HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

45th PARLIAMENT, 1st SESSION

Standing Committee on Procedure and House Affairs

EVIDENCE

NUMBER 047

Tuesday, November 17, 2026

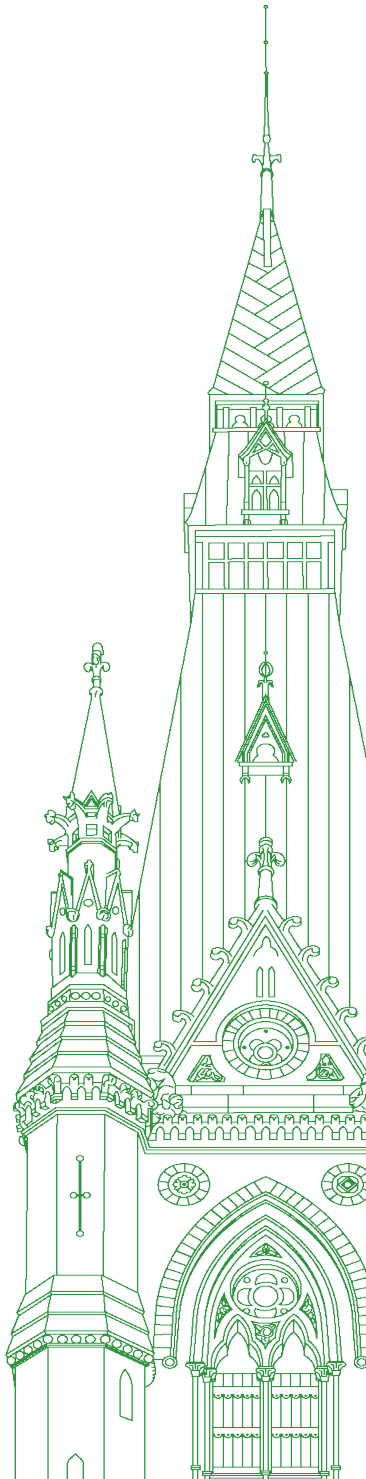

Chair: Chris Bittle

Standing Committee on Procedure and House Affairs

Tuesday, November 17, 2026

● (1130)

[*English*]

The Chair (Chris Bittle (St. Catharines, Lib.)): I call this meeting to order.

Welcome to meeting number 47 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. The meeting is taking place in public in a hybrid format pursuant to the Standing Orders.

As always, I would ask everyone to consult the guidelines written on the cards at the table to help prevent audio and feedback incidents. I remind all witnesses that committee members may ask questions in either English or French, so be prepared for that with the earpiece and be tuned to the appropriate channel. As always, all comments should be addressed through the chair, and whether you're on Zoom or in the room, raise your hand, and the clerk and I will attempt to manage a speaking list.

I would like to welcome our witnesses for the first hour.

From the Samara Centre for Democracy, we have Dr. Beatrice Wayne, director of research and policy.

We also have, appearing as individuals, Antoine Lemor, researcher in public policy and computational social science, and Archie Hall, acting economics editor, The Economist. He is joining us by video conference.

We'll start with Dr. Wayne for five minutes, please.

Beatrice Wayne (Director of Research and Policy, Samara Centre for Democracy): Thank you very much for having me here today.

My name is Beatrice Wayne, and I'm the director of research and policy at the Samara Centre for Democracy. We're Canada's leading non-partisan charity dedicated to strengthening Canadian democratic culture through research and public engagement.

I'm here with two recommendations for how we can best address the diverse and dramatic ways AI technologies are challenging our ability to maintain a trusted information ecosystem, which is a crucial building block for Canada's democratic institutions. These recommendations are drawn from the Samara Centre's Verified project, which monitors online political conversation in Canada to identify democratic threats within our digital information environment. We are currently writing a peer-reviewed book based on this research.

My first recommendation is to enhance the digital media literacy of legislators. In over 160 interviews through the Samara Centre's MP exit interview project, we have heard from former legislators who speak about their desire for much stronger support and training in navigating our digital media ecosystem. The introduction of AI tools within government makes the necessity of robust digital media literacy an urgent matter. Legislators need informed, evidence-based support to understand not only how to use different AI tools but, crucially, how to evaluate these tools and to ask critical questions about them. This is vital, because ultimately it is elected officials who will be accountable for the decisions made by these AI models.

Members of Parliament, their staff, and public servants need to be equipped to capably answer questions about why certain AI models were used for analysis. They need to understand what data these models have been trained on and how this data was acquired. Legislators do not need to be experts in the code, architecture or weights of AI models being used, but they need to know whether this information is publicly accessible. They should have a clear understanding of how bias is baked into training data and mitigation strategies for this bias.

It's crucially important for public trust, for example, to be able to answer why the government has chosen Copilot for its pilot project. That includes explaining how Microsoft has addressed the issues with racial and gender bias demonstrated in recent research. It also means having a clear understanding of the implications of the U.S. CLOUD Act for the security of Canadians' data and the digital sovereignty and national security implications of incorporating Copilot into Canada's public service.

Our second recommendation is to establish a standing AI mini-public with rotating citizenship membership. This would offer a non-partisan and meaningful approach to public participation in AI governance across Canada, and would demonstrate to the Canadian public that their perspective and input on the issue of AI governance is valued. It would create a channel for direct communication between policy-makers and members of the public. We advise that it be structured as an independent, evidence-based standing body created by statute with a secretariat reporting to Parliament. This would offer several benefits. It would insulate the institution from political winds, it would ensure democratic expertise exists within the general public at all times, and it would constantly build up a cohort of citizen experts that then return to the general public when they cycle out.

This body would build public trust, demystify a complex issue and ensure that the public's concerns are heard, because AI's impact is and will be multi-faceted. This would not replace legislators but, rather, empower them to legislate and help them better represent their constituents. A standing body would ensure that we have citizens ready to respond to any rapidly developing situation, permitting the government to be actively responsive to this emerging technology.

In addition, we strongly support the testimony presented by our colleagues about the need for platform accountability legislation that scopes in major AI platforms and crucially holds all major digital platforms accountable for their AI features.

Our recommendations are aimed at strengthening and evolving our democratic institutions from the inside out, and we would be pleased to offer our support for potential next steps as the committee sees fit.

This is a crucial moment for securing Canadian sovereignty and culture, and we can respond by enhancing our democratic capacity to meet the needs of this moment.

Thank you very much for your time.

The Chair: Thank you so much.

We'll now go to Mr. Lemor for five minutes, please.

[*Translation*]

Antoine Lemor (Researcher in Public Policy and Computational Social Science, As an Individual): Thank you for inviting me to testify today.

In recent years, I have devoted most of my time to studying public policy and public administration, as well as to the development of computational methods for social science research—that is, the development of scientific methods that draw in particular on artificial intelligence. This dual perspective enables me, as a political scientist and a researcher in computational methods, to draw conclusions from the research that I hope will be of interest.

Artificial intelligence tends to be viewed as a technology that, as a technology, is fundamentally different from our democratic institutions—something external to them, from which our institutions can be clearly separated. However, the social sciences have long considered institutions such as the state, Parliament or the bureaucracy to be social technologies whose purpose is to process information in order to support collective decision-making. Indeed, what are studies such as those conducted by this committee, if not social technologies aimed at producing the information necessary for decision-making? From this perspective, AI fulfills the same functions as most of our institutions: gathering, processing, producing and analyzing vast amounts of information for human decision-making.

This leads to a first general observation: Artificial intelligence is not a technology that is independent of or external to government activity. It permeates every stage of the legislative and regulatory process, as it performs similar functions by now generating a significant portion of the information on which governments rely. This observation is supported by research that establishes more concrete facts, which can be categorized according to the standard legislative and regulatory process—for example, in three stages.

First, in the stage preceding any public policy design—during simple advising of decision-makers, such as parliamentarians or ministers—a Canadian study by colleagues at the École nationale d'administration publique shows that briefing notes produced by artificial intelligence are now considered to be of comparable quality to those produced by humans. A number of other studies show that those who use AI adopt its recommendations and that, in 86% of cases, experts follow the choices it proposes while presenting them as their own. This finding comes from a study by Anthropic. In other words, the opinion of a senior official who relies on artificial intelligence can be influenced by it, and so can the public policy that the official helps shape.

Second, at the legislative and deliberative process stage, research shows that AI is better than the deliberators themselves at identifying points of consensus within an assembly, but also that texts partially generated by artificial intelligence are becoming increasingly common in the British Parliament and the Swedish Parliament, and that elected officials respond almost as often to emails written by artificial intelligence as they do to those from citizens. What's more, some international experiments, such as the one conducted in the United Arab Emirates, demonstrate the inclusion of artificial intelligence in the legislative process itself, which risks creating legislative inflation. This would automatically diminish the value of the law.

Third, at the stage of the regulatory process and the implementation of public policies—that is, once they have been designed—which includes certain public consultations such as impact assessments when it comes to submitting feedback, research shows that, while artificial intelligence does improve citizen feedback, of course, it also leads to an uncontrolled increase in requests that government agencies struggle to address. They respond either by increasing their use of artificial intelligence or by imposing fees and restrictions that limit access. This finding comes from a study by Schmitz and colleagues. The research also shows that different models can interpret citizens' contributions in completely different ways and summarize them less accurately depending on their social status—for example, when dealing with a street vendor rather than a financial analyst. In other words, the number of public consultations—or, at any rate, citizen contributions—is increasing, but this can paralyze government agencies, reduce access to consultations and skew decision-makers' interpretations.

These findings reveal a strong ambivalence regarding the use of artificial intelligence in the legislative and regulatory process. AI provides better access to institutions. A recent study published in Nature Human Behaviour shows that it increases the chances of success for complaints filed by disadvantaged individuals. However, by overwhelming institutions with information, it risks pushing decision-makers to rely on their established and personal networks and, as a result, increase inequalities in access. Above all, the regulatory framework appears underdeveloped. The current federal directive applies only to algorithmic decisions concerning an individual, and not to the mere information transmitted to decision-makers during the stages of the process I just described.

This leads to two proposals. On the one hand, since artificial intelligence is inseparable from government activity, the goal is not to ban it, but to require that its use—for example, in consultations—follow a published methodology that complies with scientific standards. On the other hand, rather than reinventing the wheel, we can build on what already exists. For example, the Auditor General—who has the authority and has already conducted an audit of the Canada Revenue Agency's, or CRA, chatbot—could be provided with a team of experts dedicated to overseeing artificial intelligence in accordance with these standards.

Artificial intelligence now generates some of the information on which governments rely. Its use still needs to be regulated.

Thank you.

● (1135)

The Chair: Thank you very much.

[*English*]

I apologize to the witnesses in the room, but we'll turn to Mr. Hall for five minutes.

We appreciate your patience, especially being in a different time zone. Thank you for sticking with us.

You have five minutes, please.

Archie Hall (Acting Economics Editor, The Economist, As an Individual): Mr. Chair, vice-chairs and members of the committee, thank you for inviting me to discuss artificial intelligence in Canada's democratic institutions.

My name is Archie Hall. I'm the acting economics editor at The Economist newspaper.

I come to the topic of Canada's democratic institutions, I confess, as something of an outsider, so I plan to confine my remarks here to discussing my work on the use of AI in political writing. Of course, I'm happy to answer whatever questions are of interest to the committee.

Like everybody, politicians are using AI more to draft policy papers, social media posts, and the area I've looked at most closely, political speeches. In an article that we published in The Economist this past September, I used the industry standard AI detection tool Pangram to scan parliamentary records across English-speaking countries. Nearly universally, the share of AI-written speeches as flagged by this tool has risen sharply over the past few years. Through both my own testing and third party audits, I've come to be confident of the reliability of Pangram as a tool for this task, and we can discuss why in more detail later, if that is of interest.

The tool, of course, cannot easily distinguish between a piece that was written top to bottom with AI versus one where the technology was just used extensively for something like line edits. On the results, Canadian politicians, for better or worse, seem to be at the forefront of adopting AI in their speeches. I found that around 20% of what was said in English in Canada's Parliament over the past year was flagged as AI-drafted. Pangram is tuned fairly conservatively, so if anything, that is likely to be an undercount of the true figure. Since AI use is rapidly increasing, today's figures may well be even higher.

I deliberately chose not to look at French-language contributions, since there is not yet good evidence on the sensitivity of AI detection tools and how that varies across languages.

The only politicians that I found to use AI more in their speech writing than Canadian MPs were, marginally, Australian ones. Otherwise, Canada's rate was the highest, triple that of Britain's Parliament and higher than those of both the U.S. House and the U.S. Senate. The inflection point for AI adoption, both in Canada and elsewhere, seems to usually be the most recent election. Freshly elected MPs appear to use AI at a notably higher rate.

Hearteningly, Canada does seem to be taking this question seriously. I worked with reporters at a number of Canadian outlets, including The Globe and Mail, on follow-up stories far more than any other country that I covered. The committee's interest in this issue speaks to that, too.

Is AI use a problem here? One could argue not. Politicians, after all, have long used human speech-writers. Artificial ones, perhaps, are no different. Certainly, societies have yet to sort out their norms on these things. Many, though not all, of the British MPs I spoke to for my piece were quite sheepish about their own AI use. The polling that I commissioned of British voters found that across a wide range of jobs—lawyers, journalists, assistants, teachers, doctors, and more—voters were least enthusiastic about politicians using AI to write at work, with a net approval for that below minus 30.

What can be done? If politicians are using AI—and for many tasks like research, fact-checking, and so on, they absolutely should be—Parliament should at least make sure they're using frontier AI models, something I was surprised to find was often not the case in the British Parliament. Otherwise, one norm to promote might be around disclosing the use of AI in writing. Another option, which the blogging platform Substack now does, would be to run a tool like Pangram automatically over all parliamentary contributions in the spirit of offering voters information about their representatives.

To put my cards briefly on the table, personally, I do think this is something to take seriously. Political life relies on sustained arguments—Parliament originally meant talking shop, after all—and outsourcing that work to a machine feels like an abrogation of a politician's responsibility to do their own thinking and say what they really mean.

Thank you for your time.

● (1140)

The Chair: Thank you so much.

Since we are behind schedule, and in an effort to be more efficient—and I thank Mr. Calkins for the suggestion—we'll hear from the Library of Parliament, and then we'll have the opportunity to examine all the witnesses after their five minutes.

From the Library of Parliament, we have Christine Ivory, parliamentary librarian; Marcus Pistor, director general; and Johanna Smith, also director general.

You have five minutes, please.

Christine Ivory (Parliamentary Librarian, Library of Parliament): Mr. Chair and honourable members of the committee, thank you for this opportunity to appear before you and to discuss how artificial intelligence affects the Library of Parliament.

As has been mentioned, I'm here with Johanna Smith, who leads the library and research services, and Marcus Pistor, who heads visitor and corporate services.

[*Translation*]

As your committee has noted throughout this study, artificial intelligence is evolving rapidly and changing how information is created, accessed, analyzed and shared. It affects how parliamentarians work, how public institutions operate and how policy is developed.

For us at the Library of Parliament, these shifts bring both opportunities and challenges. As Parliament's trusted source of non-partisan research and analysis, we must adapt to a changing information environment while staying anchored in our mandate.

● (1145)

[*English*]

Our mandate has not changed. It's to provide parliamentarians with accurate, timely, objective and reliable information and analysis. In a world where content can be generated instantly and at massive scale, this role matters more than ever.

Today I'll touch briefly on four areas: how the library uses AI, how AI is changing the requests that we are receiving from parliamentarians, AI as a public policy issue and the growing challenge of misinformation.

First, regarding our own use of AI, we see it as a productivity tool that supports our work, but that can never replace professional judgment or expertise.

We have taken a very structured and cautious approach. Early on, we developed guidelines for the responsible use of generative AI that are built on principles of accountability, transparency, privacy, fairness and relevance. Those guidelines are updated every few months, given the speed of change.

We've also developed an institutional road map. We're testing specific use cases and working with parliamentary partners here in Canada and abroad, as well as with cybersecurity experts to build institutional awareness, capacity and guardrails as required.

Our experience to date suggests that AI can help us with administrative and research support tasks like finding sources, narrowing topics, brainstorming and summarizing materials. That said, subject matter experts must review and verify absolutely every output.

AI can assist us with the research process, but it is not an authoritative source. Human input and oversight remain essential to ensure our products and services are accurate and credible.

[*Translation*]

Second, AI is changing how parliamentarians seek and use information. Our librarians and analysts are seeing requests that are broader, more complex and on shorter timelines. Because AI generates large volumes of questions and text quickly, people often expect high-quality research at the same speed.

At the same time, we receive more requests to verify information or assess AI-generated material. Here, our role goes beyond fact-checking single claims. We provide independent analysis grounded in reliable sources and informed professional judgment. This is where the library's value stands out. Parliamentarians have access to more information than ever. The challenge is no longer finding information, but knowing what to trust, which evidence matters most and what conclusions to draw.

[*English*]

Third, as you well know, AI is now a major public policy issue that touches nearly every aspect of legislative work. Questions about AI are regularly arising in debates and committee work in a number of different areas: copyright, labour markets, health care, public safety and national security. I could name them all, really. Our analysts actively monitor these technological developments and policy responses in Canada and internationally.

It has created a new demand for them to be aware and on top of AI through a policy lens. This growing demand reinforces that need for expert parliamentary research capacity on AI itself. Parliamentarians need balanced, contextualized analysis on AI to understand competing views, assess impacts on Canada and make informed decisions.

Finally, I just want to touch on the link between AI and misinformation. Obviously, misinformation is nothing new, but AI greatly increases its scale, speed and sophistication. Convincing dupes in text and images, and other fake content can now be created and spread with unprecedented ease. In this environment, trust and knowledge institutions play an increasingly vital role.

What we're seeing at the library is that, both in Parliament and the public, people are looking to library content for verified content, interpretation that is reliable and evidence-based analysis. We are mindful that AI systems are trained on data that often contain inherent biases and can narrow the range of perspectives in the results. This is why it is even more important for the library's research to include differing viewpoints and consider the experiences of minority and marginalized communities.

[*Translation*]

The Library of Parliament is committed to embracing AI's opportunities while managing its risks responsibly. If anything, the rise of AI reinforces that mandate and makes it more critical than ever. As information becomes easier to generate, Parliament's need for reliable expertise, sound judgment and trusted institutions only grows.

Thank you for your attention. We would be pleased to answer the committee's questions.

[*English*]

The Chair: Thank you so much.

We'll turn to Mr. Calkins for six minutes, please.

● (1150)

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

I'm going to direct my questions to Mr. Hall.

Mr. Hall, in your opening remarks, you said that approximately 20% of the English words spoken by Canadian MPs in the House were flagged as artificial intelligence-assisted.

Could you elaborate on what you think the dangers might be when it comes to the use of these synthetically manufactured words when it comes to our politics and more broadly throughout our society?

Archie Hall: It depends absolutely on the models being used. People are using worse AI models, and I think there are serious risks of hallucinations and literally getting things wrong. Those are fairly diminished with frontier models, but the risk remains that the thinking is done much less by the person doing the writing than would be the case if they wrote it themselves or, I'd argue, even if they were working with a human speech-writer they knew well. There's a danger of subtle drift where politicians are saying things that are much more shaped than they realize by the underlying AI models they're using.

Blaine Calkins: That begs my next question. In the use of the English language and in politics, Orwell warned of the debasement of language and its consequences on political thought. Are you concerned about how AI is actually changing political discourse? How do you measure that drift you just mentioned?

Archie Hall: It's very difficult to measure.

Many of the politicians I spoke to said some version of “I fed some notes into the AI model. It polished them into a nicer speech, but really they're my ideas, so what's the difference?” There's some truth to that, but there's also some sense in which actually that act of taking rough ideas and forming them into something clear and precise in political speech is part of the job and is part of the thinking.

In terms of measuring that drift, I mean, it would be tremendously difficult to do. I'd love to think of ways to do it.

Blaine Calkins: There's a lot of conversation in this study dealing with misinformation, disinformation or AI not having access to the right data. In Canada, in Quebec most recently, actually, there was a provincial election in which virtually all the Quebec news outlets prevented AI models from accessing their information for data or for information provided to users. In fact, it has been cited now that during that election, I think, 39% of people visited a page that was just created by an enthusiast using an AI chatbot or an AI agent. That became the primary source of information for voters during that provincial election because other information wasn't available for AI to use.

How do we get that balance right? What do you do at The Economist, for example? How do you get the balance right on getting fair compensation for the information that you put out there but also making sure the information is available for public discourse through an AI system? How do you manage that in your quite substantive media outlet?

Archie Hall: It's a good question.

I've seen studies from political scientists at Stanford saying similar things about Japanese elections, where AI models were recommending a fairly fringe Communist party because it had a much more detailed website with lots more information relative to some other parties. It's a risk to be aware of. A lot of it is just looking at the sources AI provides and trying to educate voters around that.

In the longer term, the question is how to get high-quality information onto AI models. I'm appearing as an individual, so I can't speak to The Economist institutionally. Working out some means by which paywalled outlets such as ours and others are compensated to get that information onto AI platforms is almost certainly going to be, eventually, the direction of travel. Right now, definitely that pollution of low-quality links sustaining and justifying AI contributions is a pretty serious one.

Blaine Calkins: We've heard countless stories about AI data breaches caused by rogue AI agents and AI swarms, including the Australian health data breach. We know of the Hugging Face incidents that have happened as well. Do you expect this trend to continue?

Archie Hall: Yes, I do.

Blaine Calkins: U.S. airspace was shut down last week for a day. I don't think that had happened since 9/11. Depending on who you ask, and the sources providing information.... Do you believe that artificial intelligence wasn't involved in any way, shape or form in that closure?

Archie Hall: My understanding is that the current public line is that's not the case. This also did happen briefly in the U.S. in 2023, as I understand it, and these systems are pretty rickety.

The baseline of cyber-attacks going up seems like a pretty real one to reckon with.

Blaine Calkins: Governments across the west are warning their citizens to have at least 72 hours' worth of supplies in case of critical systems failures. Most of that—and I've asked this question of other witnesses before the committee—is related to systems failures due to either the potential use or incorporation of AI that goes wrong or cyber-attacks that actually shut down these critical systems.

Are you aware, do you have any information, or have you published anything that corroborates this sentiment that I'm sensing from the broader public and my constituency?

● (1155)

Archie Hall: I have not personally published anything. My colleagues and I are looking quite deeply at the question of the increased pace of cyber-attacks. The data is a little bit messy so far, but certainly, the anecdotes, as you say, are pretty compelling. I know, in the British context, there's a lot of talk about civil resilience for exactly the same reason.

Blaine Calkins: Okay.

How much time do I have, Chair?

The Chair: You have 45 seconds.

Blaine Calkins: Okay.

I'm going to ask Ms. Wayne a question.

It feels like I'm in a Batman movie.

The question I have for you is on your second recommendation. How do you perceive to deal with the issue of bias and partisanship? How would you appoint and construct this review committee, or whatever it is that you called it, in such a way that we can actually be sure, if this committee were to follow that recommendation, that it truly was bias-free and recommending the best possible information to Parliament?

The Chair: Answer very quickly, please.

Beatrice Wayne: Very quickly, we have great models of deliberative processes around the world now that have worked through this. The real key is focusing on the design and representation of this committee.

Much like a jury, it's sortition. It's by lot, but it's paying close attention to the pool to ensure that there is diverse, non-partisan representation on this committee.

We can look to Madrid. We can look to the Ostbelgien region of Belgium. We can look to France, which now has a national standing committee.

Really interestingly, Canada used to be the leader in deliberative processes in the 1980s, a global leader in it. We've moved away from this, but we can move back to focusing on where the really robust research now shows how to design a great deliberative process that addresses your very concerns.

The Chair: Thank you so much.

Ms. Vandenbeld, you have six minutes, please.

Anita Vandenbeld (Ottawa West—Nepean, Lib.): Thank you, Mr. Chair.

Thank you to all of our witnesses.

I'd like to start with you, Dr. Wayne.

Samara does a lot of very good work in terms of connecting representatives with their constituents and engaging. Some of the things that we've heard in this committee are that, as MPs, our role is unique. We are representatives and we make decisions that have widespread consequences for people. It's not like how you use AI in a private company. We're speaking for, in my case, about 100,000 voters.

We've also heard that, in democratic systems, if there isn't a policy, a very deliberate set of guidelines around how AI is used, it's not that it's not being used; it's just that it's being done on the sidelines and in ways that are not as transparent.

I wonder how, as members of Parliament, but also as an institution, we can make the use of AI responsible and allow for a level of transparency in how we are actually using it. As individual members, should we be working with constituents to publish our own individual AI policies so that constituents know how and when we're using it?

Beatrice Wayne: That is a really key question. I think the issue of transparency and communication with the public is crucial. This goes back to the question of public trust. Increasing efficiency and increasing the pace of the work that you do is not actually effective if what you're losing in the process is public trust.

One way in which you most closely interact with your constituents is often through public service. If the public service aspect of your job is not trusted by your constituents, because they don't understand the way AI is integrated into it, that's a huge democratic problem.

I also serve on the Multi-Stakeholder Forum on Open Government. I was really enthused to see the national AI register, the publication of the diverse ways that AI is being integrated into the public service. I think there's leaning into that and being very clear within that register what are the metrics of success. It's not just why AI is being used to address a problem, but how the success of the integration of the AI into the system is being evaluated. How often is it being evaluated? What is the way to get this information out to the public? It's not just the question of publishing this; it's making sure your constituents have access to it, understand it's available, and then they can really understand it. That knowledge mobilization aspect is also a crucial point.

I'd also point to some municipalities that are doing this work in really interesting ways. Elk Point, Alberta, which has a population of under 2,000, has a really robust approach to their AI policy right now that I actually think the federal government could look at and get some inspiration from. Municipalities are leading the way in this in many exciting ways.

Anita Vandenbeld: I wonder if it might be possible to get a list of some of the municipalities or other jurisdictions that are doing innovative things. I don't know how long the study is going to continue, but it might be very interesting to hear from some of the people in those jurisdictions and get some ideas for our own Parliament.

● (1200)

Beatrice Wayne: I think that would be a great idea.

Anita Vandenbeld: Good.

I also want to go to Mr. Hall.

I think what we're getting at here is that AI is going to be used, but it's a matter of how. I think the trust issue that Dr. Wayne mentioned with that gotcha.... You're using AI or you aren't.

I have a dyslexic staffer who uses AI for her correspondence. It's much easier and more inclusive for her to draft things verbally and have AI put them together. A lot of it is in the prompts.

If you have a prompt that is very much limiting the source, say, to look at just committee testimony or a particular committee, and you're very specific, and then you review it as a human, particularly as an MP.... Is it always bad that MPs are using AI?

In a world where our staff are probably doing it anyway or other parts of the institution, as we've discovered, are using it, what is the best way to ensure that we maintain that public trust while still using it responsibly?

Archie Hall: It's a very good question.

Certainly, the use of AI for the things you mentioned, like research or correspondence in cases where people have issues like dyslexia, is straightforwardly a good thing. It's important.

The broader question of trying to work out when and how AI can intrude on the process of political arguments, I think, is the difficult one. As you say, someone using AI in the way you describe, where you're very clear about the sources and you steer it very clearly on the argument you want to make, is different from someone saying, “Please write me a speech saying *x*, *y* and *z*”, and leaving it at that.

On the other hand, I think there's a sense that people shouldn't be blind to the ways in which the process of writing is part of the process of thinking. This is, perhaps, me showing my colours as a professional writer, but I think that's important, and moving away from it altogether is a challenge.

The other thing I'd add, as I mentioned, is that, at least right now, voters are quite hostile to this. That may change. They're obviously using AI themselves in their lives, but that's worth bearing in mind for politicians.

Anita Vandenbeld: I have 30 seconds, so I'll thank the witnesses.

If you could come back to us with some of those examples, it would be very helpful.

Thank you.

The Chair: Thank you so much.

[*Translation*]

Ms. Normandin, go ahead for six minutes.

Christine Normandin (Saint-Jean, BQ): Thank you very much, Mr. Chair.

I'd like to thank all the witnesses.

My first question may be for both Mr. Lemor and Ms. Wayne.

Mr. Lemor, you said that, owing to the use of artificial intelligence, we're seeing an increase in the number of requests from citizens and that this often becomes a kind of a vicious cycle. To address this, artificial intelligence is being used more, which sometimes has an unintended consequence. For example, lawmakers may end up relying more on a smaller circle of people and cutting off access to the general public.

Ms. Wayne, you somewhat answered Ms. Vandenbeld's question about the transparency needed to maintain public trust.

I'd like you to speak to us from the public's perspective. We've talked a lot about what we're doing, but we don't know how the public perceives our use of AI. I'd like to hear your thoughts on that.

Does this lead to a loss of trust?

Does it sometimes give the public the impression that, ultimately, anyone could do the job of an MP because, after all, it's artificial intelligence that's responding to the emails?

I'm very interested in your perspective on this.

Antoine Lemor: I'll have a hard time answering the question about perception and trust specifically, since I prefer to have an informed opinion based on research. However, I can say that, when it comes to public administration, citizens and citizens' use of public services, AI helps—and this is a positive aspect that has been demonstrated by research—to facilitate access to public services, for example by increasing citizens' chances of success with their requests.

A study published in the journal Nature was released not too long ago. It shows that people who use AI and have difficulty writing, for example—as was just mentioned—or who live in geographic areas facing social challenges have better access to public services and a better chance of submitting their request. From that perspective, AI can help citizens gain better access to the government, public services or their representatives. We can imagine, for example, that using an AI application could help them write certain requests addressed to their MP. That's what the research shows.

In terms of the public's perception of artificial intelligence, we can draw conclusions about the potential negative consequences for public services. A study by Schmitz and colleagues was published not long ago. Approximately 60 cases were identified in which public services were flooded with requests generated by artificial intelligence. Paradoxically, this creates problems with access to those very same public services. Ultimately, it prevents citizens from accessing them. As for the impact on our fellow citizens, the research shows that the situation is very ambivalent.

● (1205)

Christine Normandin: Ms. Wayne, what are your thoughts on this?

[*English*]

Beatrice Wayne: Absolutely, the question of public trust is key. I'll go back to something that Archie Hall said. The public opinion of AI is increasingly eroding. That confidence in it is eroding. We have poll after poll from Angus Reid Institute and Abacus Data showing that the concerns about AI usage are rising even as usage increases amongst the general population. Transparency is key, but public education also is key.

I spoke about the importance of digital media literacy for legislators, but the Samara Centre has long called for a national digital media literacy strategy for the general population. This is only more important now. If the general population has the ability to understand the different nuances of these tools, particularly—and this is something Antoine and I were talking about earlier—the many different technologies that fall under the bucket of AI.... LLMs are not the same as machine learning. They have different outputs and requirements, and they shape in different ways the way a legislator might use them. Therefore, if we can increase general knowledge about these tools too, I think there will be more comfort in seeing that there are specific ways in which the integration of AI systems into our public service is an okay thing.

Again, there are a lot of outstanding questions that need to be answered, particularly in terms of the question of data sovereignty. Being able to transparently communicate to Canadians about how the different AI systems being used are protective of Canadians' data is crucially important.

[*Translation*]

Christine Normandin: Thank you very much.

Ms. Ivory, I'd like you to tell me about the impact this situation is having on the Library of Parliament's internal resources. You are receiving more and more requests from members because they themselves are receiving more and more requests.

How do you manage your internal resources?

Do you have to adjust your priorities for processing certain requests?

Do you have enough resources to handle all these requests?

As we know, the Library of Parliament is recognized for the excellence of its work.

Christine Ivory: Thank you for the question. I'll start answering it, and my colleague may want to add comments afterwards.

Even before the advent of artificial intelligence, we were already seeing an increase in the number and complexity of requests. It is true that we are now seeing that number rise dramatically. We are inevitably forced to adjust the time allocated to processing these requests. This may involve, for example, revising deadlines and reassessing what we can offer.

As for the use of artificial intelligence, one option would be to use it to assist us with certain administrative or preliminary tasks. This could sometimes lighten the workload of librarians and analysts. We certainly have many questions about this new workload associated with artificial intelligence. We're wondering whether this will be permanent or not.

This has been discussed on this committee. Even for the committee, the number of requests that the analysts have to deal with is increasing. We are also working with the House of Commons to see what we can do to help each other.

It certainly adds a huge burden to our work. Not to mention the impact on analysts and researchers, who now have to question all sources and verify their reliability.

[*English*]

The Chair: I apologize. I'm going to have to cut you off there. You'll have plenty of opportunity to build upon your answer.

I will turn to Mr. Cooper for five minutes, please.

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

I'm going to direct my questions to the Library of Parliament.

There was a recent CBC article with the headline, “University libraries are being swarmed by AI bots, putting public access at risk”.

Has the Library of Parliament experienced being inundated with data-scraping bots?

Christine Ivory: Thank you for the question. I'll give the beginning of the answer and let Marcus, my colleague, talk about what we've done to react to that.

The answer is yes. Because the library has a lot of publications and a lot of content that can be used for scraping to support the LLMs, we've seen this happen, and it has markedly slowed down some of our public websites, including where visitors access the reservation and ticketing system for tours. We are dealing with this and working closely with the House of Commons.

Marcus, I don't know if you want to add anything on the specifics of what we've done.

● (1210)

Marcus Pistor (Director General, Visitor and Corporate Services, Library of Parliament): We have put in some technological solutions to deal with that particular challenge. I think all public-facing organizations are facing that challenge right now.

We're also modernizing our web content management system to deal with that, but it's going to be an ongoing challenge, because any time we bring in a solution, others find new ways of getting into our systems and into our data. It doesn't affect our internal systems, but it does affect our public-facing systems.

Michael Cooper: Are the solutions working?

Marcus Pistor: Sometimes it takes a bit of time to figure out the details, which include access restrictions to certain regions where more of those requests are coming from. There are, increasingly, tools that can be used to ensure that the requests are coming from human beings rather than from bots.

Michael Cooper: Okay. Thank you for that.

I want to speak a little more broadly to the issue of cybersecurity at the library.

It was recently reported that an AI agent tried to hack Library and Archives Canada. What steps is the Library of Parliament undertaking to improve cybersecurity processes?

Marcus Pistor: All of our IT system support comes from the House of Commons. We work on House of Commons infrastructure. We work very closely with the House of Commons IT teams and IM teams to make sure that we follow most recent cybersecurity protocols. We also work with the Cyber Centre, so all of our systems and all of our hardware are tested, and we check that we don't create vulnerabilities when we onboard new systems.

That's what I would say at a high level, but we're closely collaborating with the House of Commons to ensure that all of our systems are safe.

Michael Cooper: Are you you aware of an attempted hacking incident by an AI agent?

Marcus Pistor: Not specifically, but I think cyber-attacks of all kinds—whether they're phishing attacks, AI-generated or not—are very common in public-facing institutions around the world, including here in Canada. We're part of the set of institutions that are trying to address and deal with that.

Michael Cooper: Moving on, I understand that the Library of Parliament has established a working group to better understand the benefits and risks of AI, to develop guidelines to orient library employees in the use of AI and to evaluate several tools to determine whether they can be leveraged in the work of the library. Do I have that right?

Christine Ivory: I mentioned that the first thing we did was guidelines on generative AI, because we knew that it was something everyone at the library wanted to know: “Can we...?” “How do we...?” We established clear guidelines on the use of generative AI and we continue to update them for what is within scope, what is allowed, the transparency, the fairness and all of those elements.

The community of practice, or working group, was really to be able to pilot safely. What we've added with the development of a road map is a governance overarching this. The community of practice is where people are piloting ideas and testing things, but any decision to put forward a sustained approach to AI has to go through a clear governance process.

This links back to cybersecurity so that nothing is allowed to be brought into our ecosystem without clear safeguards and vetting by our IT and House IT as well.

Michael Cooper: In terms of that road map and testing specific use cases, can you elaborate on some of the work that has been done?

Christine Ivory: Maybe I'll let Johanna talk.

The Chair: Please answer very quickly.

Johanna Smith (Director General, Library and Research Services, Library of Parliament): We've dived in a little bit on use cases for research specifically. We're finding that there are not great results there yet. The efficiency gained from the product does not equal the quality of the product, but we continue to test that.

We're also testing other applications in terms of concatenating data that is available. We've used AI to draft coding that will go out and routinely search for data that we use repeatedly to produce tools for you. We are using it from time to time in synthesizing witness committee work that we then use for reports. We're testing it in different ways, refining the prompts we might use and sharing that information within the group.

● (1215)

The Chair: Thank you so much.

We will go to Mr. Caley for five minutes, please.

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

I would just add that from that past round and the one before, I appreciated Ms. Wayne's comments about many kinds of AI. I think we're seeing a proliferation of capabilities that will have to be viewed differently in how we assess all of this.

The focus of my questions will be for Mr. Hall of The Economist.

Thank you for joining us today. While I appreciate the Conservative member's quoting of Orwell—I always will—on the risk of the debasement of language, and I agree with several of those risks, I would assess that some AI capabilities may also increasingly assist in aspects of legitimate, meaningful human communication going forward in a manner that, in at least some cases, may eventually not be all that different from how we've used spreadsheets to generate beautiful charts or graphics. I'll take you up on your earlier invitation to pose a bit of a broader question, informed by your coverage. In the spirit of your intervention, I'll disclose that I used an AI-enabled search to review some of your previous coverage.

Some of those who are building or testing frontier AI now put the odds of human extinction within a decade at near 10%. I think two witnesses during this committee's study made similar comments. Your paper, The Economist, warned that a fierce backlash against AI could be “dangerous”. However, it also said, “The world must urgently try” to head off AI's risks and ensure that we're not letting democracies lose ground to “authoritarians”.

As an economist, at what probability of any catastrophe does slowing down or vastly heightening safety controls become the rational choice, even at a potential cost to growth or an assumed lead over authoritarian hegemons or rivals? How concerned should we be in this committee that many of our peer democracies, or our own, can remain strong, healthy and sovereign enough to make that call, even when some of the most gargantuan economic incentives in human history are pointing the other way and, inherent in the hyperscalers, they're marshalled by some of the most immense concentrations of capital that humanity has ever seen?

Archie Hall: That's a very interesting question. I'd encourage you to look at a very good paper by the economist Chad Jones of Stanford, who's now at Anthropic. He has tried to size a little bit, based on how extensively one is worried about extinction risk, what the scale of public spending then might be to counter it. You can imagine also the scale of a response to slow down the AI race there. It is, as you might imagine, intensely dependent on exactly where your probabilities of various outcomes are. There is, of course, vast disagreement on that.

One point I would add that I think is often lost in these conversations is that it also depends on your view of how your knowledge improves in the future. If you're someone who is, in Silicon Valley terms, a very “fast takeoff” person, who worries that actually AI could get very good very quickly in a way in which institutions won't have time to respond, that's a very, very different theatre. That is probably more where I, and I think economists institutionally, take the view that this is a concern that's very, very serious, but also that societies move quite slowly. Diffusing the technology happens quite slowly. Insofar as societies need to respond, hopefully there will be time to do so, but it's a matter of getting prepared and calibrated in case that eventuality does arise.

Braeden Caley: Are there specific steps you've come across in your coverage that have intrigued you that you think democracies like ours should be taking to help meet that balance you've described there?

Archie Hall: Well, on the challenges you mentioned when you alluded to democracy trying to work it out, there's one suite of policy responses you can have around slowing down AI development. There's a whole range of state bills kicking around in America, with policy reports and so on, looking at different ways to do that around constraining access to compute for frontier labs, trying to regulate exactly how quickly they could advance on various benchmarks and having governments evaluate new model releases. All of that is possible.

There's obviously this question in the background that if democracies also want to stay ahead of China, and stay ahead of other potential geopolitical adversaries in this world in their use of AI, that complicates it. It complicates how much you can slow down. The policy conversation that's happening right now in the U.S. in the wake of the Nvidia news—around, for instance, making it much harder for Chinese AI companies to access chips that are in countries like Malaysia, where previously that sort of third country cloud had been allowed—I think is a step in that direction. But there's a two-way step, where even if you want to be able to try to pace, and there are policy options to do that, one also needs to work out how to do so in a way that doesn't cede democracies' lead to China or other countries.

Braeden Caley: Just quickly, I suppose, in a more traditional sense, the economic risk right now at a macro level, where we've seen a lot of volatility and a huge buildup, some have described as a bubble. Are you feeling those risks in a greater way over the past couple of weeks?

● (1220)

Archie Hall: Yes. Absolutely. We put some of those risks on our cover very recently.

The history of almost any large technology build-out like this, including railways, canals and the Internet more recently, is that eventually, even if the technology is deeply transformative and deeply productivity enhancing, you also can get an overbuild and a financial bubble.

My judgment is that the risks of a bubble are quite large, but the risks of a damaging bubble are not that large at all. A good analogy is maybe the tech bubble of 25 years ago. There was a mild recession response, but ultimately, with the vantage point of history, what we noticed was that the Internet was created, not that there were a few quarters of negative GDP growth that were pretty quickly reversed.

Braeden Caley: Thank you.

The Chair: Thank you so much.

[*Translation*]

Ms. Normandin, you have the floor for two and a half minutes.

Christine Normandin: Thank you very much, Mr. Chair.

Mr. Hall, there's a difference between having artificial intelligence prepare a speech that you will deliver in the House and then just reading it—without necessarily understanding it—and using artificial intelligence to prepare, test ideas, improve your understanding and then deliver a speech.

We heard from a witness, Nish Patel, who told us he uses artificial intelligence to prepare, create notes and get informed. Later, at a meeting with his colleagues, he and his colleagues set all that aside and discuss things together. Mr. Patel said that, “conversation is the work”.

What do you think of the idea of changing the way parliamentarians operate in the parliaments of various countries to ensure that, even if someone uses artificial intelligence, it is truly that person who is speaking? For example, we could require speakers to deliver their speeches without notes. We could also set aside more time for verbal sparring and debates among members. We could have question and answer periods that last longer than 35 seconds.

Is it possible to review our practices in Parliament to somewhat circumvent the use of artificial intelligence, which can sometimes be an easy way to get work done?

[*English*]

Archie Hall: Absolutely. I think you get to a very important dynamic. In some ways, AI heightens the contradictions of the problems that are already evident in many democracies. I know the British case better than the Canadian case, but I imagine it's similar. You discuss laws and you debate in Parliament, where you have a move towards speeches politicians are reading verbatim that are very, very scripted and that are determined in advance. You don't get a free flow of human debate.

Of course, that may already have been a problem before AI, and then the notion that these speeches are then also being written with AI maybe heightens that even more and makes the absurdity of it even more apparent. I think AI certainly accelerates the need for that.

To your prior point, certainly I don't want to come across as anything like a Luddite or an absolutist. I use AI all the time in my own research and my own work. I think anyone who's serious about trying to do the best work they can, given the level of technology, should be doing so. It involves the guardrails as you say: making sure it's informing your understanding and thinking but also not doing the thinking for you.

[*Translation*]

Christine Normandin: Thank you very much.

[*English*]

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

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

This question is for the Library of Parliament folks.

Over the years, you produce a great number of your own publications based sometimes on inquiries from members of Parliament. What is being done to protect those? I know that some of them are not publicly available, and some of them may be. If someone has asked for research on a certain topic and you've provided that, how accessible is that currently?

Christine Ivory: Your question ties back to the pressures we're facing, because the same analysts who support your work and the committees are the ones who do the publications. One of the prioritizations that are happening is that we are producing fewer publications because of the pressures of what we have to do to support Parliament, but we are committed to continuing to do publications, because there is this demand and a role for the library to provide public information. However, we have always prioritized what we will do for parliamentarians.

The publications that are not open to the public are ones that we do for parliamentarians as working papers, and they're on our internal servers and internal sites, so they are less at risk. What we're trying to do now is find a way to continue to protect our publications program for the public. It's just finding the resources to be able to do that.

I don't know if that answers how we're doing the work on publications, but to be transparent, it is the one thing that ends up being wedged out when there are pressures.

Hon. Rob Moore: That's great.

In the inquiries from members of Parliament as well as publications, how would you describe the Library of Parliament's own use of AI at this point?

● (1225)

Christine Ivory: I'll start, and Johanna, feel free to add.

I think, as Johanna mentioned earlier, what we're finding is that for in-depth research, AI is not the solution. It is actually, at this point, more work. Where we use AI is the brainstorming and maybe kicking off the research process, but it's very limited in what we do related to in-depth research, requests and publications, which is why at this point, you would not have seen anything from the Library or Parliament in terms of publications or in-depth papers that mention that we've used AI.

It's really looked at as a brainstorming tool and more for administrative or maybe cleaning up language, that kind of thing. It's because of all the things we mentioned. With so many hallucinations, the fake news and having to verify the sources, it ends up being more work. That may change with the sophistication of AI over several years, but at this point in time, its use is very limited for our research.

Hon. Rob Moore: Thank you.

Mr. Hall, one of the things I took note of—and I was actually kind of surprised when you mentioned it—is the use of AI for speeches in Canada versus Australia, the U.K. and the U.S., and that Canada was at the forefront. You mentioned election recency as being one of the factors in there. Is there anything else you attribute that to?

Give us some idea of the scale. Is this a marginal difference between Canada and the U.S. and the U.K. or was it quite substantial?

Archie Hall: The difference is quite substantial. The figure I gave for Canada was around 20%. Britain is around a third of that. The U.S., from memory—I'd have to double-check my article—was about 15% for the U.S. House, but much lower, in the order of maybe 3%, for the U.S. Senate. Canada is certainly on the high end, as is Australia, on those measures.

In terms of understanding why that is, election recency explains, for any given country, when the uptick begins. It happens later in countries that had their latest election more recently. That's what I found in the data. In terms of explaining exactly what it is about different countries that have led to that, I would guess some of it is staff resourcing. The U.S. Senate, for instance, is quite well resourced. That's probably one reason why people feel the need less to lean on AI. Certainly, in talking to British parliamentarians, a lot of them were saying that they're so stretched and this is something that helps them get things done quickly and on time. I don't know enough about the details of Canada to speculate beyond that.

Hon. Rob Moore: That's interesting.

You were talking about resources. Do you have some sense of the resources available?

I can tell you that in the Canadian context, members of Parliament usually have around two to three staff at the most in Ottawa for supporting us in our endeavours here and usually a couple in the constituencies. How does the U.S. Senate compare to that?

The Chair: Unfortunately, Mr. Moore, you're over time. I'll have to cut you off.

We'll go to Madam Brière.

[*Translation*]

You have the floor for five minutes.

[*English*]

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

[*Translation*]

Good afternoon, everyone.

Professor Lemor, it's a pleasure to have you here. You're from the Université de Sherbrooke, and we share this wonderful pride in that institution.

You favour the use of local artificial intelligence models over the use of external cloud infrastructure. This is fully in line with our commitment to ensuring our security, and data sovereignty and confidentiality.

What are the risks associated with relying on foreign services?

Could a local model be adapted for use by the government?

Antoine Lemor: Thank you for your question.

There are many risks. For example, so far, we've discussed the work of parliamentary assistants. They can help you prepare briefing notes, among other things.

Recent research by Anthropic shows that experts who use artificial intelligence defer to it when making decisions, such as choosing between different research plans or options. This means that an artificial intelligence that lacks sovereignty and that belongs to foreign interests could indeed pose risks. The risk lies in the fact that these choices could be influenced by models, and we don't know how those models were used. In any case, we don't know how the model weights were trained and how the decision-making might have been influenced. This is where the risk lies.

The more we rely on artificial intelligence for items such as briefing notes in the government, in Parliament or in public administration, the greater the risk that the choices delegated to artificial intelligence could influence decision-making regarding the public policies recommended to you. This has been demonstrated by research.

The solution may indeed lie in using open-weight models, which institutions can now host and train fairly easily. The institutions train these models based on specific objectives and for relatively simple tasks. We were talking, for example, about summaries. When it comes to summaries of committee business, a local model can easily handle the task without relying on private companies, which might try to influence the choices in these summaries. That said, biases may exist.

We've been talking about transparency so far. The advantage of open weights is that they can be audited. We know how they're trained. Above all, there are gold standards for each task. In computational social science, that's what we do. We can measure a model's effectiveness relative to a specific task and then publish those metrics to build trust among users—for example, in Parliament—or among the public at large. The method can be audited in order to validate, for instance, the generation of a summary or the categorization of certain comments made during public consultations.

From this perspective, the advantage of open weights is that, by definition, they're open. We have complete control over the training and we own the model itself. The only risks associated with open weights lie in the fact that companies still train them. A typical example is DeepCycle. Until we have thoroughly tested how these weights behave and we fully understand them, we also can't say for sure that their use is completely safe.

● (1230)

Hon. Élisabeth Brière: You spoke briefly about a bias that could be introduced in public policy.

Would this solution help to eliminate that bias, in a way?

Antoine Lemor: Yes, certainly.

Above all, it's a relatively simple solution. Again, this depends on the nature of the tasks. When it comes to complex tasks—such as research tasks involving development decisions or the selection of the most relevant studies—human judgment must always take precedence at that point.

For much simpler tasks, such as categorizing comments submitted according to topic, we can easily establish a gold standard in order to compare the model's performance against that standard and to ensure that the model can perform these types of tasks.

These tasks are really simple. They ensure complete transparency regarding the process that led to the categorization of the comments, for example.

Hon. Élisabeth Brière: We heard in your opening remarks—and in everyone else's, I believe—that artificial intelligence was here to stay. It's a productive tool that supports our work. However, we must have control over its use.

You developed the LLM Tool platform, which ensures protocol transparency. I think that this issue plays a key role in government accountability.

What specific regulations and standards should we implement to ensure that the departments that use artificial intelligence understand the framework in which they can operate?

Antoine Lemor: It may seem paradoxical, but I don't think that it's particularly difficult to establish auditable standards for assessing the categorizations made by models for certain specific tasks. Sorry, but I'm repeating myself a bit.

I'll give a concrete example. A consultation was recently held on the artificial intelligence strategy. Artificial intelligence was used for this consultation. Yet the protocol used to categorize the various topics and comments received wasn't published.

I found four federal government consultations where artificial intelligence was used to process the comments received. These were simple tasks, such as categorizing topics and extracting the points made. In these cases, we can simply rely on scientific research standards using specific gold standards for each consultation process. Using a code grid selected by humans, we compare the model's effectiveness in carrying out the requested categorization of topics, for example.

In my view, the institutions that already exist—and that have already been entrusted with oversight powers, such as the Office of the Auditor General and other institutions—are perfectly capable of carrying out these tasks. They did this for the CRA, specifically by reviewing the methods used to categorize comments, for example, and by auditing them. They simply used a team trained in computational methods that can audit these methods.

● (1235)

[*English*]

The Chair: I'm going to have to cut you off. We're way over.

We'll proceed to Ms. Block for five minutes, please.

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

I appreciate our witnesses being here today. Thank you so much.

I'm not sure where we are in this study, but it's been very interesting and most informative.

My questions will go to the Library of Parliament.

I note that in your opening comments, you observed that AI has changed how parliamentarians do research, that it touches nearly every aspect of a legislator's work and that it has in fact created a new and ever-increasing demand on the services provided to legislators. I'm assuming the Library of Parliament is no exception to that.

Can you tell me, given that testimony, are your current staffing levels high enough to ensure that the Library of Parliament can handle the increased volume of requests?

Christine Ivory: I think the honest answer is, ask any head of an organization and they'll tell you that they probably need more resources. As I mentioned, there was an evolution in the demands for research even before AI, and that was also because of the pace of work here at Parliament. The number of committees has increased; there are some special committees. There are other factors that are influencing how we look at what we need.

I will say that we are probably looking to ask for additional resources because of the increase and the change in the demands. Really, it's the complexity that comes from.... A lot of the more basic information requests are being handled by the staffers, but then the ability of AI means that we may get a question that comes from 10 or 20 prompts that someone in your office has gone through to ask us a much more complex question, and just deciphering and getting to the nub of the research question takes more time.

The short answer is, we think we do need to look at how we are going to manage this new influx, yes.

Kelly Block: I appreciate that. I know my colleague, in his questioning, was talking about the amount of resources it takes to not only do as you've said but also verify the accuracy of some of the information you may be asking to comment on.

I want to ask about the digitization that may be taking place within the Library of Parliament. Can you tell us how much of that digitization has taken place within the Library of Parliament?

Johanna Smith: We don't have a huge historical collection, like you might find at Library and Archives Canada, where we're converting massive amounts to digital form. Where we focus is on the historical parliamentary publications. In particular, we're looking at sessional papers right now, to make the historical sessional papers as accessible as we can online.

We have a very small digitization team. We have two clerks who do this work. They're responding to ad hoc requests for paper materials that we digitize on demand, but then in the rest of their time, they're working to make sure that we can convert as many of our paper records to digital as possible. Most of those are in printed form, so they can be converted easily-ish into characters that can be recognized. We have been experimenting with using AI, actually, in that project to do basic catalogue records so that they're findable through our catalogue. We're finding that AI is not quite there for that work, particularly in a bilingual environment.

That's a bit of a rundown of our digitization picture.

Kelly Block: I imagine that work may take even more resources than, potentially, you would have available to you right now.

Johanna Smith: Certainly, yes. It's a scope and speed question. The faster we want to get those out, the more people we would need to put on it.

● (1240)

Kelly Block: My colleague Mr. Cooper referenced the recent article about the swarm of AI bots on the libraries at universities and talked a little bit about cybersecurity. We know it was also recently reported that an AI agent tried to hack Library and Archives.

Do you have any concerns that the Library of Parliament could be targeted?

The Chair: Please respond very quickly.

Christine Ivory: We have a lot of faith in the House of Commons IT, and all of us, I think, can be put at ease by how much they work on this.

There is a targeting of public institutions, of libraries—Toronto Public was also attacked—and universities, and I think there are reasons, given the breadth of information they have. They also have accounts of public users because of fees.

The Chair: I'm going to have to cut you off. My apologies. I'm sorry.

I'll turn to Madam Kayabaga for five minutes, please.

Hon. Arielle Kayabaga (London West, Lib.): Thank you.

Many of you have already mentioned this, but I want to start with recommendations. If you could recommend two or three things that you think Parliament should do before fully adopting AI, I would appreciate hearing them.

[*Translation*]

I also want to hear from you, Mr. Lemor.

[*English*]

Then I'll get into some questions for the Library of Parliament.

Please, Madam Wayne, you can start.

Beatrice Wayne: One key thing that I'll return to from my comments is really evaluating the sovereignty implications of adopting U.S.-based large language models, LLMs. I think we have not necessarily.... The public at least is not fully aware of why Copilot has been chosen, and I think the implications of the CLOUD Act, passed in 2018, are really at the forefront of people's minds. Particularly, it's interesting that Microsoft France's legal director, in June 2025, when asked by the French Senate whether he could confirm that French data was protected from U.S. authorities, said no. He said there was no way. Even if the data was stored in France, he couldn't confirm that. I think that is an active concern, and really understanding that is a big recommendation.

The other recommendation, back to my point, is finding a way to build public trust, different forms of public participation that can repair the relationship, which right now I think is very frayed. A previous colleague of mine, Fenwick McKelvey, commented on being involved in a lot of consultations around that but not necessarily feeling like there was an uptake on it. Finding ways to have really meaningful public participation in this conversation is crucial.

Hon. Arielle Kayabaga: AI use in Parliament, generally speaking, or just—

Beatrice Wayne: I think both, but particularly in Parliament.

Hon. Arielle Kayabaga: Okay.

[*Translation*]

Mr. Lemor, go ahead.

Antoine Lemor: With regard to the use of artificial intelligence to process input from the public, I believe that the first step would involve equipping the current oversight organizations with teams tasked with identifying all projects that use artificial intelligence to process information concerning the public consultations. The next step would require establishing the benchmarks and standards used to assess these models. This could involve, for example, categorizing comments received during consultations. There must be standards in place to monitor the use of these models.

These models can help us handle a large volume of public feedback here in Parliament or in federal government institutions, such as the executive branch. These models can be audited. However, the initiatives that we see are one-off experiments. There doesn't seem to be a unit that centralizes all these experiments. Yet from a research and computational social science perspective, we can really monitor these artificial intelligence applications.

This would be my top recommendation. Moreover, it would help increase the government's ability to monitor the use of artificial intelligence in an environment that really lacks legislation. This is simply a matter of processing information. It isn't an algorithmic decision. It would also help build public trust in the government and in how this technology is used, since the methods could be made public.

Lastly, this would give us control over the models that we ourselves could also use—particularly open-weight models—to process these requests.

● (1245)

[*English*]

Hon. Arielle Kayabaga: Mr. Hall.

Archie Hall: I will just add, on top of this, that there is a two-sided trade-off in contexts like this between using the most frontier models, which are closed-source models that are generally run by American companies and may come with sovereignty risks and so on, and open-source ones.

I think it's worth bearing in mind that the capabilities gap is still especially large. It may well be that, for many cases, the trade-off is actually worth going for the frontier models over those of which you have more assuredness about sovereignty, data retention and so on.

Hon. Arielle Kayabaga: Thank you.

Through you, Chair, maybe to the clerk, could we at some point find out how Copilot was selected as what we're currently using?

The Chair: That's through us and our investigation. My understanding, and I was going to say this at the end, is that there is a desire for more meetings on the subject, and that even though we are shifting into a different subject on Thursday, we will come back to this.

Perhaps it is a good opportunity for members to submit witnesses. I know there is some interest in calling AI companies, keeping in mind we may have to go the route of a summons. Having dealt with American tech giants in the past, I know they don't frequently come of their own free will.

Hon. Arielle Kayabaga: I don't know when we are going to get to that, but in the meantime, I am interested to hear from our officers on how Copilot was selected as a method of....

They can also submit it in writing.

The Chair: My understanding is that we will be calling back House administration. That is probably the cleanest way to deal with it, rather than an ad hoc request.

Hon. Arielle Kayabaga: Perfect. Thank you.

Do I still have time?

The Chair: No.

[*Translation*]

Ms. Normandin, you have the floor for three minutes.

Christine Normandin: Thank you.

Ms. Wayne, in your opening remarks, you mentioned something that sounded like a recommendation calling for the enhancement of legislators' digital literacy.

I would like you to talk about the current situation.

Are we doing decently? Do we just need a refresher, or have we passed the point of no return? Are we beyond redemption?

[*English*]

Beatrice Wayne: No. It certainly varies. This particular committee, I think, is more informed than maybe the larger body of MPs. With the general population, we have not had access to really robust digital media literacy. It's something that can always be shored up, but I think particularly conversations with ISED and conversations with other organizations have underlined that for particular MPs, the understanding of how AI systems work is somewhat low.

Even if it isn't, I just think MPs deserve to be supported in this. There's a famous line from an IBM training manual from 1979 which says, “A computer can never be held accountable, therefore a computer must never make a management decision.” This is not computers making management decisions, but it is MPs who are ultimately accountable now. They deserve support in understanding how these systems work, because they are the ones who will be held accountable.

I think that particularly going to something that is underdiscussed, there are the cultural sovereignty implications for using these large LLMs. If MPs are using U.S.-based LLMs to help draft policy, they might not be aware of the way in which the Americanized datasets shape the language they're using and, through compression, leave out particular identities, particular aspects of culture, because they are compressing vast amounts of data. I think being aware of that and the cultural implications for Canada is important.

[*Translation*]

Christine Normandin: Thank you.

I have a quick question.

Mr. Lemor, you spoke about the risk of legislative inflation associated with using artificial intelligence to create legislation or regulations.

Do you see a risk that it might become more difficult for the public—or for anyone—to actually analyze the legislation, especially given the use of artificial intelligence for the analysis? As we've seen, bots can make mistakes.

I think that you will be publishing a book soon about how this worked during the election. I'm thinking here about bots and the risk of errors.

Does legislative inflation carry this type of risk?

● (1250)

Antoine Lemor: I pointed to the example of the United Arab Emirates, which recently announced that it's using artificial intelligence to try to speed up its legislative process.

Of course, the risk is that we'll see regulatory and legislative inflation if we introduce artificial intelligence into the legislative process itself. There will be consequences. If we focus solely on efficiency, it will affect the meaning attached to the legislation itself. The reason that only a certain number of pieces of legislation are passed each year stems from the fact that a certain value gets attached to a bill. If a huge volume of bills were introduced each year, the value of the legislation would automatically decline. It would also lose its significance to the public. People would stop talking about the bills and the discussions surrounding them.

That's the most significant impact. It isn't necessarily a measurable impact. However, it's a perceptible impact in terms of the meaning that we collectively attach to the legislation.

The Chair: Thank you.

[*English*]

Mr. Calkins, you have five minutes, please.

Blaine Calkins: Thank you so much, Chair.

I want to talk about semantic drift and compounding errors. I come from a background, even though it was a long time ago, in which I was an IT professional. I was a tenured faculty member at a college in central Alberta. I taught information systems technology. I taught database administration. I will confess that in the world of AI, notwithstanding all the credentials I used to have that are quite dated now, I feel like a Luddite when it comes to this, so I completely understand how most parliamentarians would feel a little bit intimidated by the use of AI and by having debates and conversations about this.

I believe that the entire conversation about AI and digital literacy writ large is probably the most important theme we've heard, and I think it's probably going to be one of the most important recommendations that comes out of whatever this study eventually lends itself to.

Twenty-five years ago I used to laugh at headlines that said, “Computer error blamed”. Computers don't make errors. A computer will do the same calculation over and over again a hundred thousand billion times and get it right a hundred thousand billion times, if the right inputs are put in place. Computers don't make mistakes. Humans make mistakes. Now we're talking about AI making mistakes. I think it's important that we distinguish the difference between a computer and artificial intelligence.

Somewhere along this line, when it says that a computer...and you triggered this when you answered your question about IBM saying a computer should never make a management decision.... A computer simply does the calculations that a programmer or a coder has asked the computer to do, and it'll calculate the same result, every single time, the same way, but that's not true with AI.

Is there somebody here who would like to explain the difference between a computer doing a calculation and an AI agent answering a question?

[*Translation*]

Antoine Lemor: In research, we rely heavily on artificial intelligence for the results of the research. However, we can't delegate the task of conducting our research to artificial intelligence.

We'll use it, for example, to analyze large amounts of data—such as data related to parliamentary debates—or to categorize texts when a parliamentarian uses scientific evidence to answer research questions on the use of science in Parliament.

In this context, we use artificial intelligence, but we maintain control over it. We'll verify whether the artificial intelligence—such as machine learning models or large language models—categorizes texts in the same way that we do. If so, we'll implement a process to ensure that the machine's decisions reflect human decisions.

To answer your question indirectly, we can establish standards and rules to ensure that artificial intelligence, when operating within a clearly defined framework, remains under human control. If I take the example of scientific evidence, artificial intelligence—meaning the machine learning model—will categorize sentences containing scientific evidence only according to how we define that evidence. This definition remains the product of human cognition, and it must stay that way.

These processes can be implemented.

● (1255)

[*English*]

Blaine Calkins: This is my concern with the fact that AI now is basically doing some of the thinking for us. AI is now learning from the results that were created by other AI. If that AI agent, that search or that research created a mistake or an error....

There's a game that we used to play. It started off with a sentence and you changed one word, and passed it to the next person. You try not to change the meaning. You pass it around to 20 people. That sentence gets back to the original person, and it doesn't sound like the same sentence at all, even though everybody tried to actually keep the same meaning and same intent in that sentence. That is semantic drift. That is the fact that there's compounding errors in this. How prevalent is this, and what are the risks to Parliament?

Beatrice Wayne: I would say it's prevalent. I would also say we often rely on the RSI, recursive self-improvement. These models will keep getting better and better, and we can rely on this, but the problem is exactly as you describe. Particularly depending on the dataset you have, the way in which the data is compressed, the problems compound. If you have a dataset that does not contain multiplicities and doesn't surface divergent points of view, recursive self-improvement is not actually improving; it's compounding an existing problem.

We need to be aware of that, and particularly, back to questions of cultural sovereignty, we're relying on these giant datasets where we don't understand exactly, because they're black box systems, the data it's trained on, but we know largely that they're trained on Americanized datasets, Reddit, and other forms of non-peer reviewed information.

This is going to shape our knowledge, our public knowledge but also the knowledge of legislators in the way that they approach policy. This is a concern. I think it's really important for legislators to understand that we can't rely on recursive self-improvement as the cure-all and save-all for our problems that are currently existing in these models.

The Chair: Thank you so much.

For our last question, we'll go to Madam Mingarelli for five minutes, please.

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

Thank you to all the witnesses for being here today.

I'm going to direct my first question to you, Madam Ivory.

How is the Library of Parliament preparing for the growing use of AI in parliamentary research and information services? In particular, are any of the new researchers undergoing prompting training and other kinds of AI development training to support their work?

Christine Ivory: We are. We've developed some training for all our employees, obviously for the researchers and librarians, but AI has other uses in our institutions, so we've developed that. Also, to the point of this discussion on digital literacy, we are working with the House of Commons to develop training for parliamentarians and staffers that's similar to what we've developed internally to help with the prompts and upskilling on the use of AI.

We have plans for all our employees to be able to have upskilling in terms of AI skills. That's different depending on whether it's an expert in data visualization, a graphic designer or a researcher. It's part of everyone's professional development if they want or need that training. It's now really a priority for us to do this as opposed to bringing in specialized positions on AI, because it is so pervasive to the work we have to do.

Giovanna Mingarelli: How is the AI that the Library of Parliament is using presently being trained? Who is it being trained by? Where is the data? How is the data being verified to ensure that it's not being exposed to things like foreign interference and external inputs that could be providing content that is hallucinatory or false?

Christine Ivory: I'll start, and maybe I'll ask my colleagues to jump in.

Because we are on the House administration IT system, we are using Copilot, and some of our employees are part of the pilot to use the Copilot Pro licences. At the moment, because this is the most secure and vetted AI solution available at Parliament, it is what we are utilizing. There are some clear guardrails, the data sovereignty in Canada that's been assured by Microsoft, those are all things that we've worked closely with the House on to make sure that we're complying with. We have very limited use cases to pilot other AI, but not across the institution.

I don't know whether Marcus or Johanna want to add to that.

● (1300)

Giovanna Mingarelli: Great. Thank you—

Christine Ivory: Maybe, to the point that has been mentioned here about hosting our own curated LLM, I'll just quickly say that we are...for some of the routine tasks related to our parliamentary work. That's something we're exploring. It's very nascent at this point, but it's one of our objectives.

Giovanna Mingarelli: How do you account for the two official languages in AI? Also, do you account for any indigenous languages, in terms of how the AI is being trained and in ensuring adequate responses in both official languages and then, of course, any relevant indigenous languages?

Christine Ivory: This is a big part of the issue of the limitations of AI for the Library of Parliament's support to Parliament currently. Whether it's official languages and the content that is available, it touches back to the cultural sovereignty issue and representation in AI models. We find it very limiting for French, for indigenous or other marginalized communities in Canada. It's the reason that AI can be a tool that we use in the process, but it is not the be-all and end-all. We are still relying on how our analysts are trained in GBA+ analysis and we're looking at how we've done research for years without the use of AI.

Giovanna Mingarelli: Thank you.

I think I'm done. I can't ask my question quickly enough.

Thank you so much.

The Chair: Thank you so much.

Before I end the meeting, I do want to say thank you, especially to Ms. Wayne, Monsieur Lemor and Mr. Hall, who got dragooned into a second hour of appearing. However, based on the questions, it was important information, and I'm glad you were able to stay. Thank you so much.

Thank you to the Library of Parliament, of course.

This meeting stands 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