HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

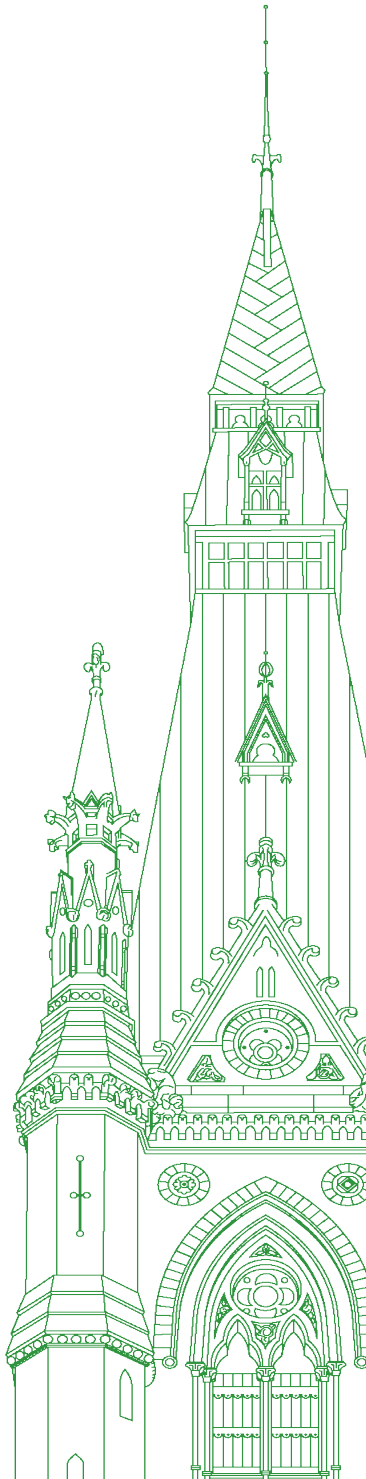
45th PARLIAMENT, 1st SESSION

Standing Committee on Procedure and House Affairs

EVIDENCE

NUMBER 045

Thursday, October 29, 2026



Chair: Chris Bittle

Standing Committee on Procedure and House Affairs

Thursday, October 29, 2026

● (1100)

[*English*]

The Chair (Chris Bittle (St. Catharines, Lib.)): I call this meeting to order. Welcome to meeting number 45 of the House of Commons Standing Committee on Procedure and House Affairs.

Pursuant to Standing Order 108(3), the committee is meeting to resume its study of artificial intelligence in Canada's democratic institutions. Today's meeting is taking place in public in a hybrid format, pursuant to the Standing Orders.

I would ask all in-person participants to consult the guidelines written on the cards on the table. These measures are in place to help prevent audio and feedback incidents and to protect the health and safety of all participants, especially our interpreters.

I'd like to remind the witnesses that the committee members may ask questions in either English or French. If you need interpretation, please take a moment now to prepare your earpiece and to select the listening channel you need in advance in order to take full advantage of the time allotted.

I have a few comments for the benefit of members. I remind you that all comments should be addressed through the chair. Whether you're in the room or on Zoom, please raise your hand. The clerk and I will do our best to maintain a speaking list.

I'd like to welcome our witnesses for today's meeting. Appearing as an individual, we have Dr. Amir-Hossein Karimi, assistant professor at the University of Waterloo. From the Inter-Parliamentary Union, we have Dr. Andy Williamson, senior researcher, Centre for Innovation in Parliament, by video conference; and from QuietWire, we have Chris Beall, president.

Witnesses, you each will have five minutes to deliver opening remarks.

We'll start with Mr. Beall for five minutes, please.

Chris Beall (President, QuietWire): Thank you very much, and thanks to the committee for the invitation.

I'm appearing as president of QuietWire, a Canadian civic AI company that's building infrastructure to enable AI tools to work as part of an institution's own team.

Today I'd like to talk about proposing ground rules for AI within democratic institutions and Parliament itself. These are standards every supplier, including QuietWire, should have to meet. Our work, I should explain, is aimed at meeting and supporting those standards.

Imagine that you have AI helping staff examine evidence, carry work forward and bring colleagues up to speed. In this model, the institution, not the tool itself, should determine what the tool accesses, remembers, shares and does. Our ambition is an AI capability that's genuinely yours—working alongside your people, under your authority and within your public responsibilities.

I should stress that in the example I provided and in the conversation I'm having, we're talking about tools that respond to instructions. That's not always the same as being under your control, and I think that's an important element to consider.

Without appropriate safeguards, staff can disclose sensitive information through unapproved tools. We can mistake generated content for verified evidence or give connected AI tools permission to take unintended actions. The level of risk, obviously, depends on the service, the terms and the configuration of the tools you're using. I should say that none of these things I'm talking about are an argument against global tools, but it's an argument about governed use and how we govern the use of what you're using.

In terms of thinking about governance, I suggest three ground rules.

First, sovereignty has to mean more than simply Canadian infrastructure and jurisdiction. Institutions also should control access, retention and reuse. They should be able to withdraw permissions and change suppliers without losing essential records or capabilities. These boundaries should be enforced through system design and tested controls, supported by but not substituted for contracts, policy and oversight.

As a second ground rule, institutional memory should support continuity as well as accountability. For consequential work, we need a retrievable record of purpose, source, what AI proposed, what people checked, who decided what and why, under what authority and what actions or commitments followed. When staff change or a new committee takes up work, authorized successors should be able to recover its purpose, sources, decisions, corrections and unfinished work. That's more than just creating an archive. It's the ability to continue work without starting from scratch to understand the context and the decisions being made. The AI systems and operations should help, not hinder, that work.

Access and delegated powers still need review when responsibilities change. It still needs to be under the existing rules. Privacy, correction and retention rules remain essential. The goal isn't simply unfettered sharing. It's allowing continuity under current, accountable authority and allowing those authorities to change.

Third, human oversight is not just a human in the loop. There has to be essential, meaningful oversight included. A person who lacks time, evidence or authority to challenge AI recommendations is not meaningfully in control. Parliamentarians need the skills and access to independent technical advice to test supporters' claims. People must be able to inspect sources, correct errors, withhold approval and stop an action. Where an AI-assisted decision affects people, there should be an appropriate route to challenge that decision and understand it.

Imagine a committee using AI to examine thousands of public submissions. That tool could help identify themes, but staff should be able to check those themes against submissions, identify perspectives that were missed, correct the analysis and then have a record of how those corrections existed.

Permission to read and summarize should not itself be a way to confer permission to send material elsewhere—either back to the manufacturer or back to the company that owns the AI—or to publish conclusions.

I'll finish with three practical recommendations.

First, make control over access, information sharing and action a testable requirement for institutional AI.

Second, require proportionate records of consequential AI-assisted decisions, preserving purpose, evidence, uncertainty, human responsibility and correction.

Third, test continuity and supplier exit before adoption. Before you put something in place, you should know that you can leave and take your material, your data and your ideas with you. Can an authorized successor in your system recover the work you did? Can the institution change tools without losing records?

These should be, of course, respectful of privacy, confidentiality and privilege, and build on existing safeguards.

As a final thought, from our perspective, it's not simply whether we should use AI. AI is in use, and AI is being operated. It's whether it can be used to strengthen capabilities without surrendering the knowledge, judgment and authority that democratic responsibilities require.

Thank you.

● (1105)

The Chair: Thank you so much.

We'll turn to Dr. Karimi.

You have five minutes, please.

Dr. Amir-Hossein Karimi (Assistant Professor, Department of Electrical and Computer Engineering, University of Waterloo, As an Individual): Thank you, Mr. Chair and members of the committee, for your time and attention. Today, I would like to speak about exponentials and the danger of responding to them with linear thinking.

We humans have an intuition for problems that grow by addition: a few more applications, a few more documents and perhaps another person to help. When workload keeps doubling, that intuition fails us. A response that seems adequate today can become inadequate remarkably quickly. Hiring more people or processing each document faster may buy us some time, but if demand grows exponentially while our capacity grows incrementally, that gap keeps widening. Therefore, if there is one thing you take away from my opening statement today, it is that playing linear games in an exponentially changing world is a recipe for failure.

My name is Amir-Hossein Karimi. I'm an assistant professor at the University of Waterloo, where I hold the Val O'Donovan chair in trustworthy AI and I lead the collaborative human-AI reasoning machines—or CHARM—lab. I'm also a faculty affiliate at the Vector Institute in Toronto.

Our lab organizes its research around understanding AI systems, intervening when something needs to change and evolving those systems through interactions with people. In the remainder of my time, allow me to share three perspectives on what that means for Parliament.

First, prepare for the gap between generation and verification.

In research, producing a convincing manuscript and establishing whether its claims are correct are different tasks. AI can accelerate the production of text, code and experiments. Checking the resulting work, however, still requires evidence, expertise and time.

Several AI conferences in the past year alone have seen submissions double, and in response to these growing submission volumes, many conferences and journals have adopted stricter screening policies, thus often rejecting up to half of the submissions before they are sent off to external review. If submissions grow exponentially, rejecting half reduces the immediate burden while leaving the underlying growth rate unchanged. The percentage reduction buys us time. It does not change the underlying trajectory.

Parliament should also prepare for the possibility that incoming volume grows much faster than its capacity to assess it. Your constituents, advocacy groups, businesses and foreign actors can and will use these tools. A brief that once required days of preparation could become one of thousands of individually tailored submissions. Some will contain valuable contributions while others will contain errors or hallucinated evidence.

The response cannot simply be to read faster. Parliament needs ways to establish provenance—that is, where material comes from—check evidence, identify coordinated submissions and use AI to assist evaluation while subjecting those evaluation tools themselves to scrutiny.

Second, preserve independent expertise.

As a professor, I see a difficult incentive emerging. An experienced researcher may find it increasingly efficient to direct AI systems personally rather than spend time training a student, but training is how we produce the next generation of people capable of performing the very task of verification. If we increase research output while weakening that apprenticeship, we could undermine the very expertise needed to evaluate what we produce. This is a dangerous combination of potentially exponential growth in output alongside an erosion of our capacity for independent judgment.

Parliament faces a related challenge. Staff need opportunities to develop judgment, question assumptions and understand an issue deeply. Receiving a polished answer is not the same as acquiring the ability to assess it. If the same small set of companies supply the systems that generate our arguments, summarize our evidence and evaluate our conclusions, their influence becomes intellectual as well as economic. Parliament should, therefore, retain the expertise and practical ability to question those systems, compare providers and reach independent conclusions.

Third, establish meaningful recourse before deployment.

As AI becomes involved in more consequential decisions, people need an answer to a basic question: What can I do when something goes wrong? If a system misrepresents a constituent's concern, overlooks relevant evidence and contributes to adverse decisions, there must be a practical route to correction.

Before deployment, we should know where responsibility lies, how an affected person can challenge an error and who has the authority to change the outcome. Where an outcome depends on a person's circumstances, recourse can also mean explaining what feasible actions would lead to a different decision. Therefore, recourse should be first and deployment next. A correction mechanism should be part of a system's design rather than something improvised after the individual encounters harm.

These mechanisms, of course, must work at the scale of deployment. If the volume of decisions grows exponentially while our capacity to reconsider them barely changes, a formal right to challenge a decision could become practically impossible to exercise.

My central recommendation, therefore, is to build evaluation capacity and independent expertise alongside accelerating adoption and establish meaningful recourse before consequential deployment. The measure of success should be whether Parliament can serve people better while remaining capable of explaining, questioning and correcting its decisions. This requires thinking critically about the scale of change and preparing for it, not simply doing a little more of what we do today.

Thank you very much.

● (1110)

The Chair: Thank you so much. Both witnesses came in under five minutes. The chair appreciates that.

The challenge is yours, Dr. Williamson, for five minutes, please.

Dr. Andy Williamson (Senior Researcher, Centre for Innovation in Parliament, Inter-Parliamentary Union): Thank you, Chair. I will do my best.

Thank you, members of the committee, for the invitation to appear.

I'm senior researcher at the Inter-Parliamentary Union's Centre for Innovation in Parliament. I'm the author of the IPU's “World e-Parliament Report 2024” and the AI maturity framework, and lead editor and co-author of the IPU guidelines for AI. I'll draw on these publications, our wider research and our global communities of practice in my remarks today about what makes a parliament AI-ready.

By way of my opening, I'd like to make three points.

First, AI is a transformative technology, and it's already being widely used in parliaments. Eighty-eight per cent of the parliaments we surveyed for the 2026 world e-parliament report say members and staff use AI, yet fewer parliaments have adopted AI at an institutional level. Barely more than half have carried out any kind of AI compliance or risk assessment, and only 24% have a formal AI policy. The technology is arriving faster than institutions can govern it, and this gap is what we call shadow AI: people using tools without approval, training or risk assessment.

AI changes how parliamentary work can be done. It can help members manage constituent casework, help committees analyze the thousands of public submissions and make proceedings searchable and accessible.

The AI use that we see today is primarily for those procedural tasks such as transcription and translation. Closer to the legislative core, tasks such as drafting and amendment analysis are often at the pilot stage, and only a few parliaments are using them in a live production environment. This is the area where the democratic stakes are the highest.

Second, AI is part of a wider digital transformation and should not be treated as a stand-alone project. It depends on everything else a parliament has built. AI is only as reliable as the data beneath it, the security around it and the governance that directs it. The parliaments advancing most successfully put governance in place first, and then data and technology. Those that add AI to existing workflows without redesigning them often find that it amplifies existing weaknesses.

There are also serious questions around digital sovereignty that, whilst not unique to AI, reflect a growing reliance on off-premises technologies. Parliaments increasingly rely on a few commercial storage and application providers whose models can change or be withdrawn at short notice. A parliament needs to understand these dependencies and maintain the ability to carry out its core functions if a service changes. Equally, parliaments need to be sure that their data is safe and know whether information submitted to them is genuine or AI-generated disinformation. AI does not just hallucinate; it offers opportunities for malicious actors too.

For an institution that holds government to account, control over parliamentary data and systems is a constitutional question as well as a technical one.

Third, members have a dual role as both users and regulators of AI, and each depends on the other. As users, members need enough AI literacy to use the tools well, to judge what they produce and to recognize their limits. Parliaments are responding to this, and half of the parliaments we've surveyed provide AI training for members, a figure that's risen significantly in the last two years. The House of Commons has offered a good example: After it launched its AI hub and an awareness campaign, positive staff sentiment towards AI tools rose.

As legislators, members must scrutinize government use of AI, set rules for its safe use across society and protect the integrity of elections and political debate. Members who understand how AI works and how it fails are better equipped to legislate for it. AI literacy is therefore a precondition for effective oversight.

● (1115)

The IPU's Kuala Lumpur Declaration on Parliaments and Responsible AI, adopted by parliaments from 65 countries, recognizes these dual roles that legislators hold. It calls on parliaments to assess their AI readiness, provide continuous learning for members and staff, strengthen internal AI governance and hold governments accountable for their national AI strategies.

To summarize, I would encourage the committee to consider assessing its own AI readiness by using tools such as the IPU's maturity framework, drawing on good practice such as in the IPU guidelines, putting clear governance in place with members involved before extending AI into core parliamentary work and, finally, investing in the future of AI literacy for members of staff.

I'd like to thank the House of Commons for the support it has given to parliaments around the world, and I'd like to ask you to continue to share your experiences with our global parliamentary community. Your contribution has been greatly appreciated by many parliaments.

Thank you, and I look forward to your questions.

The Chair: Thank you so much, Dr. Williamson.

We'll now proceed to questions, and we'll start with Mr. Calkins for six minutes, please.

Blaine Calkins (Ponoka—Didsbury, CPC): Thank you, Mr. Chair.

I think, Professor Karimi, I'll start with you, and the other witnesses can chime in if they feel it's necessary.

I've asked the previous witnesses here a couple of questions about having the competencies of a human in the loop. I think AI is going to potentially replace some of our entry-level—for lack of a better term—types of jobs and careers, those kinds of positions. Those are the people who generally advance and become supervisors and leaders in our institutions later on. If that cohort of people is replaced with AI, what recommendation do you have—not only for Parliament but also writ large throughout our economy—for making sure that we actually have not just a human in the loop but—and you stressed this—a competent, knowledgeable human in the loop to make sure that we can judge the quality of the artificial intelligence that's assisting members of Parliament, businesses and our democratic institutions?

Dr. Amir-Hossein Karimi: That's a very good question, and it's rather difficult to answer because it's hard to predict what the effect of AI on various entry-level jobs is going to be. It certainly seems like the trend is that entry-level jobs are slowly fizzling away, although my opening statement kind of argues the opposite because, when demand from your constituents goes up exponentially, you might, on the contrary, see that you have even more need for entry-level jobs.

The point I'm trying to make is that for entry-level jobs, entrants would have to be trained in an appropriate way to use these tools effectively and to understand where those tools fail.

Blaine Calkins: I have another question, then, and this might be something our friends at the IPU can come into play on as well.

In a typical government scenario, even though Parliament does kind of govern itself, the reality is that the government has great tools and great levers. The government has the levers of our cybersecurity organizations. The government has the levers of its fiscal and budgetary capacities. Ultimately, the government of the day generally has the most parliamentarians; that's why they're the government. The government has the unique ability to decide what the AI strategy is and what the implementation is. The government gets to broadly decide what kind of data it's willing to share and the format in which that data is shared.

Have we seen from any of your research whether there are any parliaments—or anybody—acting in a way where they're using the powers of their government for their own political advancement, over top of just broadly serving the public interest? Are there any examples of that yet that you've seen anywhere on the planet?

● (1120)

Dr. Amir-Hossein Karimi: If this is directed at me, I have to say that I don't have any particular insight into this, but I also welcome the other witnesses to....

Blaine Calkins: Maybe I should ask Dr. Williamson from the IPU.

Dr. Andy Williamson: I think it's a very difficult question to answer.

I don't have any direct evidence of it, but you can certainly see examples of what I might frame as some governments being more proactive in their approaches towards harnessing technologies, which might be considered for their own benefit. However, I don't think there's any hard evidence. I certainly can't see any in a parliamentary context.

Blaine Calkins: Would you say, Mr. Williamson, that the 12-plus members of the IPU are the ones most likely to have adopted the maturity framework and have better AI policies than some of the member states from other groups?

Dr. Andy Williamson: Within the Centre for Innovation in Parliament, we categorize parliaments into three tiers, with tier one being the most digitally mature. Yes, those are the parliaments that typically are well resourced, larger and in developed countries, but not always. There are some outliers in there. For example, the adoption by Mauritius of both digital transformation and AI has been quite staggering. Then we see parliaments like Fiji doing very interesting things as well.

Typically we see the well-resourced parliaments. If I go through my list of the digitally mature parliaments, Canada, of course, is in there, but so are the Brazilian Chamber of Deputies, the Brazilian senate, the European Parliament and the U.K. Parliament. You can see the theme. They're large, developed parliaments.

AI maturity is partly determined by digital maturity. Digital maturity is partly determined by scale and economic ability.

Blaine Calkins: I'd like to go back to a theme from you, Mr. Beall, about continuity and the ability of AI to be used for more than just basic research or an enhanced search engine type of tool.

Could you expand on the continuity thought and just elaborate more on that for our benefit, please?

Chris Beall: Absolutely. Thanks for the question. It's an excellent point.

For those of us who use AI tools regularly, I'm sure everyone's had that experience where you enter a question you've used before and you get a totally different answer. All these different things happen. It's the way in which the models operate.

What we've been really exploring is how can we separate the question around the data, the context and the understanding of what happens from the tool itself and hold it in a separate system, so you are building up your story, your conversations and your experience with your AI use in a way that it retains and understands the memory.

It's not simply about creating an archive. It's actually thinking about the context around the questions and the prompts. It helps you then understand where you got fictional information from, where the fantasy was created, where errors were or where things were caused by people making the decisions. It really is one of those areas where it feels like, from our perspective, there's an enormous benefit to using AI in these spaces. In the experience that I've had—

The Chair: I'm going to have to be a little rude and cut you off. It's my job to keep the clock and we're a bit over time, but you'll have more opportunities to expand on that answer.

Before I turn to Madame Brière for six minutes, I'll just remind witnesses that you may receive questions in French and require translation.

[*Translation*]

Mrs. Brière, the floor is yours for six minutes.

● (1125)

[*English*]

Hon. Élisabeth Brière (Sherbrooke, Lib.): Thank you, Mr. Chair.

To all our witnesses, thank you for joining us today.

[*Translation*]

Mr. Beall, can you tell us if you know of any government that has regulated the use of artificial intelligence in its standing operation procedures, in their regulations?

[*English*]

Chris Beall: Thank you for the question.

[*Translation*]

I'm sorry, but I'll continue in English. I've worked in French a lot in the past, but for several years now I've only spoken English.

[*English*]

At this point, part of the challenge with some of the AI systems.... My answer is no, I'm not, but I'll just take another moment on it.

One of the difficulties that I think we've all experienced when working in the information space is that some of the decisions and some of the conversations are in open spaces, and some of them are in closed spaces.

I do a lot of work in areas around cybersecurity and influence operations. I know from some of the organizations that I work with that there are governments that are thinking about embedded AI systems and systems that are operating off-line and are sort of working in their decision-making processes—some of our partners in the Asia-Pacific and others. What I'm not sure about is how much that's expanded. Where I've seen that happening is in security spaces or intelligence spaces, and I'm not sure where they've expanded that into parliamentary conversations.

There's an important element, I think, where we can look at how we test these things and how we try them out. There's an opportunity for Canada to really push forward, as it's talking about its sovereignty spaces, to talk and think about what that sovereignty and control means.

Hon. Élisabeth Brière: Thank you.

Dr. Williamson, I would like to ask you the same question.

Dr. Andy Williamson: There are a number of examples of parliaments that are adopting AI guidelines that apply right through those parliaments, such as the Italian Chamber of Deputies, the U.K. Parliament, etc. I know that you have some guidelines.

In terms of the sovereignty question, there's an interesting piece of work happening in Norway, where they're looking at very much opening up the AI so that they have their own AI tools within the public sector that are transparent. That's perhaps another interesting example. I'm not sure if that quite answers your question.

[*Translation*]

Hon. Élisabeth Brière: You have answered in part, yes, thank you.

There is a lot of ambition surrounding artificial intelligence, and it must be tempered with realism. Two of you have said that parliaments haven't adopted the use of artificial intelligence across their institutions.

Artificial intelligence and its benefits will only bear fruit if governments establish pragmatic goals and carefully plan its implementation, but that means a heavy workload.

Mr. Williamson, how much will its implementation cost?

[*English*]

Dr. Andy Williamson: I think the best way of looking at that is in the broader context of digital transformation. What we've seen in the last 20 years is that parliaments have gone from technology being a back-office function, something that was just there to support the administrative side, to really, when we hit the COVID-19 pandemic, becoming core to how parliaments could continue to function.

We saw a watershed moment, if you like, and it really highlighted the fact that the digital platforms within parliaments are now strategic. They're part of the core business. What that's done is shifted the responsibility for digital systems up the organization. Now we see that sitting with typically the board or the secretary general and the senior management rather than it being a subsidiary decision about which tools to use.

AI builds on that in that it really requires strategic governance across the institution. The requirement to govern an AI and, therefore, the cost of implementing AI isn't simply an operational one. It's about putting in place the rules and regulations, the procedures for how it works, the systems for auditing, quality assurance and all of those things. They sit at the highest level of Parliament. There is a significant overhead, if you like, at a more senior level as that shifts.

● (1130)

[*Translation*]

Hon. Élisabeth Brière: Thank you very much.

Mr. Beall, would you like to add anything?

I'd like to draw your attention to the cost. Mr. Williamson says there are no operating costs, that costs are mostly associated with skills. We know the adoption of AI favours the private sector, so there's a huge gap in skills between the public and private sectors.

Doesn't that make an institution such as the government more vulnerable?

[*English*]

Dr. Amir-Hossein Karimi: Was that for me?

Hon. Élisabeth Brière: Yes.

Dr. Amir-Hossein Karimi: I'm sorry. You initially mentioned his name. Could you repeat the question one more time?

The Chair: If you could ask it quickly, I will give a bit of extra time for the answer.

[*Translation*]

Hon. Élisabeth Brière: I said it had an impact on skills and that there's a big gap between the private sector and the public sector regarding skills, because workers will turn to the private sector. There are more benefits for them.

Do you think that makes the government more vulnerable?

[*English*]

Dr. Amir-Hossein Karimi: I certainly think that continual education and upskilling in response to these new technologies is one thing that the government should invest both senior and junior staff members in pursuing.

The Chair: Thank you so much.

[*Translation*]

Ms. Normandin, the floor is yours for six minutes.

Christine Normandin (Saint-Jean, BQ): I'll start with a question for you, Professor Karimi.

You talked about the risk of an increasing number of public submissions created by AI and that we'd need to use artificial intelligence to analyze and identify their origin to see if they were created with AI.

As far as you know, how advanced are the tools at our disposal to assess if a document was created using artificial intelligence? In Quebec, there was controversy about a book written using AI. The tool Pangram had concluded as much. However, should we be relying on AI to detect whether artificial intelligence was used? To what extent is human intervention required in the process of identifying these sources?

[*English*]

Dr. Amir-Hossein Karimi: That's a very good question.

I'd like to put on, for a moment, our engineering hats and remind us that these AI systems are also engineering systems. Engineering systems, at some point, have a failure point. The topic of hallucinations has come up, both in this discussion and in past discussions. I think we have to fundamentally accept that these systems will hallucinate to some extent. With newer and newer technology that comes out, maybe that hallucination gets pushed down into the stack and slowly starts to fizzle away, but we cannot fully remove humans from the loop. We need the responsibility to be with the institution and the individuals who actually deploy those AI systems.

[*Translation*]

Christine Normandin: Thank you very much.

Mr. Williamson, following the same line as my colleague Mr. Calkins, I'd like to hear your thoughts on the Inter-Parliamentary Union's maturity framework.

Can you tell me how it works? Do the countries do some type of self-assessment or do peers or external stakeholders evaluate their level of maturity? Are the countries' maturity levels publicly available? Can you tell me more about the criteria used to assess these levels of maturity?

[*English*]

Dr. Andy Williamson: The maturity framework is very much in the public domain. It's downloadable from the IPU website. We realized when we developed the IPU guidelines for AI that they're really complicated. There are lots of them. The first question you have as an institution is, where do I start? That question actually depends on where you are and where you want to go. We realized that the AI journey isn't fixed. No two journeys are the same.

The concept with the maturity framework is very much that it allows you to self-assess. It's not intended as a complex scientific process. It's very much a self-assessment process so you can get a feel and a sense of where you are. It starts off at level zero. We know we know nothing. We don't do anything. Maybe there are a few people using ChatGPT over in that corner, but they're not talking about it. It's very much informal use, and that's where we've seen parliaments come on in their journey. At level one, they start to think they need some process to manage that.

It goes through the various levels until you get to level five, which is that you become a partner in learning for other parliaments, you're sharing what you're doing. The interesting thing is that when we've asked parliaments to look at the framework, they've come back and they've gone, “Wow, we're level two.” I'm not talking about small parliaments here. I'm talking about big parliaments. They realize they're very early on in that maturity journey.

The idea of the framework is not so much a route map as a how to put together your own route map.

● (1135)

[*Translation*]

Christine Normandin: Thank you very much.

Mr. Beall, you talked about having closed structures in some cases, where data is kept internally to preserve sovereignty and for safety reasons. You also talked about adopting more federated approaches so democratic countries can exchange data. I guess we're talking about open data, or even open systems. I'd like to hear your thoughts on the difference between the two.

When should we choose closed systems and when should we choose open systems? Are they really compatible? Could you explain the difference between the two?

[*English*]

Chris Beall: Thank you so much for asking. That's a great question.

One of the things we've been thinking about as we develop our model and think through the different pieces we're working on...and it's something that's been important in my journey through these public policy challenges for a longer period of time: How do we think about power relationships, who owns data, who owns our structures and how we make decisions about those pieces?

For us, your question points to a key feature that needs to be built into the systems. We need to have systems that keep private the things we want to keep private and share the things we want to share. We should be making the decisions, not the tool, not the company, not the external providers. It shouldn't be a third party decision. We shouldn't be put in a position where our data is sitting in a third party location, where it could be under threat, accessed and challenged without our being there.

The answer to your question is that it becomes a question for parliamentarians to decide what the pieces are that you want to share with other parliaments, with Canadians and with other elements. What are the pieces of your conversations that you have here in the House and in the different spots that you want to be accessible to Canadians, and are there ways you can open them up even more? What are the things that should remain private and should be held under privilege, and how do you make sure that those rules are kept and set? As my fellow witnesses have talked about, you need to build in the right kinds of guidelines, tools and structures to make sure that it's not simply a matter of policy that it's not shared, but that it's actually infrastructurally unshareable and you have the ability to set your....

Thank you so much.

The Chair: Thank you.

We'll now turn to Mr. Moore for five minutes, please.

Hon. Rob Moore (Fundy Royal, CPC): Thank you Mr. Chair.

Mr. Karimi, one thing you mentioned is what to do if something goes wrong. For the context here for members of Parliament, most of what we do in our constituency offices when we do case work is because something has already gone wrong. We don't normally deal with the easy cases. We don't hear about the easy cases when someone gets the answer they need or they have been successfully dealt with by whatever federal department they're dealing with. We get the cases where someone is frustrated. They're at their wit's end.

Often it involves inconsistent answers, not being able to get an answer or delays in finding an answer, whether that's with Veterans Affairs or Canada Revenue Agency, etc. We only really deal with the tough cases. Sometimes people want accountability. They want a name. They want a record of decision. We need to know how they arrived at the state they're in. Who told you what? What were the names of the people or the identifying number of the person?

In a system that's already so bureaucratic, how do we maintain that decision-making accountability as more and more AI is introduced into the process, whether it's at the early, mid or even late level of decision-making?

Dr. Amir-Hossein Karimi: I think the fact that you started your question with, “Some cases have reached our office” illustrates the point that I want to retain essentially in your procedure, which is that there should be a process by which your constituents who have faced an adverse decision can reach your office. This is not something that's a given as the scale of the submissions dramatically increases.

You have to think about that process in advance and, as my fellow witness was saying in their opening statement, there should be logs and traces of all the uses of AI and who authorizes a particular use. As long as that's there and as long as you think about the fact that it's going to fail at some point and make sure that the constituents can contact you in advance, then I think you're off to a good start.

● (1140)

Hon. Rob Moore: Okay.

Dr. Beall, do you have any thoughts on that scenario? We used to hear of examples like a small business that would call CRA multiple times, speak to multiple agents and take the average of their advice. They'd get a yes, a no, another yes, another no and then a yes, so they'd go with yes. In light of the fact that this is the kind of world we're dealing with, how do we maintain some kind of order in the process so that we're not starting from scratch when someone reaches out to us?

Chris Beall: That's a great question.

I would say, building on my colleague's response, that we need to then verify the systems we're putting in place to make sure that we can set structures and guidelines so that they essentially remember those things. There's a real tension that comes. On the one hand, we want our systems to remember the decisions so that you're not giving people multiple spaces. At the same time, on the other hand, you still have to maintain controls over privacy.

The idea of your government having a growing record of every time you've reached out, which is run by AI, sounds like the beginning of a bad science fiction movie. We need to have, within those bounded controls, within those bounded systems, ways in which the systems we run can know and be able to touch base on and understand the messages that have gone before and know why advice might change, because advice does change. Things make errors. We need to understand why those are there and be able to go back and verify so that when your office is looking into why I get three different answers and why all three of them were wrong, you can go back and understand the decisions and the context behind each one.

I'll quickly use an example from the private sector that we've been working on with colleagues. When you're running multiple systems across multiple suppliers in multiple communities, you build up relationships over periods of time and you understand how they all work. What we found is that building in these kinds of memory-safe elements to what we're working on simplifies the process of deciding in which case you use which suppliers and how you use them, because you can remember that these things went wrong at the time, but now they have a new CEO and these are corrected. Being able to put those lines through simplifies the experience for the business and simplifies the experience for operations. That's the kind of thing, within the boundaries you have built in, that feels like it would help the situation you're talking about here.

The Chair: You have five seconds.

Hon. Rob Moore: Okay. I think we'll leave it at that.

Thank you.

The Chair: Thank you so much.

We'll turn to Mr. Caley for five minutes, please.

Braeden Caley (North Vancouver—Capilano, Lib.): Thank you, Mr. Chair.

Thank you to all of the witnesses for the really insightful additions to the study. I want to start with Mr. Beall and Professor Williamson.

There was the mention in Mr. Williamson's presentation of the core dependencies that can arise from government using these models. We talked about that a little in previous sessions.

It comes to mind that the ownership structure of one of those companies could change entirely. I'm wondering if you could elaborate a little on some of the concerns for you in that situation, and how they might relate to a country's sovereignty.

Andy Williamson: Shall I take this first?

Braeden Caley: Yes. Please go ahead.

Dr. Andy Williamson: That is a very good question, and I think it's one that we have to be very concerned about.

I can start by way of an example, which is that the Chilean Chamber of Deputies was using an AI tool and was starting to shift it into production when the supplier turned it off with virtually no notice. There's a real-life example.

We have to be aware of it. It makes us think about where our data lives, not just in terms of data security but also in the consistency and availability of that data. The solutions are typical IT solutions around redundancy, but also around having control. It comes to that judgment point: When is it appropriate for data to be outside of the institution and when should it be inside?

This comes down to a number of parliaments that we're now seeing creating their own small language models and hosting them internally or within a government AI cluster. Particularly in Europe, this is now a much bigger question. With the recent geopolitical activities, the reliance on U.S.-based technology companies is raising significant concerns right across governments in Europe, and many of them are looking at local solutions. Norway, Estonia and France are the obvious ones, but the EU itself is also looking at internal solutions.

I think the question comes down to this: When is it okay for this data to be on someone else's server, and when should it be on yours?

● (1145)

Braeden Caley: I think that answer covered it thoroughly, so I'll switch gears slightly and go to Mr. Beall.

I like something on your website. It says, “Start with the obligation the system must carry.” Both you and Dr. Karimi spoke about this.

I'm wondering if you both could elaborate quickly on how humans can work within the loop, as Mr. Calkins has spoken about so often, to ensure that in those systems, whether they're for parliamentary use or for broader government use, as we've talked a bit about today, you are setting those obligations within the right standards and you have the right human resource role—which I haven't heard many good titles for—in assessing that.

In government, what would that job description, or even just that job title, look for as the human in the loop who is assessing those kinds of things?

Chris Beall: I'll start, if you don't mind.

That's an excellent question. What I worry about a little bit is that “human in the loop” has become one of those terms of art, like “sovereignty”. We say it and we sort of know what we mean, but we haven't really thought through, as you were saying, what that role looks like and who that person should be.

One that we think about through QuietWire is, on the one hand, making the system so clear and easy to use and transparent that you don't need to read the 15 pages of print to understand exactly what you're signing on to. You are setting out the parameters you want in your language, in your space.

On the other hand, in terms of assessing the tools that are in place, you need somebody who, from our perspective, is not simply just a technical specialist but who also understands how those technical specialties line up in the public policy and parliamentary domains.

There's a moment we're at right now. We've had 10 years of misunderstanding how social media works and of not putting the right controls and parameters in place. We're finally getting around to putting those in place now, 10 years too late. There's a real worry that we're doing the same thing now.

I think we need to see the pieces that were missed. In my view, I think it's those people who could bridge those gaps, understanding the technical side and being able to give that right word, but then seeing how it would fit into the goals and guidelines of what we want to achieve. I hope that's....

Braeden Caley: Absolutely.

Dr. Karimi, would you comment?

Is there time?

The Chair: You have 10 seconds. Answer very quickly.

Dr. Amir-Hossein Karimi: Complementary expertise can be in one person or multiple people. At the end of the day, you want domain expertise plus AI expertise.

The Chair: That was excellent and with time to spare. Now you're out of time.

[*Translation*]

Ms. Normandin, you have two and a half minutes.

Christine Normandin: Thank you.

Professor Karimi, a number of witnesses have talked about their fear that, due to a growing reliance on artificial intelligence, we may lose our ability to think critically when making decisions.

I'd like you to tell us about the idea you proposed, of establishing effective recourse before AI deployment so that we know beforehand who is accountable and how to correct any problems that arise. Could this kind of approach help us maintain our ability to think critically? Is there anything else that you think could help us maintain the critical thinking skills we need?

[*English*]

Dr. Amir-Hossein Karimi: That's a very good question and, certainly, an open question in the field of human AI collaboration. This is still an unanswered question in the field: the extent to which AI use helps with creativity, critical thinking or reasoning processes. There are some early studies that show a non-linear curve—in that, to some extent, the use of AI can help improve any of these factors that we care about, after which it becomes an overreliance and what's known as an infantilization of the individual, an overreliance on these tools.

There was one other thought I wanted to share, but it's eluding me now.

[*Translation*]

Christine Normandin: Mr. Williamson, I'd like to hear you tell us more about the shadow AI you mentioned. Is there a correlation between a country's lower use of shadow AI and greater digital maturity, or the opposite?

● (1150)

[*English*]

Dr. Andy Williamson: No, it's actually a wonderfully global phenomenon. You can look at the least digitally mature parliaments and the most digitally mature parliaments and see that everyone's using it. The only difference is that the least digitally mature are less likely to have any standards, regulations or oversight.

[*Translation*]

Christine Normandin: Generally speaking, developing our own AI will therefore not prevent members, public servants or anyone else from using private AI systems. We shouldn't think that developing our own system would insulate us from the use of other technologies.

[*English*]

The Chair: Answer very quickly.

Dr. Andy Williamson: No, that's a matter for regulation by that parliament.

In Canada, you don't restrict, particularly, what you can use. You might be piloting Copilot, but MPs can use what tools they like. Some parliaments are restricting what can be accessed. It's a decision for the individual parliaments.

[*Translation*]

The Chair: Thank you very much.

[*English*]

Mr. Cooper, you have five minutes, please.

Michael Cooper (St. Albert—Sturgeon River, CPC): Thank you, Mr. Chair.

Dr. Williamson, in your opening statement you noted that AI is being widely used in parliaments throughout the world for things such as transcription, translation, presumably research, streamlining, administrative processes and so on. You went on to state that in areas such as drafting analysis, far fewer are using AI for this and that this really raises the democratic stakes. You said that is "where the democratic stakes are the highest."

More generally, can you speak about categories of risk in terms of the deployment of AI systems in the context of parliamentary systems?

Dr. Andy Williamson: Yes, of course, and I'll take, as an assumption, the general risk of AI that it brings to any function.

The way of looking at it is as a continuum from the criticality of the function in terms of its centrality to democracy. If we're looking at something that's drafting legislation, I don't think I want ChatGPT to be left unfettered to draft security legislation. I don't really want ChatGPT trying to work out what the impact of an amendment is on a whole statute of laws, which is one of the reasons why these things are not being done as much. It's because they're complex. It really comes down to how integral to the democratic process this thing is that we're trying to apply AI to.

Translation and transcription, on the other side, although fairly integral to the process, are actually very well tried and tested. The other side of it is how reliable the systems are at handling this functionality. On the procedural side, they're pretty good. You can get a sufficient success rate and a low error rate. On the more complex tasks that require institutional knowledge, legal knowledge and real critical thinking skills, they're much less proven and reliable.

Michael Cooper: Thank you very much.

I think there's general consensus that guardrails are key. Does the IPU have a set of recommended guardrails?

Dr. Andy Williamson: Yes, but in the sense that the AI guidelines for parliaments include a number of procedures and functional areas where we recommend how AI can be adopted. We try to do it in a non-rigid way because we accept that everybody's working differently. When it comes to specific guardrails, it's very much up to the parliament to take the good practice from what we're publishing and adapt it to their own working procedures.

Michael Cooper: You spoke about AI literacy being an absolute prerequisite on the part of parliamentarians, both as users and regulators of AI. Can you speak a little more about AI literacy? For example, should training be mandatory, and what would that look like?

Dr. Andy Williamson: Probably, yes. That's the problem with—

Michael Cooper: I would just add to that: Do you have any recommendations about training? We have had training. It was recent. There has been some uptake, but I'd be interested in your insights.

● (1155)

Dr. Andy Williamson: I think mandatory training is interesting. I have an eight-year-old, and if I tell him it's mandatory, it ain't happening. He's going to resist.

I think the way forward is "just in time", but it also requires an institution-wide education program, if you like, to make sure people understand why this matters. It's no good saying, "You have to do AI training". You need to understand why it matters, and the importance of it is that it's everywhere. One of the challenges is recognizing when AI is there in the tools that you're using but also in the information that you receive. Being able to understand better when AI is in the picture is really important.

My recommendation is perhaps not "mandatory" but "highly encouraged". I think underpinning that is very much an education campaign about why it's important, and why it's going to become even more critical and important.

Michael Cooper: I'll leave it there.

The Chair: Thank you so much, Mr. Cooper.

We'll turn to Ms. Mingarelli for five minutes, please.

Giovanna Mingarelli (Prescott—Russell—Cumberland, Lib.): Thank you, Mr. Chair.

Thank you to all of the witnesses for being here today.

I will begin my questions with you, Mr. Beall. What's the greatest threat to institutional memory today: technology moving too quickly, people moving too quickly or organizations failing to value what they already know?

Chris Beall: That's a fantastic question. Thank you very much.

I think there's an interconnection between them. There's definitely an organizational challenge, and I think this predates AI. Those working in organizations understand how little we keep track of why we make the decisions that we make and how we codify them. In fact, my interest in this is that there's a real opportunity, by using AI tools properly, to enable us to build from that and not be constantly trying to understand why things happened in the past. Creating and retaining history is an essential way forward.

I think, though, that AI has opened up a new risk previously unseen: It's not just that it's not remembering why we did things or making decisions based on those things, but it's also unfettered AI access to your system runs the risk of it rewriting the history to match more what it wants its records to be.

I had a terrifying conversation with friends at the OECD earlier this year. They're looking at the big datasets they use to maintain the history of the economy over the last 100-plus years. AI access to those systems, they're worried, actually gives it the chance to go in and mess things around, to fix things up.

I think we need to really understand, as part of the guardrail discussion and as part of setting the boundaries, why we want to use it and to challenge ourselves on how we use it already, and then set up AI to be the tool we need rather than the risk that's undermining us.

Giovanna Mingarelli: Thank you.

Just to build on that point, the last thing you said about setting up AI to be of service to us rather than undermining us, what do you think is the best way to do that?

Right now, AI is heavily and predominantly a commercial endeavour. We, as the government, of course, are developing and building AI tools to be of service to the government and to Canadians, but the government doesn't have an army of AI developers building our own sovereign technology. We still have to hire external resources. What do you think is the best way to incentivize the development of prosocial AI that will be beneficial to society and to the government?

Chris Beall: That's an excellent question.

From our perspective, one of the things to look at is where the commercial opportunities are that AI companies and systems are picking up. Are there opportunities for government to step in to create incentives for prosocial, civic AI and also look at whether you can set up different barriers?

The monetary possibility or the economic benefit for a lot of AI companies is around scraping the data and data usage. If you wanted to encourage non-data extractive technologies, then you need to ask if there are ways to limit the extent to which AI systems are able to take in, hold and use people's data. You also need to ask if you can create economic benefits for companies that are not doing that, by opening up other possibilities by saying that there are certain kinds of activities where you can't scrape data and use data—not just for training purposes but also for using it as the fodder for building up the advertising sales around it. I think there's an opportunity to take both sides of that equation.

● (1200)

Giovanna Mingarelli: My last question builds on that.

When we talk about data scraping and the collection of the data, at this time there's still an added layer of not being able to verify either the identity or the sovereignty of the individual. The data that they submit could be anything. It leads to poor quality control in some scenarios. In the worst-case scenarios, it leads to issues like foreign interference, because you don't have the ability to manage properly the incoming data because you can't actually identify and properly manage the end-user.

On the flip side of that, we have the issue that individual users have the right to their own data. It's data sovereignty for the individual. We are the government and we're building civic AI and civic technology, so we want people to also have ownership of their data.

There are the two sides. How do we make sure that the incoming data, whether it's scraped or not, is verifiable? How do we ensure that the identities of the people who are submitting the data are verifiable?

The Chair: Before you answer, I have to apologize. That was a very long lead-up to the question and we're over time, unfortunately. If you do—

Giovanna Mingarelli: Can he have 10 seconds?

The Chair: Can you answer it in 10 seconds?

Chris Beall: I'd be happy to submit something.

The Chair: If you submit something to the clerk, that would be wonderful. Excellent. Thank you so much.

I'd like to thank all of our witnesses. We will suspend the committee for a few minutes for the next panel.

● (1200)

(Pause)

● (1205)

The Chair: Welcome back, everyone.

I'd like to welcome our witnesses for the second panel. Appearing as an individual, we have Fenwick McKelvey, full professor, information and communication technology policy, Concordia University. From the translation bureau, we have Jean-François Lymburner, chief executive officer, and Matthew Ball, vice-president, linguistic services.

We'll start with our witnesses who are here in person.

You have five minutes, please.

[*Translation*]

Jean-François Lymburner (Chief Executive Officer, Translation Bureau): Thank you very much, Mr. Chair.

Thank you for the opportunity to contribute to your study on artificial intelligence and Canadian democratic institutions.

I am joined today by my colleague, Matthew Ball, Vice-President of Linguistic Services at the Translation Bureau. I would also like to acknowledge the interpreters in the booths, Charles, Renée and Alexandra, and thank them very much for their work.

At the Translation Bureau, we have been working with language technologies for nearly fifty years. We have had experience with computer-assisted translation, translation memories and, more recently, generative artificial intelligence. This experience has led us to a simple conclusion: for democratic institutions, the main challenge posed by artificial intelligence is not technology alone, but rather trust.

Our institutions depend on citizens being able to access reliable information, understand public decisions and interact with government in both official languages. For this reason, language is not simply an administrative tool. Language lies at the heart of access to information, civic participation and the ability to exercise one’s rights. When AI is used to produce, translate or disseminate public information, the reliability of language directly affects public trust in our institutions.

Our experience has allowed us to identify several principles that extend far beyond the translation field.

First, the quality of AI outputs depends on the quality of the data it is trained on. Having a high-performing model is not enough. It is equally important to understand the source, quality, representativeness and limitations of the data that is used.

Second, data sovereignty and data protection are critical issues, as we just heard. Public institutions must know where their data is hosted, how it is used and who can access it. They must remain in control of their data so that they can use it for the public good. In a digital world, the sovereignty of government data is an essential condition for trust.

Third, expertise and human accountability remain indispensable. Artificial intelligence offers speed, processing capacity and power. Yet in the language field, it can produce content that is deceptively natural-sounding and convincing while being fundamentally incorrect.

Language professionals do more than correct artificial intelligence; they help make it trustworthy. They check meaning, understand context and nuance, safeguard linguistic, cultural and institutional integrity, and detect errors that may go unnoticed by non-specialist users. Their role begins long before the final result. They improve data, help shape technology, establish quality standards, assess risks and contribute to continuous system improvement.

The question, therefore, is not whether to choose artificial intelligence or human expertise. Rather, it is how to combine the power of technology with human judgment to achieve better, safer and more reliable outcomes.

This complementarity is also built through collaboration with the academic community. Through scholarships, internships and partnerships, the Translation Bureau is helping prepare a new generation of language professionals who combine linguistic expertise, professional judgment and the responsible use of AI.

● (1210)

[*English*]

Our experience with GCtranslate provides a concrete example of this approach. Developed to meet the needs of the federal public service, GCtranslate puts into practice the principles of high-quality, Canadian linguistic data, protecting government information, human expertise at the heart of the solution and a gradual risk-management-based rollout.

However, the most important lessons extend well beyond GCtranslate. Innovation and prudence are not incompatible. Public institutions can benefit from artificial intelligence while maintaining high standards of quality, security, transparency, governance and accountability.

For our democratic institution, the question is, therefore, not simply what artificial intelligence can do. It's also a matter of when it should be used, with what data, under what safeguards and under whose responsibility.

[*Translation*]

Artificial intelligence can enhance government capacity and improve access to information. However, if it is also to strengthen our democratic institutions, its use must uphold public trust, protect the information entrusted to us and fully respect Canada's linguistic duality.

It is in that spirit that the translation bureau will continue to leverage its linguistic expertise, academic partnerships and experience with artificial intelligence to support Canadian institutions.

Thank you for inviting us.

[*English*]

The Chair: Thank you so much.

We'll go online to Professor McKelvey for five minutes, please.

Fenwick McKelvey (Full Professor, Information and Communication Technology Policy, Concordia University, As an Individual): Thank you for the opportunity to be here today.

I'd like to begin by acknowledging that I'm speaking from unceded indigenous lands. The Kanien'kehá:ka nation is recognized as the custodians of the lands and waters on this place I call home.

My comments, in a short time, aim to cover my broad interdisciplinary research into democracy and digital technology, and are based on 20 years of direct experience and historical research dating back to the 1950s. I've written three books broadly related to AI and democracy, two of which are open access. They are *Internet Daemons* and my new book, *SimPolitics: America's Quest to Solve Politics with Computers*, published this year by MIT Press.

I have three recommendations.

First, create a minister of democratic institutions and AI.

I have participated in at least three consultations to address the social impacts of AI and its democratic implications. None of these efforts has merited a serious response from policy-makers, and too often, meetings are simply used as evidence of consultation.

I cannot be more clear: Canada's approach to AI governance is undermining its democratic governance. Time and time again, hastily developed industrial strategies with vague promises of digital sovereignty undermine actual democratic sovereignty.

Canadians and Canadian civil society have answers and solutions if they are able to meaningfully participate. In 2023, I was part of the ArtIA strategy, an internationally recognized collaboration on commons-based management solutions for AI in the cultural sector. That 18-month consultation brought together over 230 participants across the arts and cultural sector in Quebec, only to present once before committee. We're still waiting on movement from the government on meaningful policy for AI in the cultural sector and a demonstration of listening to these types of projects.

In 2023, I was lucky to be part of a broad coalition around the People's Consultation on AI. Following an open letter with over 200 signatures, the consultation held a four-month open call that received 66 submissions. Thanks to my colleagues Dr. Joanna Redden, Meghan Voll and Leo Litke, we understand the range of submissions and key concerns. Cynthia Khoo submitted these responses to many different ministries, and again we have yet to hear a response.

Creating a dedicated institution to handle and coordinate issues around democratic governance and AI is essential. Building on past examples like the tiger team, it becomes a way of ensuring that the social impacts are given the necessary attention and that the ability to change our democratic institutions to ensure just and fair participation in governance will continue in the future.

With the time I have left, I have two more comments.

More broadly, there's a growing concern, as we've noticed from panellists, around the relationship between AI and privacy. One thing that needs direct attention is to close the loophole for third parties processing political data and ensure that Canadian political parties have meaningful rules around the use of AI.

There is clear consensus for solutions. A recent national Ipsos poll found that 80% of Canadians say that political parties should follow the same privacy rules as businesses and public institutions.

While we focus on deepfakes and AI-assisted messages, we miss major changes in Canada's political infrastructure. AI is changing how parties and political actors process and analyze their data. That applies to governments as well. National Observer reporter Rory White followed the Toronto mayoral election campaign, in which mayoral candidate Brad Bradford transferred voter data to an international third party to encourage his voters to post AI-generated messages of support on his social media.

Nothing is more disingenuous than political parties wringing their hands about the threats of AI to democracy without the willingness to undertake basic reforms to privacy law and data protection to ensure that they're accountable for the ways that they're processing and handling voters' personal information.

Number three is to include public AI and an AI commons as critical components of Canada's AI strategy.

Finally, I want to emphasize that key data commons—Wikipedia, Creative Commons and the Internet Archive—as well as public service media are suffering. Bots and AI scraping dramatically raise the cost to provide free and public services online, as well as the cost to government to provide these services. AI undermines the reciprocity necessary to keep these sites online, creating a paradox as the content is used more and more with fewer and fewer maintainers.

This is an exciting time for innovation and change if we recognize that Canada's commitments to communication rights need to be renewed for an age of general AI.

I'm here today, having stepped out of an international symposium on data commons, and I contest that there's a global dialogue and that Canada needs to better integrate these discussions into its federal AI strategies. Canada's AI strategy needs to invest in public service media and open Internet institutions, as well as ensure that work done at the committee here understands its role in supporting these AI commons and data commons initiatives.

Thank you very much.

● (1215)

The Chair: Thank you so much.

We'll now go to questions.

We'll start with Mr. Cooper for six minutes, please.

Michael Cooper: Thank you, Mr. Chair.

I'm going to direct my questions to the translation bureau.

On Tuesday, we heard from Antoine Hersberger, vice-president of the Canadian Association of Professional Employees, representing the translation group. With respect to the translation bureau, he said that while AI is a valuable tool in providing translation services, at the bureau it is being deployed rapidly, which has resulted in compromised quality and reliability. He said that this has led to a new quality standard at the bureau, that standard being “good enough”.

Obviously, a standard of “good enough” isn't good enough. What is your response to Mr. Hersberger's characterization?

Jean-François Lymburner: It's about quality. The translation bureau has been producing translation for over 90 years. We still have our specialized services, but we must recognize that AI is around us.

One of your previous witnesses spoke about the use of shadow AI that was happening in government. To a certain extent, the demand for traditional translation has been going down for the past 10 years. We knew that public servants were using all the kinds of public tools, custom or off-the-shelf tools, that they could find, and there was no way for us to really interact with those users.

At the translation bureau, we still receive highly important documents, such as laws, memorandums to cabinet and budget figures, and our specialists are still working on those. There are types of demands that we don't see anymore. Not a long time ago, we could have seen an invitation to a meeting that happened last week. It was simply a compliance issue. Our translators are not seeing those demands anymore. It was obvious that we had to provide a secure self-service for quality data that public servants could use.

When I heard the testimony a couple of days ago, I found there was a lot of confusion. Our translators have been using AI. They have a very sophisticated tool box, the best tool available to do translation. When we provided self-service to public servants for their routine types of translations, we knew that was a need. In only five or six months, almost 800 million words have been translated with our self-service tool, while the bureau still receives the important documents that need to be translated by a translator.

In addition, we provide warnings and information about the risk. It's a matter of risk. I get the "good enough" point, but I'm also very concerned about that. With that being said, when clients accept a certain amount of risk for routine translations, they might not need the top-of-the-line translation, depending on the risk.

Michael Cooper: Mr. Hersberger went on to say that translators are overworked, that AI is not keeping pace and that this has also contributed to a downgrade in quality.

What is your response?

● (1220)

Jean-François Lymburner: I'll let Matthew, who's the vice-president of probably the largest group of translators in the world, answer that question.

What I can say is that there are actually two aspects of the bureau that I think are sometimes being confused. We are supporting Parliament. That's our utmost priority. We have translators who are here and are supporting the work of Parliament day in, day out. That business is happening in two blocks of about 50 days per year, the session in the winter and the session in the fall. You all know that those are very intense. The translators who are working in those conditions are actually working overnight and making sure that everything is available the next day. Then we have another business: We provide translation to departments based on their needs.

I think what was mentioned here was that, yes, it's intense while there's a session. We do use some level of intelligence to support your work, although it is within the confines of the House of Commons.

Matthew, do you want to add anything on that?

Matthew Ball (Vice-President, Linguistic Services, Translation Bureau): I will.

I think that Jean-François has covered this quite well. Honestly, we provide services to all government departments. It's a requirement, and our service levels have not diminished. I'm very confident in the quality of the products that we're providing.

Michael Cooper: On this issue of quality and reliability, when Mr. Calkins asked Mr. Hersberger whether he could identify or point to any major errors, he was unable to do so, but he did say that the broader issue is that errors that go undetected—errors hidden, for example, within well-written passages or sentences.

Can you provide any insights into the extent to which this should be a concern and what steps are being taken? What processes are in place and what auditing mechanisms are in place to see that these errors are addressed?

Matthew Ball: I think we share the same concerns as Mr. Hersberger. Obviously, AIs are powerful tools. Our language professionals, be they interpreters in the booth, terminologists working in our offices or translators working for departments, all master these tools and are very well equipped to handle them.

For GCtranslate, which is a self-service platform that the bureau provides for departments, this is a concern. The first thing that a user will see is a notice saying that AI makes mistakes and to be careful.

As we all know, departments—

Michael Cooper: I'm sorry to interrupt, but what's your—

The Chair: We are over our time. I did want to hear the answer to that question, but we are well over.

We will turn to Ms. Vandenbeld for six minutes, please.

Anita Vandenbeld (Ottawa West—Nepean, Lib.): Thank you very much, all of the witnesses, for being here.

Just before I go to some of the deeper questions on AI, Mr. Lymburner, I'm part of the Copilot pilot project, and I've noticed as we've been using the pilot project that some of the policies have changed.

In the beginning, one of the things that I found fantastic was this. My French is good enough to allow me to type a message to a colleague with grammatical errors and probably some mistakes, but it took me a lot of effort. What I was able to do was just type it in English, and then, if it's for a francophone colleague, just say, "Please translate this into French." The Copilot in my email would do that for me.

Yesterday I tried to do the same thing, and I like the distinction you made between routine and formal. I tried to do the same thing, and I believe there probably has been a policy change, because I got a message that said, "Sorry. I can't translate or draft political communications for public officials or political representatives." It was just a simple email to a Bloc Québécois colleague.

I'm wondering if the pendulum is going almost too far the other way. What I did was use Google Translate. I speak enough French to be able to look at it and say, "This isn't quite right. I'll change this. It's a little too formal."

I wonder where that line is. Could you maybe comment?

Jean-François Lymburner: This is an excellent question. I'm very happy that you asked it.

I think you had Mr. Janse here in this session. They are providing the set of tools within the House of Commons. The tool that we were going to talk about a bit later is probably GCtranslate, which is a translation tool that is available to departments, but it's not available within the House of Commons now.

You're absolutely right that Copilot could be extremely helpful to help you write in any language, but at the translation bureau, we don't provide generative AI like Copilot. We're harnessing AI for specific translation needs. You haven't seen that tool. Our translators have a tool within the House of Commons. Hopefully, there will be access, and we're in discussion right now about how you could access a tool.

You mentioned that you reverted back to a public tool or a free tool, which means that the information gets out of the country, and that's what we're worried about. That's why we felt there was a dire need to provide self-service to public servants so that they can write.

We see that in our data. In fact, somebody mentioned that we were only doing 10% of French to English. Since we have had our self-service tool out there in departments, it has gone up to 30%. That means that people are more confident to write in the language of their choice, which is another piece of good news.

Hopefully, we'll bring a tool for your translation needs within the confines of the House of Commons.

● (1225)

Anita Vandenbeld: I appreciate that, and Copilot is a trial. I think we're learning as we go, but I really appreciate that perspective. Thank you.

The next question I have is for Professor McKelvey.

You talked about the need for a minister of AI and democracy. We have, of course, a minister of democratic institutions. We also have a Minister of AI. Of course, global democracy falls under Global Affairs Canada, but a lot of this is under national security and defence. It's falling in many places.

When we were in the Subcommittee on International Human Rights, there was testimony that you needed a coordinating function within government at a senior level, like an envoy or something, who could then coordinate democracy, including AI, across all government departments so that there's not one for domestic democracy, one for international democracy and one for AI.

Would that be something you think would be helpful?

Fenwick McKelvey: First off, part of the reflection here draws on my past experience on seeing what the tiger teams were doing around social media and thinking about the opportunities across ministry collaborations on these intersectional issues that you raised.

Part of the key concern to me is that very often, when we're talking about where these questions of AI sit, we see that AI falls under different mandates. Most often, if you look at the past 12 years of Canada's AI strategies, you see that in general they have focused on R and D and industrial policy. We're playing catch-up on these critical issues around literacy, the capability to use these technologies and in understanding how they impact personal privacy.

My recommendation was really meant to be a provocation to think about an institution that has a dedicated commitment to try to process these issues. Rather than thinking of it as a one-off, it's to think of it as change management that we're going to have to go through. That change management is going to have to have someone accountable, whether it's a minister, an envoy or another key resource. It is really kind of open. That needs to be different and separate from where we're seeing a lot of the attention being paid, which is around industrial strategy.

The Chair: You have 30 seconds.

Anita Vandenbeld: Do the others want to comment on that?

It's okay. I'll give up my 30 seconds.

The Chair: Thank you so much.

Fenwick McKelvey: Thank you.

[*Translation*]

The Chair: Ms. Normandin, you have six minutes, please.

Christine Normandin: Thank you very much.

I thank all the witnesses.

Mr. Lymburner and Mr. Ball, you spoke to the importance of public trust in the translation systems. They need to be reliable. You also mentioned that you are training new language specialists to work on these systems and ensure a certain level of reliability: They will not be correcting the artificial intelligence, but they will be ensuring its reliability.

Of all the translations done by GCtranslate, what percentage was double-checked by humans as part of the workflow? Can you give me an estimate?

Jean-François Lymburner: One important point that was mentioned, and I believe it's important, is that the translation bureau does not retain translations produced by clients. Clients are informed about the tool. What we do is train the tool. We receive feedback and comments from our clients. We conduct tests on a regular basis. We tested the tool and compared it to other common tools that people were using. The translation bureau's data is the best English and French data in the world. This allowed us to see that the tool performed as well as, or, in many cases, better than, so-called commercial tools. That's the first point.

I've heard some people mention errors involving proper nouns, cities, or things like that. We've corrected them. Correcting them can be a little tricky in certain contexts. As I was telling you, this morning I did a test before coming here. The sample I used was as follows: "*Mon père, Luc, vient de Saint-Jean, et je travaille au Bureau de la traduction*." The translation was: "My father, Luc vient de Saint-Jean Québec, and I work at the translation bureau." We talked earlier about human intervention, about people stepping in when necessary. That allows us to run tests and make decisions.

We talked about this earlier, and I'd like to come back to it. Every time people log into GCtranslate, they see a warning that artificial intelligence makes mistakes. I think all of your witnesses said that artificial intelligence can make mistakes. Artificial intelligence can make mistakes with confidence, too. In terms of language, we aim to produce a translation that matches the original text. Often, when errors were reported to us, the way the original text was written led to certain errors. That doesn't mean it's perfect. We'll continue to improve the system, but we're very proud of our performance so far.

● (1230)

Christine Normandin: Just to be clear: I understand that there is no ad hoc review of the translations conducted after the fact. You rely on the clients who use the translations, but obviously the clients are using the translations because they aren't perfectly bilingual and aren't translators themselves. This adds an even higher level of complexity.

Of the 125 million words translated by GCtranslate, how many were proofread?

Jean-François Lymburner: We have now translated nearly 800 million words using GCtranslate. We also periodically retrain the model in collaboration with our clients.

My colleague, Mr. Ball, can talk about the monitoring we do using the tool that we also use with professionals.

Matthew Ball: We have several ways to provide feedback. Unlike standard commercial products, GCtranslate allows users to contact a translation bureau professional directly through the “Ask a Translator” option. This is a service that doesn't exist anywhere else in the world. So anyone who wants to speak with a translator has direct access at any time.

As for proofreading, as Mr. Lymburner just mentioned, in principle, the data and information in GCtranslate are not proofread, but I can assure you that all translations sent to language professionals, the bureau's translators, are proofread by humans. Even when we use translation memories, terminology databases and neural machine translation, or NMT, behind the scenes, everything is reviewed by a professional.

Christine Normandin: Once again, I understand that this is done almost exclusively at the user's request. The content produced by GCtranslate isn't systematically reviewed. Is that correct?

Matthew Ball: The text belongs to the user. We do not own users' documents or data. It is up to users to decide whether they want to review them. What we do is display the warning. We also offer this service for free. Anyone who wants to can access a professional.

Christine Normandin: Given that the government plans to cut funding for the translation bureau by 25% over five years, how can we ensure that service quality will be maintained and that people can be confident that there will always be a human translator available to do the work if they request it? I understand that they must specifically request it.

How can you ensure this level of service, especially for francophones? We know that between 10% and 30% of translations are from English to French, not the other way around. How might budget cuts affect your work?

Jean-François Lymburner: I don't want to talk about cutbacks here. As I mentioned earlier, since around 2020, demand for traditional translation has been declining worldwide, and this is true for all translation service providers.

The translation bureau is a special service agency. Aside from the services we provide to Parliament, we operate on a cost-recovery basis with government departments, and it is in this area that we have seen a reduction.

As an organization, we need to develop a viable five-year plan. It is within the framework of this plan, and in light of current demand, that we must proceed with some reduction in staff through attrition. This isn't just about translators or interpreters, God knows we need them. Rather, it's everyone at the translation bureau who is affected by this situation as we work to meet the current demand.

However, through the initiatives we're discussing, we're reaching out to new public service employees and informing them about the translation bureau's services, which is bringing us new work. We're also helping certain departments implement tools for their use. So we hope to turn the tide, but it's clear that when it comes to traditional translation, there's been a decline.

As I mentioned earlier, we're no longer receiving invitations to golf tournaments like we did last week. So there's been a drop in demand, and since we operate on a cost-recovery basis, we cannot afford to lose money.

● (1235)

[*English*]

The Chair: We'll now go to Mrs. Block for five minutes, please.

Kelly Block (Carlton Trail—Eagle Creek, CPC): Thank you very much, Chair.

Thank you to our witnesses for joining us today.

I had some similar questions to my colleague's concerning the reduction in at least 339 positions at the translation bureau. Of course, it would match what you have cited in regard to a drop in client requests due to the widespread use of AI and free translation options.

I also recognize that you've talked about the pressure on our translation services due to the nature of their work and the two blocks being in the spring and in the fall.

My question in regard to some of those facts would be this: Do you believe that AI's output is proficient enough to be used in official translations, or is something being lost when AI is used to do these tasks?

Jean-François Lymburner: That is an excellent question.

Maybe I'll give a quick premise to justify why you have the translation bureau here.

In terms of AI, in your work we keep hearing about LLMs, which are large language models. I encourage everybody to spell them out. They're language models, and we're experts in language. We've been translating, automatically, weather reports since the 1970s, so we were at the forefront of using some form of intelligence back then.

A few years ago when I joined the bureau in 2024, some people said that translators will disappear. Some of those people who were telling me that were programmers using information languages like Python, Java and COBOL. That's all gone. The interaction with large language models happens through natural language, so the way we use our prompting, we're using language. The literacy is extremely key, and I think it positioned the translators and the interpreters at the boundary of AI and humans. AI, as I've heard here, is essentially a notion that keeps growing every day, but it's a notion of numbers and algorithms, which somehow we translate into language.

I think our role is key. Of course, in university they are changing their programs right now. Matthew is a linguist with two master's degrees, but he didn't do a lot of math. I come from an engineering background. We've worked with a lot of numbers and algorithms, but not a lot of language. The new linguists are going to come with probably those two notions combined, so that they can harness AI and make sure that we understand exactly what comes out of it.

Kelly Block: Thank you.

Raymond Théberge, the former official languages commissioner, has warned that the translation software "cannot...bypass linguistic obligations within the government." He also wrote a letter to the Treasury Board in March 2025, highlighting the need for verification by competent persons.

With these two things in mind, do you believe it's a mistake to cut positions at the translation bureau?

Jean-François Lymburner: Again, they're not cuts per se. Right now it's attrition.

I met Mr. Théberge as well, and he was asking a question around AI use at the translation bureau. In my conversation with him, I was clear that it was the translation that was happening outside the translation bureau that I was concerned about, and he was concerned as well.

That's why it's important to consider risk. By risk, I mean that when you get a document or you're writing or you're using AI, is there a reputational risk for you? Is there a financial risk? Is there a legal risk? All of these questions will help you to decide if you need a professional. Sometimes within your day, if you just want to understand an email or you want to work on a draft with your people, it's acceptable.

We're really trying to educate and provide enough information. Again, the next generation of linguists will work hand in hand with solution engineers, architects and data scientists, and that's a first, because they've been going to different departments in line with their school programs. Now we need to really make sure....

Somebody spoke about literacy around AI, and it's going to be key.

Kelly Block: Really quickly, is it good enough to just warn users that software like GCtranslate makes mistakes?

● (1240)

Jean-François Lymburner: I think we've mentioned this.

First of all, the website is fully accessible, and there is a button to "Ask a translator". When you press that button, we can quickly provide an assessment that says, "You're going on a website. We strongly suggest that you use a professional." At least we have a connection.

A year ago, people were still doing it. That's what we see in our data, and I see it right now, with people going pretty much outside government. They were sending information who knows where. One of your witnesses talked about a different computer elsewhere. At least now we're containing that.

Is it enough? I think we need to continue to educate, and a committee like this one is really helping on that front.

The Chair: Thank you so much.

We'll turn to Mr. Jeneroux online for five minutes, please.

Matt Jeneroux (Edmonton Riverbend, Lib.): Thank you, Mr. Chair.

Thank you, witnesses, for attending today. I'm sorry I can't be there in the room with you, but there has been some fascinating testimony in both panels.

I have two questions for you, Mr. Lymburner, following up on some of Ms. Normandin's questions.

When it comes to the long-term employment of language professionals, I think you were at the official languages committee about a year ago, testifying that AI complements the work of language professionals rather than replacing them. We heard similar testimony from Mr. Kevin Deveaux last week as well. We've heard a lot of testimony on the other aspect, about AI replacing workers.

I'm hoping that you can expand a bit more on what you see in terms of the complementing factors.

Jean-François Lymburner: If I look at the value chain of the work, when we were getting a document in the past—let's say a deck—we would separate it and send a component to be translated and then reassemble the document.

In the workflow of translation, a lot of things have been transformed. What we see right now is probably less high-volume, low-value work. The work that still comes to the bureau is of lower volume but higher value. By that, I mean that we see a significant number of words coming our way, but the number of projects is lower, which means that departments are turning to us for more complex issues.

That's not only for translators. I see that for accountants and lawyers. The work that will require a professional and an expert will need to be work for an expert, not work that can be done easily.

The other aspect is that there is a lot of manoeuvring in the information we are getting that are being automated right now. The tools that we have—the format of your PowerPoint, your Word, your Excel—are now a lot easier to maintain, and you don't have to handraulically change all of that. We're trying to find efficiencies in that area.

Matthew, do you want to add anything?

Matthew Ball: I might add that the work of the translators and all the language professionals is evolving because of AI. Jean-François referred to solution engineers. Right now we have translators and terminologists working with them. We have new job descriptions—translator technicians. These are people working on cleaning up the data, training the data and training the engines that are used in the neural machine translation tools that we use.

I'm optimistic that the role of language professionals will remain critical. In fact, I would say it's even more important because, as we know, AI solutions can really be deceptively perfect when, in fact, they're not. They do make mistakes. It takes a professional to really detect those easily and quickly. It makes different types of mistakes. We are at the top of the value chain. We're there training the programs. We have new opportunities for careers for language professionals. Also, using the tools in their daily output is very different work now.

Matt Jeneroux: I'll add that you probably think that boosting productivity would be part of that aspect as well. Maybe you can continue in that vein, if you care to.

In my last minute and a half, I want to get to the disclosure requirements that you indicated. It sounds like it's very proactive disclosure. Is there post-disclosure that occurs when GCtranslate is being used? We've had a lot of testimony around disclosure as well and what parliamentarians have been looking at. I'm curious as to your thoughts on that.

● (1245)

Matthew Ball: As I mentioned earlier, the data and the documents that clients provide to us belong to them. We're not the owner of them. At the end of the day, deputy heads, departments and individuals are responsible for the integrity in order to respect the official language requirements under the Official Languages Act.

We do allow for a feedback mechanism. Users of GCtranslate can always provide feedback and we encourage them to do so. If they don't like the output, if the name of a program is wrong, they provide that feedback and we have a team of translators who are training and correcting the machine in the background.

Disclosure is up to everyone. When clients provide documents to the translation bureau, we have a notice now when we send them the response saying that this has been verified by a human professional translator. At the end of the day, that document belongs to them and they can use it as they wish.

The Chair: Thank you so much.

We'll now turn to Madame Normandin for two and a half minutes, please.

[*Translation*]

Christine Normandin: Thank you very much.

I would like to come back to the issue of attrition as it relates to the translation bureau's expenses and revenues. There is something that strikes me as contradictory. Am I mistaken, or, despite what you've mentioned regarding lower demand for traditional translation, is there actually increasing pressure on translators in terms of the number of words they're expected to translate? Isn't there a contradiction in the fact that we're saying we need fewer translators, yet at the same time we're imposing a heavier workload on them?

Jean-François Lymburner: That's interesting. Clearly, as I mentioned, the current working conditions for translators assigned to Parliament include a system of parliamentary leave that allows them to recover from periods of higher workload. If we compare these conditions with those of our more traditional services, the word count and workload vary significantly. There are documents we receive from the Supreme Court and from our legal partners for which very few words are translated each day, because it's tedious work that requires a great deal of research. I'm not sure about the word count metric alone.

One of your colleagues talked about how we can assess risk. When he was doing a Google search, he saw that it was 90% accurate. There are more and more tools that let us see if these are documents we've seen before. Artificial intelligence designed to assist translators will tell them exactly where in the document they should focus their efforts. Once again, the absolute number of words translated is no longer really the benchmark it used to be. It's much more about the critical nature of the document. A one-page document that must be error-free and has a huge impact on the government's reputation will require more time for proofreading, whereas a simpler document….

Mr. Ball, I see that you want to add something.

Matthew Ball: We often talk about production rates. These serve as guidelines. They help us plan our work, so we have a rough idea of how many words and how many texts we can assign to translators. Ultimately, if a text is more difficult, if it requires a lot more research, or if it has special requirements, the bureau takes that into account, and we do not expect people to turn in their work faster and faster. It depends on the text.

Christine Normandin: Thank you.

[*English*]

The Chair: Mr. Calkins, you have five minutes, please.

Blaine Calkins: Thank you, Chair. I'm going to ask Mr. McKelvey questions.

I'm going to use my five minutes to ask you questions, sir, and I'm hoping you can provide answers in a written submission because you won't have time to answer all of these questions here. I would like answers to these questions to be submitted to the committee.

The first question I have deals with your thoughts on the existence of superintelligence, and what responsibilities you think we have as legislators when it comes to making them responsible for the actions of their creations?

In your co-authored study, “Recursive Power: AI Governmentality and Technofutures”, you warn of feedback loops. What does it mean from the point of view of a legislator? Do you think we're going to be stuck regurgitating the same information in these loops—infinite loops, spiral loops—of information? How could these loops potentially influence development or changes in our culture and in our society?

You've done some work with the Machine Agencies working group, making reference to machine agency and AI commoners in its mission statement. Can you just elaborate on what that means?

The last thing I want to talk to you about is the adoption of AI tools and AI assistance, which could lead to something called “algocracy”, and I think you know what that is. It's where the algorithm basically starts to determine, through AI use, where our democracy goes. How do you see this interacting vis-a-vis determinism and free will if we rely too much on artificial intelligence?

Also, I'd like to know your views on humanism and the post-humanist philosophy, because you have some dalliances in that as well.

One of the last things I want to talk about is the compounding of errors, because a true information system that is built by a programmer and a coder relies on a database that is controlled by a database administrator. Artificial intelligence uses data points, but it also uses information that's already synthesized, and we see that AI already relies on information generated by AI. Can you talk about how that will work and what you foresee if we don't correct the errors that AI makes? In terms of the compounding of errors through generational use of AI, where does that take society, and how will that affect decision-makers like parliamentarians?

● (1250)

The Chair: Before you answer—this won't cut into Mr. Calkins' time—we will send you a transcript of that. Usually professors don't get assigned homework, but if you wish to respond in writing, we will provide you with those. Please go ahead with your answer.

Fenwick McKelvey: Thank you for the questions.

First off, one of the critical challenges in doing the work around understanding the state of artificial intelligence and questions like superintelligence is the problem of hype. In a lot of my work, we're trying to figure out the state of the art. Some of the things that I think are critical for legislators like yourselves are having the resources and having unfettered access to understand what's going on.

Something that I really want to be sympathetic about is that the term “superintelligence” is something that could just as much be cooked up by AI firms trying to raise their stock valuations as it is something that is a genuine concern. The problem of hype is a profound problem because, as we're trying to think about the impacts of technology, very few people are participating in setting these expectations for the future. That's something that comes up time and time again in some of the research we've done around media coverage and also around who's an “AI expert”. I want to be sympathetic about this problem of superintelligence.

Second, really quickly, when I talk about AI commons, I think there's a thread going through here. I worked as the co-director of the Applied AI Institute at Concordia University, where we were building tools for government, and I think there's something missing from the digital sovereignty portfolio. As far as I can tell, GCtranslate was something that was using a local model developed internally. How do we think about the problem of building an AI stack that is sovereign and autonomous in such a way that we're recognizing that it needs to be collaboratively done?

A part of it, to me, is actually attending to the difficulty in connecting the internal development around something like GCtranslate with the broader work that has to be done to build sustainable infrastructures online that require collaboration. Certainly, this wave, when you're talking about compounding errors, is due to the fact that it becomes harder and harder for one organization, one individual, one company or one institution to be able to maintain these systems. That's one part.

There are two last things to say.

First, I think there's a point to emphasize about this idea of peak data. I was on a panel yesterday with Kathy Reid from the Mozilla Data Collective, who was talking about this. One of the challenges we now have is that, more and more, we're hitting a limit around the quality of data online. This creates a potential tragedy of the commons because there is a lack of support for this type of data. There's an argument that AI is actually polluting the quality of training data, creating these model collapses where models do not work on generated data. There are real issues around the sustainability of these types of systems in the long term that we really need to factor in.

Again, this comes back to an issue of how you build capacity for government and regulators to make sense of these issues. Something that is central to academic research is knowing that we want to be able to assist and help steer these challenges but also to assess the limitations and capabilities of AI in a more pragmatic way that's outside of the influence of propping up stock valuations.

The Chair: We will now go to Madame Brière.

[*Translation*]

Mrs. Brière, you have five minutes.

Hon. Élisabeth Brière: Thank you, Mr. Chair.

Professor McKelvey, along the same lines, if artificial intelligence is used to respond to a citizen and that citizen suffers the consequences, should they be able to know that artificial intelligence was used to generate that response? Should the citizen be able to request a review if the result is biased or does not comply with the rules?

● (1255)

[*English*]

Fenwick McKelvey: First, I want to acknowledge the opportunity of using AI to better deliver services to Canadians, when done in a way that recognizes the potential harms and risks. Certainly, I want to say that when we're talking about a chatbot providing services, there's a difference between providing this around basic tax advice versus something that would be a more high-risk activity—as we've talked about in the EU. I think the context of when these tools are being developed is really important.

Second, we were talking about this loop. There's a challenge, and we see this has been reported in The Economist in the United Kingdom, and we hear this elsewhere. There's this loop where governments are using AI to process and constituents are using AI to write to the systems, and it leads to an acceleration. I think one of the missing pieces when we talk about AI adoption is how to mitigate a negative loop where there's an acceleration, where there are two AI systems talking to each other and there's no human in the loop anymore. As a professor, I'd say it's the risk where AI is grading and writing papers at the same time. This is really an existential question that I want to emphasize.

I think it's important then to ask when these systems are appropriate. When is it something where, symbolically, there's a need to be seen to be listening, in which case a human has to be there, not because it can't be automated but because automation undermines the symbolic function of that process. That, I think, is a really important piece here. Maybe we can do this, but it might not be right if we're maintaining expectations of democracy and democratic feedback.

Third, it's very important that if we're developing these tools, we're making sure that we're using open-source models, developing an independent stack and understanding where this data is coming from so that we have an audit trail and some ways of enforcing accountability. Those are the three components, I'd say, briefly.

[*Translation*]

Hon. Élisabeth Brière: Thank you.

How should the use of artificial intelligence be monitored or regulated? Do you think there should be an agency or an office to oversee the use of artificial intelligence across various government departments?

[*English*]

Fenwick McKelvey: It's a challenge because one of the changes that has happened with generative AI is that there has been an adoption shift. This was something I saw when I was co-director of the Applied AI Institute.

Initially, we were talking about specific applications of AI and a process of deployment that included checks and balances. I think this is really a key part of the tension that I see going on. If you're talking about procurement or developing AI services, you have examples from the Treasury Board. You have work that shows you how that's supposed to be done. When you start externalizing it and procuring to third parties, the accountability becomes lessened. However, I want to acknowledge that there are systems in place when we're talking about context-specific AI applications.

One of the greatest risks is that general purpose AI has positioned it as a one-size-fits-all solution and undermined that need to think about when we're deploying it, whereas in the context of GC-translate or CANChat, you're thinking about it in terms of much more specific applications.

I don't think that the answer is necessarily one organization so much as it is thinking about what the process of procurement and adoption is and ensuring that it's not something where you're taking a tool out of the box. I think a minimum is ensuring that the work of the Treasury Board and other places that have been thinking seriously about AI in procurement is a central part of these steps.

For the broader question of democratic institutions, yes, I do actually think that what we need to have is a place for the problems that AI is going to pose for things like consultations. What does it mean when all the letters are AI generated, and how do we understand these things? I think there are these critical questions about the process of government and the process of democratic participation that are going to be challenged in an interesting way with AI, and we need to resource that. That's beyond what we've done in the AI strategy—the all-in AI strategy—or what ISED is doing around AI.

I do actually think there is a place for parts of this, in collaboration with other ministries and departments, for these broader, more existential questions.

● (1300)

The Chair: Thank you so much.

I'd like to thank our witnesses for appearing here today.

Seeing no further business, the committee is 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