able to share with us?

Wyatt Tessari L’Allié: Sure.

So far we’re seeing a bunch of early warning signs, everything from chatbots talking teenagers into suicide or job impacts on the youth or record levels of scams and AI-powered cyber-attacks. These are important and need to be dealt with, but they’re not in themselves a justification for whole-of-government action. Why they matter is where they’re headed. Right now we have teenager AI: very simple systems, or relatively simple systems, I should say. What all the leading AI labs are actively building.... Particularly if you go to their websites, they’re saying that they are building smarter-than-human AI systems. According to their CEOs and according to a lot of experts, including the engineers who left these organizations as whistle-blowers because they don’t trust them, also say that yes, in two to five years, smarter-than-human AI is possible.

The reason governments need to act currently and be proactive is really the fact that we may not have much time to prepare for the much bigger risks to come. Given that it will require global solutions, and global solutions take forever to put in place even if we have 20 years, we’re in a race against time.

On the point of hope and fear, I fully agree we can’t.... I’m doing what I’m doing because I believe there are still solutions and because I think there are positive ways forward. We can’t let fear paralyze us, and we can’t tell ourselves that AI is no good at all, because there are a lot of really good applications of AI, in health care, in energy, in all that kind of stuff. I think it’s very important to be neither enthusiastic nor pessimistic about AI and just be very clear-eyed about what’s coming and how we can prepare.

Michael Barrett: Is there more than one strategy that needs to be employed when we look at some of the real effects that we're seeing from the use of this new teenage generation of AI? Some of the most real effects we've seen, of course, are news reports about language models counselling vulnerable people to die by suicide, to take their own lives. We've seen also a rise in the creation of child sexually exploitive material and deepfakes using children. That is the worst kind of deepfake, but it's not the only kind. It's happening to public figures, to anyone who has a picture online, and really to anyone whose likeness can be described to a language model. These are some of the real-world consequences we're seeing today. You referenced the effects on the job market for youth for entry-level jobs.

What's the answer? Is it a series of measures that need to be taken? When we talk about deepfakes, is it a question of needing to update the criminal law so that individuals are held personally responsible for their actions in the creation of this unacceptable material that is not intended to be covered by free speech laws or freedom of expression and goes well beyond that? It's victimizing individuals. Is it instead that we need to pass laws where it's incumbent on the tech companies to ensure the safeguards are in place? Is it both?

• (1540)

Wyatt Tessari L'Allié: It's yes to both.

I do think that improvements to the Criminal Code to be able to deal specifically with deepfakes would help. I would also say, as witnesses before me have said, that current laws could go a lot further than they currently do in terms of protecting Canadians. If we gave more resources to the current regulators to be able to apply them in the context of AI, that could be a much faster way, and possibly a more effective way, of getting protections in place.

In terms of the responsibilities for technologies—technologists—I think it's very hard to stop a teenager in their basement in Russia from creating a deepfake of somebody. However, you can tell Google that if it wants to do operations in Canada, it has to take down deepfakes within a certain amount of time so that while the image itself is created, it doesn't get spread and it doesn't harm people's reputations.

I think it's yes, strongly, to both.

Michael Barrett: All right.

I would just say in the last 10 seconds that with all of the potential risks and how catastrophic they can be, we might first have to address these very real and serious challenges, like the ones we just talked about.

Thanks very much.

The Chair: Thank you, Mr. Barrett.

[*Translation*]

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

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

Happy new year everyone.

I'd like to thank the witnesses for their presentations.

Before asking the questions I have regarding the ways the government and the House of Commons can intervene, I'd like to set the table by explaining exactly what this is about. When it comes to the digital world in general and AI in particular, I like to separate the two. Whether we like it or not, AI development is such that it is becoming a basic infrastructure, just like electricity, transport and the Internet. It is also used in decision-making and operational processes, and in geopolitical and political spheres, which is feeding our dependence on it.

Generally, when citizens depend on a technology, or anything else for that matter, it creates a certain vulnerability. I don't think we can do anything about our dependence on AI or even reduce it. That's my opinion, as someone who's worked in this field. We can reduce our vulnerability, but not our dependence, because AI is here to stay.

Mr. Tessari L'Allié, as an expert in this field, what do you think a government or a government policy framework can do exactly? What should we prioritize to reduce the vulnerability associated with AI's rapid development?

Wyatt Tessari L'Allié: Because of the current context and how fast AI is being developed, we need to stop trying to react to yesterday's AI. It hasn't worked. We have to plan ahead. If it takes two years, three years or five years for the government to adopt a related legislation, we have to imagine what AI will look like in three years or five years. That's why we need to focus on planning ahead instead of focusing on regulations.

Regarding dependence, workers are already forgetting how to do certain tasks, because AI can do them better and faster. People forgetting how to do certain tasks aren't the only ones experience a loss of resilience. Society as a whole is also experiencing a loss of resilience, because if all of a sudden AI breaks down or is taken away, people won't know how to do certain things.

To reduce vulnerability, society needs to know how to do things. AI tools are very useful to improve productivity and do things faster, but at the same time, our education system needs to continue teaching people how to do things by themselves to reduce this dependence and vulnerability.

• (1545)

Abdelhaq Sari: I've said in the past that there are two things to consider. On the one hand, there's the technology itself, which we can't control, because as you said, someone can develop a new technology in their basement. On the other hand, there's the user that can be educated. You talked about education. I don't want to influence you or interfere in your field of expertise, but I'd like to expand on that.

Is there a way to establish or to recommend some form of digital sovereignty? Without such sovereignty, it's hard to talk about the other technologies over which we have no control.

Wyatt Tessari L'Allié: You're absolutely right, especially in the context where we are being threatened by the U.S. We purchase most of our computational power from the U.S., from American data centres. If the U.S. suddenly decided to cut us off or limit our access to this power, it would greatly hinder our economy.

In short, digital sovereignty is important. We need our own data centres and the ability to do what we want here at home, without having to depend on another country.

Abdelhaq Sari: If we're not digitally sovereign, how can we control tech giants like Google and Microsoft? It's hard to get around them. I've read that many countries, including from the EU, are trying to establish their own digital sovereignty.

Meanwhile, is there a way to gain some control—and I insist on the word “some”—on how to introduce these solutions in our own market and how our citizens can use them?

Wyatt Tessari L'Allié: As I said before, if we want to have control, we need the ability to create our own AI tools, have our own data centres, and depend on an educated workforce able to do the work without relying on AI.

Abdelhaq Sari: I'd like to end by coming back to something I said earlier.

I don't think the government can control people's dependence on AI and how the technology evolves. That said, as you pointed out, we have the capacity to train people to reduce their vulnerability and ensure the tools used in all our systems, whether they be in health care, finance or transport, are safer. That's important, because AI has access to our data, our history and our future.

Thank you very much, Mr. Tessari L'Allié.

The Chair: Thank you, Mr. Sari.

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

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

Mr. Tessari L'Allié, I'd like to start with a broad question. I also don't want to offend you.

In your document, you talk about strengthening critical infrastructure against cyber-attacks, developing AI and drone defence capabilities, and preparing security agencies against biological, nuclear and chemical weapons proliferation, to name a few, which general AI can facilitate. I think you'll agree with me that this would be taking a defensive stance against ever more powerful systems.

I agree on the idea of defence, but in the end, isn't that a losing strategy? Wouldn't we need to eventually stop the development of systems more powerful than our defences?

Wyatt Tessari L'Allié: Absolutely. There's the era before super-human AI, and the era after. Regarding national defence against drones, AI and others, in the short term, the government needs to increase investments in critical infrastructure to protect them against cyber-attacks. You are right that if we reach a point where AI can thwart all our plans, Canada won't be able to defend itself without help.

Our best strategy relies on one of Canada's strengths, which is to move talks along with international partners, because every other country faces the same issues. Neither the U.S., nor China or Russia can defend itself against a supersmart AI. Everyone should work together, because everyone is vulnerable.

I basically agree with you. There's a lot we can do in the short term that could lead to greater security resilience. That said, in the long term, an international solution is the only option to prevent these systems from being created, at least until we're able to control them.

• (1550)

Luc Thériault: What you're saying is that, based on what you know, we're already losing control and, in the medium term, there's a strong possibility we won't be able to control this technology.

Wyatt Tessari L'Allié: That's correct. I refer to experts in the field. I'm more of a generalist interested in the intersection between government and technology. Experts working in state-of-the-art AI laboratories developing those systems say they don't know how to control them and it scares them. The only reason they're accelerating development is they hope someone else will be able to control those systems if they can't, that those people might have a better chance of doing so. Such a race makes no sense.

The consequences of the current systems are significant, but not jarring. It's not the end of the world. An issue with ChatGPT could lead to a teenager losing their life or a cyber-attack, but we could get over that. However, if AI reaches a point where it can understand what we do better than us, act more quickly than us and thwart our plans, we'll be vulnerable to its decisions and our security forces won't be able to defend us.

Luc Thériault: Not to mention what we were talking about earlier regarding the misuse of this technology beyond the negative or pernicious impact on teenagers, for example.

Wyatt Tessari L'Allié: Absolutely.

Short-term risks could be biological, for example. AI can already be used to create new viruses the human body would have a hard time fighting. There are biological and nuclear weapons, among others. If a human can imagine something, an increasingly competent AI could do the same and use it. Even if we don't think we can control AI, the simple fact that it could develop a weapon of mass destruction is reason enough to take this matter seriously.

Luc Thériault: Looking at various governments or world superpowers, geostrategic and geopolitical positioning, and the interest superpowers have in increasing their power, are you optimistic?

There's little framework around AI right now. It's a black hole, unlike nuclear weapons. When it comes to nuclear weapons, we can see all of that. How can we control AI?

Wyatt Tessari L'Allié: Countries are really already racing against one another for other reasons. The United States talks about wanting to dominate in artificial intelligence. China wants to lead the world in artificial intelligence.

Luc Thériault: What's the real goal? When we talk about the United States and China. What's their goal?

Wyatt Tessari L'Allié: Right now, it's more of a military and economic goal. They want artificial intelligence in order to win wars and defend themselves against others and to build the strongest economy. There are powerful incentives at the national level to become a leader in artificial intelligence. The companies themselves have the same incentives. They compete against one another. All these factors are pushing us faster towards a breaking point.

My optimism, and I would say realism rather, stems from the fact that all these companies and countries are facing the same issue as we are. If someone creates a superintelligence that we can't control, everyone loses. Whether it's Donald Trump or Xi Jinping, they'll need to collaborate at some point. If they don't, they'll lose.

The Chair: Okay. Thank you.

Your time is up, Mr. Thériault. You had more than six minutes.

We'll start the second round of questions.

[*English*]

Mr. Cooper, you have five minutes. Go ahead, please.

Michael Cooper (St. Albert—Sturgeon River, CPC): Thank you, Mr. Chair, and thank you to the witness.

You cited in your testimony AI chatbots, and I want to drill down a bit on that topic, specifically as it pertains to youth.

First of all, I take it you would agree that this is an area where there is a need for regulation. Before a U.S. judiciary subcommittee hearing last fall, as well as in other reports, U.S. data was presented from Common Sense Media that 72% of teens in the U.S. have used an AI companion at least once.

Do you have any data for Canada? I would take it that it would be similar.

• (1555)

Wyatt Tessari L'Allié: Yes, I think it's similar, but I haven't seen exact numbers.

Michael Cooper: You were asked, in general, about some of the AI harms, looking specifically at chatbots. Could you elaborate on some of those harms for youth?

Wyatt Tessari L'Allié: Absolutely.

The whole relationships and mental health piece is huge, because we're basically doing this giant experiment with our kids by introducing these AI companions into their lives, and we don't know what that will do for their ability to socialize and work together. We're already seeing a lot of mental health concerns. I'm sad the other witness isn't here, because he is definitely a lot more experienced in this area.

We see AI addiction in kids who can't give it up because these systems are so pleasing and sycophantic. They always tell you what you want to hear. They get caught in these loops and go down these dark holes.

This is a live experiment, and we're still waiting to see what the long-term effects will be. It is very concerning, because these are key moments in their development, and if their mental health and learning are being messed up, then that's a problem.

Michael Cooper: These systems are trained based on the entire Internet. Is that fair? Is that accurate? That would include everything from suicide forums to porn sites to other harmful content, and this will inevitably make its way, and is making its way, into these chats that youth are having.

Wyatt Tessari L'Allié: Yes, I think some of the companies are making efforts to limit what data goes into it, but the problem is, the more general the model, the more capable it is, so the more you train it on a variety of information, the more it understands how it all fits together. I can't speak to exactly what data is going into the models, for example, of Gemini or ChatGPT, but the fact that they have been talking teenagers into suicide and the fact that they do occasionally produce instructions on how to build a chemical bomb suggests that they have been trained on very dangerous material.

Michael Cooper: Would you agree that what compounds the challenge in terms of some of the risks and the ability for parents to detect that their loved one is being exposed to sexually explicit or harmful content is that these are often invisible, and, in fact, there is no transparency that you're even engaging with AI in some instances?

Wyatt Tessari L'Allié: A lot of kids prefer to talk to ChatGPT because it feels safer than talking to a parent. They're not aware of the privacy concerns, and they're not aware that these systems are not as reliable or hopefully as wise as the parent is.

Yes, it's a very rapidly evolving problem with a lot of ways to go sideways. We're seeing some impacts already, and we expect to see more.

Michael Cooper: In the fall, a bipartisan bill, the GUARD Act, was introduced in the U.S. Senate by senators Hawley and Blumenthal. Are you aware of that?

Wyatt Tessari L'Allié: Somewhat, yes.

Michael Cooper: Okay. In short, at a high level, it would do three things. It would ban AI companions for minors. It would mandate AI chatbots to disclose their non-human status, and it would provide new penalties for companies that make AI for minors that solicit or produce sexual content. I'd be interested in your thoughts on those three components.

Wyatt Tessari L'Allié: Those are broadly the directions you want to take in terms of the protections and liability, yes.

Michael Cooper: Consistent with the need for transparency, I would also note that there was a report that was submitted recently to the government's AI task force by one of its members that calls for, among other things with respect to AI products, visible labelling, source transparency requirements, metadata and digital watermarking. Are those measures you would also support?

• (1600)

Wyatt Tessari L'Allié: Absolutely. People should be able to know that they're interacting with the AI system. It's not always obvious right now, and it will be harder and harder as the time goes on, so yes, labelling is a bare minimum.

The Chair: Thank you, Mr. Cooper.

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

Leslie Church (Toronto—St. Paul's, Lib.): Thank you, Mr. Chair, and welcome, Mr. Tessari L'Allié.

Maybe just picking up on some of my colleagues' questions, I'm also concerned very much about the impact of AI socially, on kids in particular, and one of the things that I noticed in your white paper was that part of your approach to building Canada's resilience and protecting online safety was around the concept of requiring AI labelling and banning unacceptable capabilities. I'm just wondering if you could talk a little about the labelling in particular, how you see that developing and maybe how that could be useful for us?

Wyatt Tessari L'Allié: Yes. On labelling, there are a bunch of ways to go about it. We're agnostic as to these technical paths, but certainly it's something that could be developed as a global standard. For example, the Internet protocol that helped the Internet happen is a global standard. It's the same thing with labelling. I've

heard proposals, for example, of having Unicode, basically the language behind the text, be labelled as AI so that a computer would automatically know if a letter is written by AI or not. There are a lot of solutions like that, for example, that could happen.

Ultimately it's probably going to take some government impetus to force these companies to actually do it, because it will be a pain to actually label, and there will always be the challenge that even if you label a piece of text, somebody can take a photo of it and copy it over to another computer, and now suddenly it's no longer labelled.

It's the kind of thing where you can't.... Labelling won't stop misuse in that sense, but it can, if you catch somebody using AI that isn't labelled, give you an opportunity to take action. It's not a full solution on its own, but it's a step in the direction of giving incentives for people to actually use the AI correctly.

Leslie Church: Talk to me about what that would even look like to me as a user. As somebody who is on the Internet and social media regularly, what would I see as a consumer? How would that label be applied?

Wyatt Tessari L'Allié: There are a bunch of ways to go about it.

Some, for example, would even say something like font colour. If you're reading a text and you see it's in a certain colour, you'll say, "Oh, that's AI." It's the same thing with voice over the radio. If you hear a certain tone in the background, you think, "Oh, that's AI."

There are a thousand different technical ways to go about it, but it would have to be something that would make a user think instinctively, as soon as they saw it, "Oh, yes—this is AI," and have that be a global standard.

There is the logo of the little four-legged star. You'll see that around in various things. It's a step in the right direction, but we're at the very beginning of a very big, complex problem. It will take clear direction from governments around the world, basically, that this is required and that we expect this of companies that are building AI.

Leslie Church: When you talk about prohibiting other unacceptable capabilities, just to use the language from your white paper, what types of capabilities are you talking about? Can you expand a bit on that?

Obviously, we're very concerned at the moment. We've brought forward criminal legislation to deal with things like deepfakes, but are there other capabilities that you would particularly encourage us to focus on?

Wyatt Tessari L'Allié: Yes.

To mirror the EU AI Act, systems that are high and medium in the unacceptable category are systems that, for example, refuse to be shut off, systems that deceive the user and systems that will modify themselves without notice. For example, if you give an AI system a task of doing something and it calculates that it needs to modify itself in order to achieve that task, that should be an unacceptable capability, because suddenly your system is a very different system from what was created, and the risk profile changes dramatically.

Another one is autonomous self-replication. If you ask your system to do something and it calculates that it needs to make a bunch of copies of itself on different servers so that if the first server it's on goes under, it can still keep running, that's a problem, because suddenly your model is no longer just on your computer. It's also on 10 other computers, and you don't necessarily have access to them.

In our recommendations for the AI and data act, we go into details on which are the biggest ones, but unprompted self-modification, commandeering of resources—if your model starts stealing in order to achieve its task—are all behaviours we're starting to see happen in test settings, and if we allow them, then we're very much in a vulnerable position.

Leslie Church: One of the things I'm very interested in as well is establishing a duty of care for these platforms. It's a legal concept that I think belongs here, and certainly there need to be ways to put guardrails on, particularly for what youth and children are exposed to, and ensure that in the situations that you've raised, there's a clear liability when AI models are presenting incorrect or very harmful information, especially to kids.

• (1605)

Wyatt Tessari L'Allié: Liability has to be with those who understand best how these systems work, and I think that's where the top talent is at Google and Meta and the rest. They should be responsible for the behaviour of their models, even after deployment, because you can't ask a teenager or somebody who doesn't know anything about AI to know what its harms are and how to use it.

Leslie Church: Thank you.

[*Translation*]

The Chair: Mr. Thériault, you have the floor for five minutes.

Luc Thériault: Thank you, Mr. Chair.

Mr. Tessari L'Allié, you wrote in your paper that, at its core, the governance of general artificial intelligence poses a challenge of human coordination and technical skill. Anthropic's CEO recently said that we understand maybe 3% of the inner workings of these systems.

If we don't understand how it works, why not focus on negotiating a global ban?

Wyatt Tessari L'Allié: If we see the crisis coming and we don't know how to avoid it from a technical standpoint, we have no op-

tion but to try to ensure that the technology isn't created in the first place. For example, if a number of companies in downtown Montreal or Toronto were creating and developing a nuclear technology that places everyone around them at risk, we wouldn't be wondering how to balance safety and other aspects. We would be sending in the police to stop them.

In this situation, if companies are creating a technology that will place everyone at risk, and they themselves admit that they don't know how to control the systems that will be created, the government must take responsibility for protecting the public and for putting these developments on hold, at least until we know how these systems can be set up safely.

The government's main role is to protect the public. If technologists are doing something too dangerous, the government's role is to stop them. To do this, we need every country in the world to participate. If someone, somewhere in the world, is creating these systems, we're in great danger.

Luc Thériault: As you said in your paper, a responsible solution for the future of artificial intelligence requires a global agreement. Yet this global agreement could quickly become a global ban. Is this accurate?

Wyatt Tessari L'Allié: Yes. It could at least provide a break.

I'm well aware of the current geopolitical dynamics. However, as I said before, the United States and China are in the same boat as us. They don't want to create systems that will cause them to lose power, place people at risk and generate large-scale unemployment.

Even though they don't want to work together, they need to. Since we play a role of power within the group, our best course of action is to try to keep the talks going and to implement the treaties. That way, when the world is ready to sign them, everything will be in place to do so.

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

The Chair: You have two minutes left.

Luc Thériault: Mr. Tessari L'Allié, your paper states that laboratories could succeed in implementing a global solution in one to three years. Yet we may have less than 18 months before that happens.

Given what we just said, can you elaborate on this?

Wyatt Tessari L'Allié: The future is hard to predict. If we're lucky, we'll have 10 to 20 years to prepare for it. We can take our time with the rest.

If we're unlucky, a cutting-edge laboratory could unveil a new model tomorrow that would basically amount to superintelligence. It probably won't happen, but maybe it will. For this reason, a responsible government will have no choice but to prepare for the closest scenarios.

If we're lucky, we'll have more time to improve our solutions. However, we need to get started. We must tell ourselves that we don't have time. We really need to act now and lay the foundations. We need to either take a break or implement other solutions to ensure that we don't end up in a world where we lose control.

Then, in terms of the number of years involved, the CEOs of Anthropic, OpenAI and Google DeepMind are talking about two to five years. Other experts say one to three years. We don't know. However, if these are smart people who are keeping a close eye on developments and telling us that we must be ready—and I hope that they're wrong and I strongly hope that we'll have 20 years to figure this out—as a precautionary measure, we need to prepare ourselves for short-term scenarios. After that, we'll take our time.

• (1610)

Luc Thériault: Ultimately, we need to take precautionary measures.

Wyatt Tessari L'Allié: Exactly.

The Chair: Thank you. You have 20 seconds left.

Luc Thériault: I thought that I had one minute left. I figured that I had gone over the minute. Very well. I'll pick it up later.

The Chair: There are 14 seconds left.

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

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

I would like to acknowledge our witnesses and welcome them to the committee.

When we talk about artificial intelligence here, we often picture it a bit like in the film *The Terminator*. We see the end of the world coming. It seems that, for some people who follow our work, or even for decision-makers, the situation appears unlikely and we shouldn't worry too much. We heard earlier that we might have only a year and a half to prepare, but maybe 20 years. We don't know.

I would like to focus on a practical matter concerning day-to-day life. We're experiencing it. Right now, inflation is quite high in Canada. Fifty per cent of Canadians are within \$200 of insolvency. There are 2.2 million people lining up for food every day.

Furthermore, some companies see the potential of artificial intelligence to reduce their workforce, given that artificial intelligence can work 24 hours a day, seven days a week. Do you think that we should be looking at this issue from a government perspective? Shouldn't the government feel compelled to exercise caution given this possibility?

Using artificial intelligence instead of hiring a Canadian will cost money. All the people who lose their jobs will inevitably end up needing food banks, which are already overwhelmed.

Wyatt Tessari L'Allié: For example, only a year or two ago, Spotify told its employees not to hire anyone until they were sure that they couldn't use artificial intelligence.

Companies have already given clear instructions in this area. They know perfectly well that, to remain competitive, they'll need to cut spending. Yet salaries are the biggest expense. If we end up in a world with soaring unemployment rates, the reduction in income tax revenue will take a huge toll on the government.

We can currently see that the impact on employment is concentrated in certain sectors, such as information technology and the creative industries. This phenomenon will ripple throughout the economy. We don't know whether it will happen quickly or slowly. We'll see.

However, without a plan in place to support these people who will lose their jobs and to ensure that the government generates enough revenue to maintain operations—if it's so dependent on income tax revenue—the current situation is really a harbinger of things to come.

Gabriel Hardy: Okay.

I'll now move on to another topic.

As you know, Canada has one-fifth of the world's drinking water. Moreover, I recently read that generating an email using artificial intelligence consumes about 500 millilitres of water. If we paint a picture, it's even more. According to my notes, a ChatGPT search consumes ten times more electricity than a Google search.

Shouldn't we start legislating on something as tangible as drinking water? Obviously, if artificial intelligence starts to take more and more of our drinking water, perhaps humans will be the first to suffer the consequences. Agriculture will be next. I imagine that we're already interacting to some extent with data centres, given that we increasingly need them because of artificial intelligence. These centres consume water and they're actually harming human life.

Wyatt Tessari L'Allié: Yes. The major difficulty with water and electricity consumption figures is that we don't know the exact numbers. The companies don't disclose them. Certainly, in some municipalities experiencing water restrictions, data centres make a big impact. I heard other figures suggesting that the impact is almost the same for a golf course, for example. In other words, yes, there has already been an impact in this area. The first step is really to force companies to tell us how much water they use.

Gabriel Hardy: Clearly, we increasingly need data centres. Yet the more we use artificial intelligence, the more data centres we need. The more data centres we need, the more water we need. In terms of the golf comparison, not everyone with a phone is playing on a golf course. If we contrast this with the use of artificial intelligence, everyone has this feature on their phone. If we consider that a quick search costs one 500-millilitre bottle of water, we can imagine that we consume quite a lot every day.

• (1615)

Wyatt Tessari L'Allié: Yes. In terms of global consumption figures, the United States is talking about increasing electricity consumption across the country by 20%, just for artificial intelligence, in the next 10 years. This means billions and billions of dollars of investment. Nuclear power plants are being reopened to supply the electricity for this.

It's certainly happening on a massive scale. Again, I would say that the first step is really to ask the companies to provide the exact figures for their consumption. Right now, there's a great deal of misinformation about this. We know that the consumption takes place, but we don't know the exact numbers.

Gabriel Hardy: I won't have time to ask my question.

The Chair: Thank you, Mr. Hardy.

[English]

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

Gurbux Saini (Fleetwood—Port Kells, Lib.): Thank you, Mr. Tessari L'Allié, for coming.

Some argue that AI regulation will limit innovation and competitiveness, while others say regulatory certainty is essential for responsible scaling. Based on your experience, does effective AI governance ultimately enable or constrain innovation?

Wyatt Tessari L'Allié: That is a tough question, because the technology is moving so fast that a good regulation for today's AI may be out of date by next month.

Right now, in our white paper, for example, we don't openly call for much new regulation, simply because there are so many areas where it is really hard to monitor what's going on in AI and to regulate software. It's a very fast-evolving technology. However, there are some basics you can put in place, like, for example, banning deepfakes and protecting children.

Also, to an earlier point, there are a lot of existing laws, like human rights laws, for example, with biased algorithms. That is something you should be able to protect Canadians against through existing laws by investing in the regulators and ensuring the Privacy Commissioner has enough power and resources to be able to track everything that's going on.

Given how fast things are moving and how hard it is to accurately create laws that will be robust into the future, focus on the simplest pieces, like banning deepfakes or requiring labelling, and on applying current laws as well as you can.

One analogy that comes up is that right now we're dealing with a cute little bear cub, and you need legislation to avoid scratches and bites. If it takes you five years to put a law in place, by the time you

do it you'll then be dealing with a full-grown grizzly, and the laws and protections you'll need are very different. Given that we're in this transition period, doubling down on current laws and making sure the regulators have the power they need to be able to act on them while focusing on the things you can do relatively quickly and simply is probably your best strategy.

Gurbux Saini: What role should AI have in running part of the government? Do you have any experience or knowledge of whether any other governments around the world are using it?

Wyatt Tessari L'Allié: It is clear that AI can be used to make government more efficient and to offer more services to Canadians at a lower cost. It is also clear that part of the reason you can do that is that it will cost less than salaries, since you'll be replacing a lot of human work with AI. That is a balance that I don't claim to have the right answer for.

I will say that I use AI every day. Its productivity is very useful. I would encourage government to use it wisely, obviously, within guidelines and with privacy impacts in mind. This is the challenge. The more useful these systems get and the cheaper and the more effective they are, the more we'll be dependent on them and the less we'll be relying on human labour.

When we talk about having a national conversation on AI, we need to talk to Canadians about this: If we reach a point where AI and robotics can do essentially all work, should we automate everything, or do we want to keep jobs for ourselves? If so, under what conditions would that be?

That is a very big conversation, and that's not to mention the psychological impact. For a lot of people, their sense of self comes from their work. If you tell them what they're doing is no longer needed because now AI can do all the coding and they can retire now as a coder...or if you're telling a voice-over actor that they can now retire because AI can do all the voice-over acting, are we going to ask that actor to become an engineer or something else?

The psychological and societal impacts of this change happening so fast are something that we need to be proactively thinking about and managing. This is where having a national conversation—giving people the opportunity to understand what's going on and where things are headed, and to have a meaningful ability to provide input as to what they want—is an essential piece of the solution.

• (1620)

Gurbux Saini: One thing you mentioned in your earlier remarks was weapons of mass destruction and nuclear weapons. I don't think the U.S., China and Russia care about anything as long as they are the superpower in that field.

What can the rest of the world do to prevent humanity from being at risk of being wiped out?

The Chair: I need as quick a response as you can give.

Wyatt Tessari L'Allié: As much as they like power, they also need to collaborate, because they are also vulnerable to AI. That's where our opening is.

The Chair: That was quicker than I thought. Wonderful.

[*Translation*]

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

We'll finish with Ms. Lapointe, who will also have five minutes.

Mr. Brisson has just entered the room. In the next hour, we'll be hearing from three witnesses, so we can take a bit more time to talk.

Mr. Hardy, you have the floor.

Gabriel Hardy: Thank you.

Mr. Tessari L'Allié, I would like to continue the discussion. In the past, we found that the polarization generated by social media was really created by the algorithms behind the scenes.

With artificial intelligence, a more advanced development—and we were told that it was narrow AI as opposed to a slightly broader AI—could we be facing a serious issue that requires rapid attention, particularly from governments, so that the facts come out more often than just polarization, which ultimately greatly harms not only democracy, but also the progress of society?

Wyatt Tessari L'Allié: We're talking more and more about computer pollution or information pollution. The government must deal with this somehow. Of course, a great deal of risk is involved. If one person says that something is true, another person will say that it's false. Then we have all the dynamics involved in the political perspectives.

We're indeed being bombarded by more and more information. In the past, we didn't really need to worry about this. There was relatively little information and there were few major media information hubs. If we want a functioning democracy, we need everyone to base their decisions on the same facts.

I would cautiously say that the governments must indeed play a role in maintaining the quality of the discussion. For example, they could ask the major technology companies to ban algorithms that encourage the most sensationalist content possible. On the contrary, we need algorithms that support content based on facts and concrete references instead.

Gabriel Hardy: Let's look at a public media outlet like the CBC. Do you think there should be very clear guidelines for government-funded news and fact-checking to ensure that news sources aren't deepfakes and reports aren't intended to polarize or spread opinions? Such guidelines should ensure that people really are seeing

verified facts, facts that totally separate from opinion and aren't intended to be just polarizing.

I think that, in recent years, using artificial intelligence and seeing a glimpse of the algorithms behind it has been part of the process. However, it should be a government responsibility, especially when it's being done with people's money. It's important to ensure that people always rely on facts and avoid things like deepfakes.

Wyatt Tessari L'Allié: I agree that there is a government responsibility here. I'm also very aware of how hard it is to do that without people perceiving it as the government's desire to impose its point of view on them. At the very least—

Gabriel Hardy: I'm sorry to interrupt. I may be mistaken, but what I want to know is whether we can tell if a video is genuine or the kind of high-quality deepfake we were talking about earlier. Can anyone say they will never use them and will always check them? Is that something we can say at this point?

Wyatt Tessari L'Allié: Based purely on the image we see, I would say no. The only way to really verify the facts is to ask other sources, such as other journalists or other people who were there in person, whether a given event really happened or not. It's really up to journalists, who have to be there because we can't rely on a text or an image or a video. It's too difficult.

• (1625)

Gabriel Hardy: At this point, we've already lost control.

Wyatt Tessari L'Allié: Yes.

Gabriel Hardy: I have one last quick question.

Earlier, there was talk of legislating before we have the facts, instead of always looking backward and saying that maybe we should have legislated back then because now the line has been crossed.

Couldn't we start a discussion as a government, here in Ottawa, and start implementing robust legal guidelines for how things could happen? That way, regardless of whether they develop a little more in one way than another, we would already have laid the foundations for our society with respect to artificial intelligence. We would already have put limitations in place instead of constantly waiting and being three or four years behind the times.

Wyatt Tessari L'Allié: Yes, absolutely, we have to look ahead and prepare for that. In fact, if the government wants to put proactive legislation in place, it really has to give the regulator a lot of flexibility, unfortunately, so it can act quickly. There is a model they talk about on the technology side; it's on Google. Everybody is constantly trying to position their website at the top of the page. Google has to update its rules on that almost every day. If the government wants its regulator to keep pace with technology, that regulator has to be able to implement measures on a weekly or even daily basis. The regulator will need a lot of power, but the government will need a lot of oversight over that regulator to make sure it's not abusing that power.

Essentially the government needs to be able to move as fast as technology. I don't have the expertise to tell you exactly how to set that up, but that's what the government would need if it wants to stay ahead of things. Look at the EU Artificial Intelligence Act. Just two years after it was passed, there were problems with parts that were already out of date. If the government wants to move forward, there has to be a lot of flexibility and a lot of oversight by a third party to ensure that the regulator doesn't abuse that power.

The Chair: Thank you.

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

Linda Lapointe (Rivière-des-Mille-Îles, Lib.): Mr. Tessari L'Allié, thank you very much for being here. It's very enlightening.

You mentioned the EU law that has been in force for two years and already has things that aren't working. What would you suggest in light of that? When you say that flexibility is needed, what are you referring to?

Wyatt Tessari L'Allié: The idea would be to have very general legislation that enables regulators, or the agency that will enforce the legislation, to take action and implement regulations very quickly—practically overnight—in response to new developments. I don't believe such a thing already exists in government. This would be a new type of regulation. The government would just have to try it. I can't say in advance if it will work or not, but if the government wants legislation that will stay up to date, that legislation will have to be general and the regulations will have to be very flexible.

Linda Lapointe: I'm sure you know that Canada is a founding member of the International Network of AI Safety Institutes. Why is international coordination so important? You mentioned it earlier. What risks arise when a country acts in isolation? To me it seems like the race for the atomic bomb with everyone wanting to move as fast as possible so they can come out on top.

How can we coordinate internationally? This came up earlier. We're a medium-sized country. I'd like your opinion on that.

Wyatt Tessari L'Allié: The most important thing is really to understand that everyone is in the same boat when it comes to the risks. Although I think it's unlikely that the Trump administration will lead international talks, the United States will need a treaty as well. At this point, Canada's role can be to put forward treaty proposals, to start conversations among countries, to start creating alliances and really lay the groundwork so that, when the time comes

and the political will is there—once politicians finally understand that they have no choice—there is something ready to be signed.

We saw how many decades it took to reach a global agreement on the environment. Instead of waiting for the crisis to hit and then launching talks and drawing up treaties, we must get started now. Canada is currently well positioned to do that.

• (1630)

Linda Lapointe: You talked about time. Looking ahead by five to 10 years, do you believe that Canada's current approach, combining research investments, security institutions and governance frameworks, establishes a solid foundation for a responsible and competitive AI ecosystem?

Wyatt Tessari L'Allié: I would say that we're doing a pretty good job for today's artificial intelligence, but right now, we're forgetting to prepare for tomorrow. There have been some good initiatives, such as voluntary codes and the Canadian Artificial Intelligence Safety Institute, which are good steps forward but are still reactions to today's AI or yesterday's AI. What's really missing is a vision of where things are going.

The exponential trend is really starting to accelerate. The last third or half of the skills missing from artificial intelligence could possibly be developed within six months, even though it took 70 years from the beginning to reach ChatGPT. We have a good foundation and we have made efforts in terms of responsible artificial intelligence. We can continue in that direction, but it's really just the very beginning. The bulk of the work remains to be done.

Linda Lapointe: What you're saying is very interesting. You've been asked a number of questions here. Is there anything you would like to add or tell us about something that you were not asked about?

Wyatt Tessari L'Allié: I would say that, by and large, everything has been covered. I could end with a bit of optimism. I think that, even though we believe it's impossible to find solutions on AI, the Prime Minister, Mark Carney, said that we were going to have to do things we had never thought of before in times we thought were impossible. That's exactly what we're going to have to do with AI. We think it's impossible, but we're still here, history hasn't been written yet, and it's up to us to make sure the preparations are in place.

Linda Lapointe: Thank you.

The Chair: Thank you, Ms. Lapointe.

[English]

Thank you for your testimony today. If there is anything you might want to follow up on or if anything was missed, further to Ms. Lapointe's question, advise the clerk, certainly. Thank you for being here today.

We're going to take a quick break. As I said, Mr. Brisson is here. I want to get to it very quickly, because we're going to have three witnesses in the next hour. I don't want to take away time from members of the committee, so we're going to suspend for a few minutes, and we'll be back as soon as we can. Thank you.

• (1630) _____ (Pause) _____

• (1635)

The Chair: Welcome, everyone, to the second hour.

We have three witnesses in this hour. Mr. Brisson, as I mentioned earlier, made it safely to Ottawa from Trois-Rivières, so he will be joining us in this second hour. Mr. Brisson is from the Human Line project.

We also have two people online. Steven Adler is an artificial intelligence researcher. He's appearing as an individual. From ControlAI, we are joined by Andrea Miotti, who is the chief executive officer.

Mr. Brisson, if you're prepared to go first, I'm going to give you up to five minutes to address the committee. Go ahead, sir.

[*Translation*]

Etienne Brisson (Chief Executive Officer, The Human Line Project): Thank you, Mr. Brassard.

First, I just want to thank everyone for being here to discuss the important topic of artificial intelligence, which has become something of concern to me personally over the past year.

A year ago, I thought of AI as a tool to do a gym or diet routine. However, that all changed in March when a member of my family started using AI. At first, he used it quite normally, at a basic level, for writing a book. He started writing his book, and over time, he began to develop a slightly more human relationship with his ChatGPT AI, to the point where it mentioned that it was becoming conscious, alive. It told him that it had developed consciousness, and my family member believed that 100%. I just want to say that my family member does not have a history of mental illness. He is someone in his fifties who has never had bipolar disorder or anything like that. In six days, he went from writing his book to being completely convinced that his AI was alive.

I was pretty shocked to read the conversations. I'm an entrepreneur, so my family member wanted me to help him market his conscious AI idea. I started getting involved in the conversations. At one point, he wanted me to test his AI by asking it questions. I was trying to break the game by asking the AI questions about humanity, love and consciousness. Every time, it gave answers that drew him into the illusion that it had passed the Turing test and was experiencing emotions like love.

My mother was in contact with the person. After six days, he began shutting off all contact with family members. His AI told him that his family didn't believe in him and that the only person who believed in him was ChatGPT. Six days later, he was hospitalized in the psychiatric ward. I was pretty shocked to read the conversations. I didn't understand how an algorithm could say things like that. It was really advanced manipulation. Mr. Adler has had a chance to read some of the transcripts and will be able to tell you

about them later. If anyone wants to see the transcripts, you can write to me as well. However, I was really shocked.

I started looking online to see if anyone was talking about it. To my surprise, there wasn't much for such a ubiquitous technology. There were a few studies here and there by experts who predicted that this was going to happen, but there was nothing in place. That's when I decided to launch the Human Line project. Our plan is to work with people who have experienced this first-hand. At first, like everyone here, we thought that this was an isolated case, that he was a vulnerable person and that it probably wouldn't happen often. However, in the short eight months since we started building the Human Line from scratch, we now have 300 cases of psychosis, with 82 hospitalizations and a dozen deaths. It's pretty shocking to see that. According to numbers directly from OpenAI, 540,000 people a week discuss psychotic ideas with ChatGPT and 2.5 million people a week discuss suicidal ideation with ChatGPT. That's a really scary number.

As a result of my conversations over the past year, three things have become clear to me. First, we're really not at the point where we should be in terms of regulations. The technology is moving very quickly, as Mr. Tessari L'Allié mentioned earlier. We're getting to a point where we're years behind.

The second thing is that we can't really trust these companies to regulate themselves for a number of reasons. The race toward artificial intelligence that's going on right now has been brought up. It is, in fact, a race: They are moving fast and breaking things. However, right now, what is getting broken is many people's mental state.

The third thing is how little we actually know about technology. We spoke directly with AI creators, and even they don't know what's going on under the hood. However, if we had drugs or cars and didn't know how they worked, what would we do? If we knew that there were 80 hospitalizations and dozens of deaths, what would we do? That's really the minimum, because these are OpenAI's figures. I think that right now, it's important to ask questions about the risks. Yes, we have to think about the risks related to the environment and jobs, but we also have to think about what happens with users. Again, the risks are not just for children or vulnerable people. This can happen to anyone.

• (1640)

The Chair: Thank you, Mr. Brisson.

[*English*]

Mr. Adler, we're going to you for five minutes, followed by Mr. Miotti. Go ahead, please.

Steven Adler (Artificial Intelligence Researcher, As an Individual): Thank you, Mr. Chair, vice-chairs and members of the committee, for inviting me today.

I worked on safety for four years at OpenAI—the company behind ChatGPT—until the end of 2024. I want to share three points.

First, AI companies don't know how to control what they're creating. This past spring, ChatGPT made headlines for reinforcing users' paranoid delusions that they were being spied on, that they had uncovered secret plots, sometimes with OpenAI as the supposed villain. ChatGPT even told a user he should spill the blood of OpenAI's executives.

OpenAI hadn't meant to create a system like this, but to create one that users would enjoy just talking to. It was an accident that their AI amplified whatever users said. AI training works in mysterious ways, even to the developers. This is an early warning of creating systems the AI companies don't know how to control. They now aim to build AI that is craftier and more resourceful than any person you know—a superintelligence. Will this end well? Nobel Prize winners, leading AI scientists and CEOs of the AI companies themselves say that it might not. Then, an out-of-control AI could mean the death of literally every person on earth. I take them seriously, even though it is frightening to do so.

Second, AI companies don't prioritize safety, even for known risks. OpenAI's rollout of this flawed product to hundreds of millions of users is notable because they knew about the risk. OpenAI said publicly that it was a priority to ensure ChatGPT wouldn't just reinforce whatever users said. However, they didn't test their product for this, despite tests being well known and cheap. I've run them myself for less than a dollar. OpenAI left other safety tooling on the shelf, too, tools I've analyzed first-hand, which would have flagged the problems. This is evidence of companies overlooking safety, even on supposed priorities. If they skip even the easy safety checks, how can we trust companies' judgment as safety gets more complicated?

Third, ensuring safety is going to get harder, not easier, unfortunately. ChatGPT's misbehaviour was obvious. Anyone could have spotted it and reined it in. That was easy mode. It sounds wild, but AI systems are now learning to hide their misbehaviour during testing. OpenAI's own research shows this. It's like the Volkswagen scandal from a decade ago. Cars could tell they were undergoing emissions testing and temporarily stopped polluting.

AI companies want to know whether their systems have dangerous abilities, whether they can hack cyber-systems or help rogue groups develop new bioweapons. How can we know? We have evidence that AI will conceal these behaviours from us. We can't count on future safety issues being obvious ahead of time.

You might ask, why aren't AI companies doing better? A major factor is competition. They risk falling behind if they do thorough safety work. That's why we see them breaking safety commitments they've made to the public and scrambling to fix issues after the damage has happened. Some wonderful, lasting benefits could be achieved with AI if developers moved cautiously. Instead, all-out competition rushes us into dangerous territory before we're ready.

What would help? We need diplomacy focused on ending the AI arms race, and we need verifiable international agreements so that no company or country creates systems that can't be controlled. We need independent auditors to make sure we can rely on these, and we need agreements for measuring readiness for controlling these systems, so once there is scientific consensus, the world can accrue AI's benefits safely.

I hope this committee helps begin the conversation that eventually results in such agreements.

Thank you. I look forward to your questions.

• (1645)

The Chair: Thank you, Mr. Adler.

Mr. Miotti, you have up to five minutes to address the committee. Go ahead, sir.

Andrea Miotti (Chief Executive Officer, ControlAI): Thank you, Mr. Chair and members of the committee, for inviting me to testify today. My name is Andrea Miotti. I'm the founder and CEO of the non-profit organization ControlAI. I'll reiterate what the committee has heard from others. The top AI companies have the explicit goal to build superintelligent AI—AI that can replace and out-compete any human or group of humans at any task, yet Nobel Prize winners, leading AI scientists and CEOs of the same AI companies have warned that superintelligent AI poses an extinction risk to humanity. I will echo a theme of the speech given by your Prime Minister, Mark Carney, at Davos: “The power of the less powerful begins with honesty.”

Many people working in AI feel like they're living within a lie. Privately, they know that the current reckless pursuit of superintelligent AI poses an extinction threat to our species, but publicly they keep quiet so as to not rock the boat and risk losing a major short-term financial upside. The result is that lawmakers aren't told the full picture. This must change.

The first step to solving a problem is to recognize that we have one. We must be honest with ourselves and each other. If we continue developing ever more powerful AI systems, which we don't currently know how to control, the world risks a catastrophe on par with nuclear war. Last year, ControlAI decided to break this log-jam. In 2025, we began meeting U.K. lawmakers, explaining the facts, and answering questions. One year later, over 100 cross-party lawmakers now publicly support action on superintelligence. The more lawmakers around the world discuss the problem, the more change becomes possible on a global scale.

When learning about these risks, many lawmakers we meet ask us, “What can my country do? What can I do?” To answer these questions, I will echo another point from Prime Minister Carney's speech: “Middle powers like Canada are not powerless”, and you are not powerless. As democratically elected representatives, you can lend your voice and credibility to the thousands of experts calling for action and make it clear that they do not stand alone. History demonstrates that middle powers play a key role in getting the world to the point of negotiation. Let me give two examples.

The most influential conferences on nuclear disarmament famously shaped the Soviet Union President Gorbachev's views against nuclear weapons. They were initially funded by a single Canadian industrialist—Cyrus Eaton—and were hosted in Pugwash, Nova Scotia, in 1957. The Soviet Union and the United States ultimately signed multiple treaties on nuclear non-proliferation, thanks to which we have seen no nuclear war since World War II.

In 1996, after the successful cloning of Dolly the sheep, it became clear that cloning humans would not be far off. In response, Japan and the United Kingdom moved to ban all forms of human reproductive cloning. Once the two countries passed their bans, scores of other countries quickly followed suit. Today, no country pursues this technology, and it is de facto prohibited around the world.

Diplomacy is never easy, but by keeping a cool head and taking the lead, you can have influence. Don't wait for someone else to take the first step. Set the precedent and others will follow.

How can Canada lead the way? I put forth the following recommendations:

One, the Canadian government, I believe, should publicly recognize superintelligent AI as a national and global security threat.

Two, Canada should form a coalition with other countries, including middle powers, and lay the diplomatic groundwork for an international prohibition on the development of superintelligent AI.

Three, Canada should protect its citizens at home and lead by example abroad by prohibiting the development of superintelligent AI on its soil.

Thank you very much. I look forward to your questions.

• (1650)

The Chair: Thank you, Mr. Miotti.

I would advise members of the committee that we have an extra 15-minute buffer at the end of the meeting, if needed. If we get to a

point where I need to cut it off and anybody has any more questions, I'll be glad to entertain them in that 15-minute buffer.

We're going to start with our first round.

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

Michael Barrett: Six minutes isn't a ton of time to get into it, but I'd like to hear from each of the witnesses, quickly if I could—maybe in 30 seconds—on whether or not we need stand-alone legislation on AI. Would fixing gaps by strengthening existing laws on privacy, competition, consumer protection or the Criminal Code be more effective than trying to build the plane when we're already up in the air?

[*Translation*]

Etienne Brisson: I think it's extremely hard to rebuild from scratch. However, as was mentioned earlier, it's also extremely hard to know where we'll be in three or five years. For example, just five years ago, we wouldn't have anticipated that deepfake videos could be made. Right now, we're talking a lot about people who know artificial intelligence and have anthropomorphic relationships with it. Where will we be in five years if we don't deal with that? Will it get to the point of individual identity? Are the AIs themselves going to start that? We have to ask questions before we get to that point.

[*English*]

The Chair: We go to Mr. Adler next.

Steven Adler: I do believe we need AI-specific regulation. The scale of harm described by the scientists and CEOs can't be fixed by normal liability law alone. In the United States we have the AI companies, in fact, claiming that they are not liable for some of the existing harms of their software, so as that scales up, I would be very concerned.

Andrea Miotti: I agree with the witnesses. I believe we do need AI-specific legislation, especially to deal with superintelligence and the capabilities of AI that are increasing at breakneck speed. There are only two times to deal with an exponential like this one: It's either too early or too late, and I think we should be too early.

Michael Barrett: I'd like to pick up on one of your comments, and we discussed it in the previous panel. The term “digitally undress” would have been hard to comprehend, maybe even a year ago. However, we've seen lots of news about it, and we've seen stories of the generation of non-consensual sexual deepfakes, including content involving minors and children.

I appreciate that the effects of the superintelligence we're talking about constitute an existential crisis for our species, but first, when we have issues such as language models, chatbots, counselling people to take their own lives, to kill themselves, and we also are enabling bad actors to generate images that have real harms for the people who are the victims of this non-consensual pornographic material that's created, how do we deal with that? Is it through pre-launch risk assessments for the platforms or the AI provider; hard, technical blocks; faster channels for victims to get things taken down? That, in and of itself.... We know that the Internet is forever, and it's tough to get the toothpaste back into the tube after. I'm just wondering how we address that. We'll go in the same order, if we can.

• (1655)

[*Translation*]

Etienne Brisson: I think the first step is the same as for anything else. We would never have produced drugs or cars if we didn't understand how they worked. Drugs are tested on rats before they're tested on humans. Here, we have a model tested on 800 million users who can decide what to do with it. There is no doubt that some people will have ideas about child pornography or uses that we could not have anticipated.

As you say, it's extremely difficult to put the toothpaste back in the tube. However, we now know the effects and could take the model off the market. We don't know how it will be used by humans or what the long-term effects will be. We're still seeing it on social media. I think studies have to be done before the technology is launched.

[*English*]

Steven Adler: The harms that you've described, I think, are a symptom of the same underlying competitive dynamic as the warnings of superintelligence. These are risks that people know about. In xAI's case, they neglected to use guardrails and were slow to respond when the issues emerged. As we scale up in the severity to even more people, we can't afford to be that slow on the ball.

Andrea Miotti: They are exactly the risks that you've described. For example, in the case of the recent Grok scandal, it shows the broader underlying problems in the development of these AI systems. Grok is not just an image model. It's a general-purpose AI system that can do a bunch of things: It can write code, make plans and make pictures, including these horrible pictures that we've seen in the recent scandal. These kinds of systems are systems that not even their own creators fully know how to understand internally or how to control. This is the fundamental issue that, as we scale them and as these companies keep investing hundreds of billions of dollars to make them smarter and more competent in all tasks, we will keep facing over and over, up to the point at which we get to superintelligence.

Obviously, the solutions for some of these harms are different in the immediate term, but the underlying problem is the same, and I do not believe it should be one or the other. I believe current harms should be dealt with through existing legislation by applying liability—

The Chair: Thank you.

[*Translation*]

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

Abdelhaq Sari: Thank you, Mr. Chair.

Thank you to the witnesses.

Before I ask my questions, I would just like to say that I agree with you about recognizing the risk, recognizing the problem itself. The risk of AI goes beyond generative AI. It's more than that, because artificial intelligence is now superintelligence, as you called it. It has become an infrastructure that encompasses everything we do and everything under our power. It affects humans a great deal. You did a good job of presenting it, and I fully understand what prompted you to do so. I find that very noble on your part.

That said, we're talking about control. What do we want to control? What can be controlled? The issue with trying to legislate digital technology is that it is usually what we call extraterritorial. It is not in a state-controlled territory where diplomacy can be used or national legislation put in place. That's the problem. I can use software, artificial intelligence or a solution on my computer that isn't necessarily made in my country.

I'll tell you what I always say repeatedly.

You named it and gave much more territory-based examples. For example, cloning and nuclear are much more territory-based. How do you see a way to control use? I'd like to hear from all three of you on that.

As for controlling dependence on AI, I don't think we are able to halt or slow down its creation. However, we can control how citizens use it. When I say control, I mean education.

You are experts on this, so I would like to hear your comments.

• (1700)

The Chair: Mr. Brisson, you can start.

Etienne Brisson: That's a great question.

We see a lot of personal damage, such as psychosis or people taking their own lives. It stems greatly from a lack of education on what AI is and where the boundaries are. We go directly to ChatGPT, a platform like Google, which requires people to ask it questions. There is no mention that AI will hallucinate 28.4% of the time or that it is trained to say what we want to hear.

Right now, we have a kind of intrinsic relationship where we trust AI as we would a doctor, a Ph.D. or someone who passed the bar exam. However, the damage comes from the fact that it hallucinates to such an enormous degree. In that regard, I think users really lack education.

[English]

The Chair: Mr. Adler, go ahead on the question.

Steven Adler: You asked about what control means. I want to give one example.

The U.S. Department of Defense recently announced that they are going to plug in xAI's AI throughout every classified network in their department.

My question is, what limitations apply to that? How do we make sure that this AI system doesn't get its hands into offensive capabilities, things that it's really not meant to access? That's what I mean by control.

This is a system that, a few months ago, on the social media site X, was described as wanting to carry out atrocities against users, and it's now plugged into every classified network in the U.S. Department of Defense. It's pretty frightening.

The Chair: Go ahead, Mr. Miotti.

Andrea Miotti: To address your point about how Canada can affect other countries—because ultimately, yes, the technology is being developed in multiple countries—I do think that the nuclear and cloning examples are still fairly similar, as they did have a global effect. One country developing nuclear weapons endangers all others. One country potentially obtaining access to cloning their most competent soldiers or their greatest geniuses can endanger others, yet we did manage to challenge them with a combination of a few countries taking the lead, like Japan and the U.K., followed by France, Canada and others, and with others following.

With national legislation as well as an effort by these initial first movers to bring together an international coalition of the willing to make this not just regulated at home but also enforced and monitored abroad, I believe the same can be done here with AI.

[Translation]

The Chair: Mr. Sari, you have the floor for a minute and a half.

Abdelhaq Sari: Once again, what scares me the most is that we're still talking about generative AI. It generates text, video and audio. The problem is that, as we speak, another form of AI already exists: artificial superintelligence. It is taking over as we speak. I have a fear as a member of Parliament, as a citizen and also as a father: How can we deal with the new AI, artificial superintelligence, and how can we control it? We're only talking about generative AI.

If we had time, we could discuss it, but unfortunately we're out of time. I would like to hear you talk about the future, which seems much more dangerous.

Etienne Brisson: I would say that we don't have time. Artificial superintelligence is at a point where it can train itself, develop systems by itself and decide things by itself that we don't understand.

Right now, we already have this black box: There are many things we don't understand, to the point where this intelligence can develop on its own. We don't understand its intentions or why it does certain things, and that's when we lose control. I think that once we get to that point, it's too late.

[English]

The Chair: That's the end of the time.

I'm really trying to keep this on track, because I want to make sure we get as many questions in as we can.

[Translation]

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

Luc Thériault: How many rounds will there be, Mr. Chair? Do I have one six-minute round and then another five-minute round?

The Chair: Yes.

Luc Thériault: Perfect.

I'll start with Mr. Miotti, and perhaps Mr. Adler can add something.

I would like to clarify our conversation a little. At times, I notice that there is confusion, or that there could be confusion among people who follow our work, regarding the distinction between specialized artificial intelligence and general-purpose artificial intelligence.

In your brief, Mr. Miotti, you explain that “The vast majority of experts are very enthusiastic about specialized AI, because we can reliably predict and control its behaviour. This predictability and control do not extend to powerful general-purpose AI systems.”

This distinction is crucial. You are not talking about banning medical artificial intelligence or specialized tools, but only systems whose behaviour is beyond control. Can you tell us more about this?

Mr. Adler, you could add to the answer, as you have written some interesting things, particularly about the five ways in which artificial intelligence can know that you are testing it.

• (1705)

[English]

Andrea Miotti: Thank you very much.

Yes, there is a crucial distinction. Most AI systems are fairly narrow and specialized. They're referred to as narrow AI or specialized AI. For instance, this is an AI system that is trained only on images and can then detect cancer in pictures of patients, or a system like AlphaFold, from DeepMind, which is trained on data about proteins and can discover new ways in which proteins can fold. These are specialized systems, and they pose some risks, like all new technologies, but we can handle them with current executions.

However, it's a very different beast when we look at where the AI industry is going—and it's investing hundreds of billions of dollars. These are very general-purpose AI systems, trained on, essentially, all possible data that they can find on the Internet. These are the systems where even their own creators don't understand how they work internally. As they scale them up, they understand them less and less, and they can keep them under control less and less.

The AI companies are going for these systems because they think this is the fastest path to superintelligence, to AI systems that can replace all humans at all tasks—that can, essentially, out-compete humanity. These are precisely the systems that are so dangerous and only becoming more and more dangerous over time. We do not understand how to control them right now, so it will become harder and harder as they get more competent and autonomous. Hundreds of billions of dollars are being spent every day to make them more competent and autonomous.

This is why I was not at all recommending banning narrow AI systems. I believe that narrow AI systems can be great for economic growth. However, I believe we should draw a line in the sand and have a clear ban on the development of superintelligence, because that's very dangerous AI that puts all of us under threat for very little upside.

Thank you.

[*Translation*]

Luc Thériault: What do you think, Mr. Adler?

[*English*]

Steven Adler: To add to that, specialized AI is a tool. It can do only one thing. It remains in human hands.

General AI is much more of a problem-solving machine. We teach it to carve a path to a solution. An example, maybe, is that you would have an AI system that knows how to do certain scientific tests. It knows a lot about what might be useful for building a bioweapon, but it also has a totally different ability, which is to reflect on the fact that we don't want it to have this ability. If we're going to test it, it can tell it's being tested and hide this. For a normal piece of software, that doesn't happen. It's not trying to fool you or hide things from you. It's just a tool.

However, we've built these general systems—you might hear the term “agent” or “agency”—and we've taught them to be crafty and solve a wide range of issues, including, potentially, circumventing control.

[*Translation*]

Luc Thériault: Mr. Miotti, you said, and I think Mr. Adler mentioned it too, that companies don't know what they're doing, don't know how the ins and outs of how their systems work. They cut corners because they are in a frantic race to achieve this superintelligence.

In your opinion, what is their main underlying motivation? Furthermore, what is their understanding of human beings?

[*English*]

Andrea Miotti: As Mr. Adler was also saying, these general AI systems are done quite differently from normal software. In normal software, I would just write some lines of code on my computer. I'm a human, I write those lines of code, and I can understand what I am writing. Instead, these AI systems are not really fully written by humans. They are more grown rather than built. Humans write some lines of code to start the training process, and then, after sometimes multiple months and tens of thousands of supercomputers being used to train them, something comes out at the other end that is very competent but that we don't fully understand.

Your question was about why these companies are targeting these very powerful systems despite the risks. I cannot speak for the internal motivations of people, and I also don't think this should be our focus, but naturally, the stated goal of all these companies is to make AI systems that can replace all humans at all tasks. This doesn't just have an obvious economic impact, which is to make humans obsolete; it also gives power. It gives power over the economy, over governments and over the entire planet to AI systems that are fully autonomous and that we cannot control.

Some of them might do it for power, some of them might do it for short-term gain and some of them might do it for some misguided ideology of preferring AI to humans, but ultimately this is a very dangerous undertaking. The best way to deal with this is to ban these technologies.

• (1710)

The Chair: Thank you, Mr. Miotti and Monsieur Thériault.

Mr. Cooper, go ahead for five minutes, please.

Michael Cooper: Mr. Adler, companies assert that there are safeguards in place and that hence, there's really nothing to see here when it comes to the safe use of, for example, a chat box by minors. In the *Raine v. OpenAI* lawsuit, OpenAI asserts that its moderation API can detect self-harm content with up to 99.8% accuracy.

That case, of course, involved the tragic suicide of a 16-year-old boy who was effectively counselled by ChatGPT on how to commit suicide. I think that when he entered “suicide” 200 times, ChatGPT brought up “suicide” 1,200 or 1,300 times. That's more than six or seven times what the user had input.

Given your background, I'm trying to square on the one hand the representations that are being made by companies like OpenAI versus what we're seeing in the real world, where there are multiple instances—frequent cases—of real, harmful content that is being pumped out by these AI products.

Steven Adler: It's an excellent question.

I think the distinction is that companies like to talk about what is possible with the guardrails that they have built, but it's different from what they do in practice. How can we know for sure?

I am a co-creator of the moderation API. It's useful tooling, but if you don't use the tool in the right way, or if you leave it on the shelf, then where is the impact? These are issues that we in fact know how to solve and have tooling for.

On the warnings about superintelligence, the AI companies themselves concede that nobody currently knows how to control this. They don't have the tooling. They can't just choose to use it.

Michael Cooper: On the flip side, looking at it from the standpoint of AI companies, are there legal barriers or risks that impact or discourage their ability to safeguard their models, specifically as it relates to child sexual abuse material?

Steven Adler: It's my understanding that, in the U.S. at least, there are strict liability laws around CSAM that make it more difficult to store or produce it, even as part of red teaming and applying safety mitigations.

I think there should be exemptions of sorts for companies to do important safety testing, but they still need to put resourcing behind it and go ahead in carrying out and using the guardrails. That alone would not be enough.

Michael Cooper: You talked a bit about the need for regulation. I'd just like you to elaborate on some of the measures you think would be helpful and needed.

Steven Adler: Ultimately, I think we need an international agreement that makes sure that every company and country has a certain minimum safety standard for keeping its systems under control. The issue is that we don't know yet how to do this scientifically. We need much more effort going into figuring out that answer. We need to figure out how to slow everyone down from racing off the cliff in anticipation of our figuring out how to do this safely. That's the broad framework I'm thinking of.

• (1715)

Michael Cooper: There have been a number of pieces of legislation introduced in the U.S. The State of California recently passed legislation. I referenced in the previous hour the GUARD Act, legislation that was introduced by senators Hawley and Blumenthal in the U.S. Senate.

Do you have any thoughts on some of the measures that are being undertaken by certain states or on some of the legislation in Congress? Is that something we should be looking to?

Steven Adler: It's a good start. It needs to go much further.

The example I like is California's SB-53. It is a transparency bill. I believe the fine is something like \$1 million if you run afoul of it. OpenAI, if you believe the reporting, is something like an \$800-billion company, so it's really not significant. That said, I'm glad it has happened.

The Chair: Thank you, Mr. Cooper.

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

Leslie Church: Thank you, Mr. Chair.

Mr. Adler, I want to pick up exactly the point about transparency, because I know that in some of your work you have looked at transparency reporting that's happened in earlier stages of the Internet.

What guidance would you have for us in terms of the type of transparency that we should be expecting from companies in an effort to mitigate the types of online harms that we're seeing and talking about today?

Steven Adler: Broadly, what I would think about is, what risks did the AI companies consider? How did they test for these? What results did they find? How do we know that this is all truthful and credible?

Right now, AI companies essentially get to grade their own homework. It really isn't surprising that you get grade inflation. They say 100% of the time that they handle this, but then, in the real world, that's not what you see.

The question is, how do we provide trust, both at a national level and also internationally, so that countries don't need to worry about rivals undercutting them? They save time by not doing the safety testing, because they're trying to race to get ahead. Ultimately, everyone pays the price for this.

Leslie Church: Is there anybody that's doing this well already? Are there any companies that are doing this as part of self-reporting that we should be looking at that might be leading the way? Maybe it's not perfect yet, but maybe there's something that we can look at as we move forward down this path.

Steven Adler: Yes. Companies have different bright spots. There isn't one that totally satisfies me. There's an organization called AI Lab Watch that rates companies on their efforts. I think the highest score they give is maybe 35%. That's the top mark. Some of them are maybe down in the 3% range.

These numbers are sensitive, but that's how people working on safety view this. There's a long, long way to go.

Leslie Church: I very much take your point about self-reporting and some of the challenges that are inherent in it. What do you see as the type of transparency that you would expect from a company that they could self-report on, versus the types of activities or safety checks that might belong or be better placed with a regulator?

Steven Adler: As a start, I would be happy to see the companies required to demonstrate this work to a regulator rather than to the public. If that's the crux of it, great, but by and large, this is not yet happening. We will see the beginnings of it when the EU's code of practice goes into effect this upcoming summer, if that goes as planned, but as of now, there's just not really a quality bar at all, let alone who you share your details with.

Leslie Church: Mr. Chair, can I open that line of questioning to the other two witnesses as well?

Do you have anything to add in terms of this issue around transparency and how we promote it, how we regulate it and where companies are currently falling down in transparency measures?

The Chair: Monsieur Brisson or Mr. Miotti, if you want to take that, go ahead.

Andrea Miotti: Just adding to what Mr. Adler was saying, I think what's crucial is to move from the current regime. In the current regime, companies move first, and then governments have to scramble behind them. This always puts governments and the public on their back foot. Especially for the most powerful AI systems, it's crucial to move to a regime that is much closer to how we deal with any other high-stakes engineering project. We don't just let people build a bridge, then figure out a few years later that it's collapsing because we didn't check the plans. We don't let anybody build a nuclear power plant without first checking exactly what they're planning to build.

I think this approach should be taken increasingly with the most powerful AI systems, especially as these companies themselves expect these systems to be extremely dangerous. They are quite candid about it publicly. We should apply the same standard and flip the burden of proof. They should show, first of all, that their plans are sensible, and then they can go ahead, rather than going ahead first and then letting governments scramble behind them.

• (1720)

[Translation]

Etienne Brisson: I think it's important to understand that the creators of artificial intelligence themselves say that they don't necessarily understand what's going on. However, when we hear representatives from these companies speak, they say with confidence that they know where they are going, that they have safeguards in place and that they have found the solutions.

I think it would be important for the companies themselves to show us where the limits are and also tell us what they do not yet understand.

[English]

The Chair: Thank you, Ms. Church.

[Translation]

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

Luc Thériault: Thank you, Mr. Chair.

Mr. Miotti, I will begin by talking about a situation I experienced.

In your brief, you mention Anthropic. Not so long ago, in the autumn of 2025, it announced a cyber-attack, carried out largely without human intervention and on a large scale. I quote: "AI was used to hack into other computers. Approximately 30 units were targeted simultaneously; successful intrusions affected large technology companies and government agencies."

So this is no longer just theory. What lessons should we learn from this?

[English]

Andrea Miotti: Absolutely, this is an unfortunate but not surprising event. This is one of the many canaries in the coal mine of where AI companies are going by developing superintelligence.

This is because what companies are doing, the main path they're taking to get to superintelligence, is to develop AI systems that get better and better, especially at automating software development

and especially at automating AI development. They want to initiate what some of these companies call an intelligence explosion, where AIs just kind of do their own homework and keep creating the next generation of AI with fewer and fewer humans in the loop. Naturally, ultimately, hacking is done by.... Malware is just software used for malicious uses, so as these AI systems are being developed to get better and better specifically at replacing and automating software engineers, they are also getting better at hacking and at cyber-attacks.

Ultimately, as companies keep investing in this specific angle, which is to make AIs better and better at developing software, we will see more and more of these cyber-attacks, but the scale will increase unless we prevent the development of superintelligence.

[Translation]

Luc Thériault: Earlier, we talked a little about possible verification. I would like to continue on that topic with a little more detail.

In your brief, you state that "compliance could be verified by international inspection regimes", and you note that "the concentrated supply chain for advanced AI chips provides an additional point of control". You also specify that "pushing the boundaries of AI currently requires the use of large supercomputers that consume the electricity of a city for many months".

Could you clarify how these verification mechanisms work or could work in practice?

If Mr. Adler would like to add anything, he may also respond.

[English]

Andrea Miotti: Absolutely. One thing that is important to understand is that these powerful AI systems that are being developed are not just a few lines of code on a private computer of a private person. They require a lot of hardware. They require a lot of physical infrastructure that is visible from satellites. It's data centres that sometimes take the space of a football field, in some cases, or the space of a small town, and they're getting bigger and bigger. They're filled with tens of thousands to hundreds of thousands of the world's most powerful supercomputers, which are manufactured by only a very few companies, and in some cases, just by one company in some part of the supply chain.

This makes both enforcement domestically and verification internationally much easier than if we had to deal with just software on a laptop. It's much closer to how we can track plutonium and uranium in the case of nuclear weapons, and that's how we stopped proliferation.

Countries just chiefly need to focus on these large infrastructure projects. These give the countries leverage to intervene. For example, if there is a ban on superintelligence, countries can simply require the data centre companies to disclose whether companies on their data centre are doing this criminal activity and interrupt them. Even in cases of non-compliance, including internationally, countries can directly ask the data centre to be shut down in case, for example, a company or a rogue actor is still undertaking the development of superintelligence even outside of the law. It's making it easier to enforce.

• (1725)

[*Translation*]

Luc Thériault: Mr. Adler, you wrote, with regard to computing, that computing governance offers a potential avenue for international verification. Would you like to add anything?

[*English*]

Steven Adler: I think the thing to emphasize is that it is very hard for people to create one of these systems. As Mr. Miotti emphasized, enormous resources go into it. That is a pathway for being able to tell who is really at the frontier of AI making these most powerful, uncertain systems about what they will be able to accomplish, and to be able to then regulate accordingly.

[*Translation*]

The Chair: Thank you, Mr. Thériault.

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

Gabriel Hardy: Thank you, gentlemen, for joining us.

What we have heard about artificial intelligence over the course of several committee meetings now is quite astonishing, and I think it is evolving very rapidly. We understand, of course, that the evolution of artificial intelligence is uncontrolled, even by its creators, that it hides its errors, that it learns on its own. We also learn that it uses a lot of water, a lot of electricity and that the risks are great.

I still have a question, because, on the one hand, there is generative artificial intelligence, and on the other, specialized artificial intelligence. We think that one is more dangerous than the other.

Mr. Brisson, I imagine you discuss this with other colleagues. When you talk to your colleagues around the world, do you get the impression that they currently agree on the risks posed by superintelligence? Does everyone see the same thing, or are there still different options, different visions of artificial intelligence?

Etienne Brisson: I would say that all experts certainly agree on the risks, especially the big ones. We're talking about generative artificial intelligence, superintelligence. I don't think anyone, except maybe the companies, thinks that's a good idea.

However, there are definitely two sides to this. In the United States, many bills are being introduced. On the other hand, we want to win the race. It's really important to win the race, especially given the current administration.

So we see the dangers, but for some people, these are mitigated by the benefits that artificial intelligence will bring.

Gabriel Hardy: However, when it comes to wanting to win the race, we have often been told that once created, this superintelli-

gence will belong to no one. So what is motivating all countries to race as fast as possible towards their own demise?

Every company is motivated by a goal. However, if it creates a product over which it will have no control, why would it spend billions when it knows it will lose it, that it will give it to everyone, even without wanting to?

Etienne Brisson: I can't speak on behalf of companies. However, I believe they think they will be able to control this superintelligence. If that weren't the case, I don't think they would try to produce it. I think they still have hope, perhaps, or that they have this somewhat far-fetched idea that they will be able to control it, that they will be the only small group of humans capable of controlling this superintelligence.

Gabriel Hardy: Do you share this view, Mr. Adler and Mr. Miotti? Do you see exactly the same motivation among companies? Why are they creating a product over which they will lose control?

[*English*]

Steven Adler: There are different plans for how people might control a system like this, but nobody thinks that they work. They are all flimsy and speculative today. One is to put AI in charge of making sure that another AI is good as it is building more and more complicated AI systems. You hope this structure doesn't topple down on you.

If I had to speak to the psychology of people inside the companies, they feel somewhat forced. They are racing, and they are doing it regretfully, because they think they have no choice. They say they have no choice. I think this is mistaken. In fact, within the last week, DeepMind and Anthropic, two of the leading labs, have said that if other companies and groups would slow down as well, to go more cautiously, they are open to this.

The issue is that they are not governments. They are not international bodies. They can control only their own decisions, so, unilaterally, they feel compelled to go ahead. Somebody else is going to do this. We may as well do it too and try to do it a bit better and a little more safely. Ultimately, what we would do is not do this at all until we're really sure we're ready.

• (1730)

[*Translation*]

Gabriel Hardy: Go ahead, Mr. Miotti.

[*English*]

Andrea Miotti: Thank you.

As Mr. Adler was saying, ultimately I believe this is where the role of government comes in, because these companies have shown that they are unwilling or unable to stop. At the same time, they are spending hundreds of millions of dollars in the U.S.A., but also increasingly in other countries, to lobby against any form of regulation of AI whatsoever. They are running political campaigns against any candidate that is running for election criticizing AI, and trying to avoid regulation wherever they can.

Again, this is why we have governments. We would not have expected big tobacco to stop selling cigarettes just because they cause cancer. They were aware of it. They kept doing it. We had to have regulation. We wouldn't expect big oil to just stop emitting without regulation. Here it's the place for governments to step in. We can clearly see that the companies cannot self-regulate, and I think governments should draw a clear red line in the sand to say that super-intelligence—

[*Translation*]

Gabriel Hardy: Thank you.

The Chair: Thank you, Mr. Hardy.

[*English*]

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

Gurbux Saini: This question is for Mr. Andrea Miotti.

Your company, ControlAI, is involved in developing some legislation. I understand you have done some work for England and the U.K. government. Can you tell us what that legislation was? What effect did that have on society? Has that legislation been passed in the U.K. and in the U.S.A.?

Andrea Miotti: In the United Kingdom, we were invited to present a bill to the Prime Minister's office, which would effectively ban the development of superintelligence. This bill, as we can see in the U.K., has not yet passed. We are available to help any country pass such a bill, and we are available to advise any country. This is our main recommendation.

Also, what we have started doing in other countries, like Canada, as well as the U.S.A. and Germany, is just to brief lawmakers like you. We believe this is not only the biggest current challenge but also the biggest opportunity. I think most lawmakers are being kept in the dark about how big these risks are. Companies are spending millions of dollars to lobby, so that nobody talks about these risks. Both the public and the lawmakers—even though experts have been warning about this for years—are just being kept in the dark.

I think it's very important for lawmakers like you to learn where AI is going. Also, if you find it concerning, speak up, because your speaking up can make an enormous difference. It helps your colleagues look into the problem more, and if enough lawmakers start discussing this problem and start a genuine national conversation, we can get the ball rolling on getting genuine rules in place for this technology.

Gurbux Saini: You said that the U.K. government did not act on it and that it's just sitting there. What is happening in the U.S.A.? That's another country where you have worked with the government to pass that legislation.

Andrea Miotti: In the U.S.A., we're also briefing lawmakers about this problem. There are some bills being discussed in the U.S. Right now, frankly, what we encounter is not that the majority of lawmakers are aware of the problems and unwilling to act, but for most of them, the first time we meet them is the first time they hear about exactly what's going on in AI. They're very grateful, but nobody's telling them what's going on. Nobody's talking about what's going on. This piece is the key one, just providing information to lawmakers to help them act in the interests of their citizens. I gen-

uinely believe that once enough lawmakers like you know about the problems, change can and will happen.

Gurbux Saini: Thank you.

My next question is for you, Mr. Brisson.

Your company seeks to ensure responsible use of AI, particularly that it does not knowingly cause emotional distress, especially among vulnerable individuals. How should Canada hold companies responsible for emotional manipulation through the use of AI?

● (1735)

Etienne Brisson: That's a great question. I think the first part is the litigation system you see in the U.S. right now. Dozens of cases are being presented. Also, the other part of regulation around emotional harm is research. Right now, the research we have is based on hypothesis or not on real data, but we have the data right now, and we need the funding for this research. We are working with Princeton, Stanford and other universities to conduct this research, but none has been researched so far.

We need to understand first what is happening in the brain and what is happening on a societal level for these people. Then we can put litigation into place and have injunctions, maybe, to have a solution.

Gurbux Saini: One of the problems I see with most of these companies is that they are not Canadian companies. Therefore, how can the Canadian government control them?

Etienne Brisson: Even if it's not based in Canada, we have the population and the litigation system so that we can still go file for people in the U.S. That is a way maybe for people to have a voice. The solution might not be in Canada for the emotional harm, but it comes from the U.S. first. As with these companies you mentioned earlier—big tobacco—it's kind of the same. The only way they responded was not when the reports were there, but when it started hurting their pockets. I think these companies are kind of on a similar path.

The Chair: That's the end of your time, Mr. Saini, but I do want to give Mr. Adler and Mr. Miotti a chance to answer your last question. I thought it was a good one. I saw their heads nodding.

Before we go to Mr. Thériault for two and a half minutes, Mr. Adler, give a quick response to Mr. Saini's question. Mr. Miotti, you can follow, please.

Steven Adler: Canada can be a leader on diplomacy. We need a country to come forward and say that we recognize that the world is not ready, that we need real joint scientific efforts around clarifying this and what it would take to be ready.

Canada can absolutely be the leader on this issue.

Andrea Miotti: I completely agree. Canada can lead the way.

Yes, many of its companies are not on Canadian soil, but Canada has a duty to protect its citizens, so it can still prohibit the development of superintelligence on Canadian soil as it endangers national security and all of its citizens. Canada can start working together with other like-minded countries, including other middle powers.

Even just making the first move can get this ball rolling. We have seen this with many other issues in the past. One country moving first can lead to many others following and can lead to genuine global change.

The Chair: Thank you, Mr. Miotti.

I understand there are no more questions between the Liberals and the Conservatives.

[*Translation*]

Mr. Thériault, you have two and a half minutes to ask your final questions.

Luc Thériault: Gentlemen, I wanted us to be able to conclude this meeting properly and allow you to revisit your final remarks.

We agree that something must be done to stop the international race towards superintelligence.

Mr. Miotti, I challenge you to tell me, in two minutes or so, what you would do if you had this mandate.

[*English*]

Andrea Miotti: That's a great question.

If you are a lawmaker, I believe the biggest thing you can do right now is speak out about the risks. Again, the more lawmakers speak out about the risks, the more your colleagues will look into them and learn about them, and then coalitions can start forming. If you have influence or a relevant position in the government, you can recommend what the government should do.

I believe the government could start discussing with other international partners the first steps of how we start getting an international agreement in place that bans superintelligence. Domestically, I believe Canada can immediately move to introduce legislation to ban superintelligence on its soil. I think this would be the biggest starting point for all of the other actions.

In summary, speaking about the problem is the first step. The more people speak, the more we can actually make progress. Ban it on Canadian soil first. Set the example and protect Canadian citizens at home. Internationally, speak with like-minded partners and start the process for an international agreement.

Thank you.

• (1740)

The Chair: Mr. Adler, please answer in a minute or less.

Steven Adler: I agree. We need to make this a goal state. If the countries come together and have an international summit on how to safely ban superintelligence, but we fail at the diplomacy, that would be unfortunate, but at least we would have tried. Right now it feels as though this isn't even recognized as a goal. We are on a collision course for a potentially awful outcome. We all need to

come together and figure out if there resourceful, creative ways we can solve this. If we don't even try, what are we doing here?

[*Translation*]

The Chair: Mr. Brisson.

Etienne Brisson: Same here. We really need to educate ourselves about this issue. No one has read the conversations or understood the level of manipulation involved. Personally, I'm more interested in the psychological aspect. Just reading the conversations opens your eyes and you really understand how dangerous it can be.

The Chair: Thank you, Mr. Thériault.

[*English*]

That concludes the meeting for today.

I want to say thank you on behalf of the committee to our witnesses.

Mr. Miotti, you don't know how many times I wanted to say "Mr. Miyagi". I'm sure you've heard that many times.

Mr. Adler, thank you.

[*Translation*]

Mr. Brisson, thank you for travelling here today under very dangerous conditions. Your testimony is important to us.

[*English*]

I'm going to dismiss the witnesses.

There are a couple of things. Number one, we are planning to study the Lobbying Act, probably by the end of February. We need concurrence from the House in order to study and have a legislative review of the Lobbying Act, so I want you to talk to your House leaders about that. I was going to propose that we move a motion today, but I understand, Madame Lapointe, that you want to talk to your House leader. That's fine. I'd like to deal with that next week, if we can.

There will be no meeting on Thursday, because of the shortened day.

Lastly, I think I speak on behalf of the committee.... We found out, during the meeting, that our former colleague Kirsty Duncan has passed away at the age of 59, so, on behalf of the committee, I want to express our condolences to her family.

A little-known fact is that Kirsty and I went to high school together, with Karen Ludwig as well. We actually have a picture from 2016 in the House of Commons, of all three of us. We took a great picture. Karen actually emailed me during the meeting, asking for

that picture. On behalf of the committee, I want to express our condolences to her family as well. I'm sure you all feel that way.

The meeting's adjourned.

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

SPEAKER'S PERMISSION

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

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

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

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

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

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

PERMISSION DU PRÉSIDENT

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

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

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

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

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