



Enhancing trust in AI through increased transparency

Submission of the Canadian Association of Research Libraries
September 22, 2026

This consultation took place through an online submission form with respondents choosing which sections they wished to comment on. The topics below are those the CARL submitted a response.

AI-Generated Content

Would it help you trust what you see online if you could tell whether something was created by AI or by a person? When and for what kinds of content (e.g., image, video, audio) is it most important to know if something was created or modified by AI?

All AI generated content and material modified by AI must be clearly identified as such and treated with equal priority. This is essential to ensuring transparency, combating misinformation and disinformation, and providing clarity in terms of copyright. Academic libraries actively engage in combating disinformation and misinformation and educate students, faculty and staff on how to identify fabricated and misleading data and resources. Furthermore, academic libraries navigate complex licensing agreements to provide legal access to research and educational materials and support institutional copyright compliance and measures. In order to provide access to and disseminate vast collections of scholarly works, academic libraries must be able to recognize and detect AI generated and modified content. This is essential in verifying the integrity of sources and maintaining credible collections. Likewise, libraries must be able to train patrons on how to distinguish between AI generated content and human creations. Clear transparent identification of AI generated and modified content is essential in building trust and ensuring integrity in works.

What would best help Canadians determine when content is AI-generated (e.g., invisible or visible watermarks, disclaimers, provenance metadata)?

Identification of AI content should be both human and machine-readable. No single method, whether it be watermarking, disclaimers, or provenance metadata, is sufficient on its own, since each has different

strengths and vulnerabilities. The most effective approach is therefore a layered one, combining multiple mechanisms rather than relying on a single safeguard. This layering matters for two reasons. Firstly, it raises the bar for bad-faith actors by addressing content that has been stripped of a visible label to carry embedded metadata or a watermark, making it harder to fully obscure AI involvement. Second, it gives good-faith actors - platforms, journalists, researchers, and the public - more tools to verify authenticity, rather than depending on a single point of trust that can fail or be circumvented.

Provenance metadata would serve as a gold standard as it not only identifies a work as AI-generated, but includes information about the work's creation, including what tools were used, and to what degree. This approach acknowledges the complexity of AI-enabled creation, in which works are not either completely human made or AI made, but are often a mixture of the two.

Invisible watermarks (machine-readable) are another valuable technologically-enabled tagging method for AI as they can persist even when content is copied or re-uploaded. This can be helpful in preventing disinformation, as it is information that cannot easily be stripped out or removed.

Visible watermarks are also highly important, especially to the human reader as they serve as the most immediately legible evidence of AI usage and do not require any special skills or tools to interpret.

Critically, none of these tools work without public infrastructure verification. Canadians need accessible, trustworthy tools (e.g. browser-level indicators and platform-level checks) to confirm authenticity, not solely a claim attached to the content. Libraries and public institutions have a role to play in digital/information literacy education, helping people understand what these signals mean and how to check them.

Who in the AI value chain (developers, deployers, or others) should be responsible for providing transparency around AI-generated content? Why?

Responsibility for transparency should be distributed across the value chain, with obligations calibrated to each actor's degree of control and knowledge, rather than assigned to a single point.

Developers are best placed to embed technical provenance mechanisms at the point of creation, since they alone have visibility into training data, model architecture, and generation processes. This includes implementing robust content provenance standards and disclosing the general nature of training corpora, including whether copyrighted or licensed scholarly material was used. Developers are also best positioned to ensure that provenance signals are durable and cannot easily be stripped out of content as it is disseminated.

Deployers (those who integrate models into products or services, including publishers, platforms, and institutional systems) bear responsibility for surfacing that provenance information to end users in an accessible way, and for disclosing when AI has been used to generate or substantially modify content presented to the public. In a research or library context this matters directly because AI tools are embedded by deployers into discovery layers, research databases, or publishing platforms. Scholars and students should know if content they are being presented with, be it AI-generated summaries, abstracts, or other content is produced by artificial intelligence so they can assess its reliability and provenance before relying on or citing it.

The government should establish a regulatory baseline that makes distributed responsibility workable. Voluntary, industry-led standards risk inconsistency, leaving users unable to rely on transparency signals. Regulation should mandate minimum disclosure and provenance standards, ensure interoperability, and provide oversight to verify compliance rather than accept self-monitoring by companies. These standards should be calibrated to avoid unduly restructuring legitimate text and data mining or the development of a sovereign academic AI infrastructure.

Do existing market practices, technical tools, and legal frameworks make it easy enough to know when content is AI-generated? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?

AI generated content is not easily or reliably identified as evidenced by growing issues with deepfakes, fake news and other AI powered scams. While some technical capabilities exist to label AI generated and modified content they lack consistency in terms of process and implementation. Canada's Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems encourages labelling of AI generated content but is optional and unevenly adopted across the sector both in terms of if and how AI generated content is identified. More must be done to achieve transparency with respect to AI use in creating and modifying content and a multi-pronged approach to transparency most likely to succeed. As noted in the previous question, government regulation is needed to address this.

While tech companies are currently taking the lead in terms of how content can be labelled and in developing counter measures that assist in identifying AI generated and modified content, experts across stakeholder groups should be engaged to ensure transparency measures are attainable and understood by all parties. Academic libraries can play an essential role in literacy initiatives needed to understand these signals.

In addition to national stakeholder and expert consultations, Canada can benefit from monitoring AI disclosure and transparency requirements in other jurisdictions. Observing the successes and limitations of the European Union's phased transparency obligations in the EU AI Act, and the mandated AI labeling requirements in South Korea's AI Basic Act can help inform and establish robust processes and practices in Canada.

AI Interaction

When and why is it most important to know that you are interacting with an AI system? Are there contexts where you don't need to know that you're interacting with AI?

Disclosure is most critical where the AI system is generative and produces content that could be presented, cited, or relied upon as if it were human-authored. This is especially important in an academic context. Disclosure is also important when an AI system is being used to make or materially inform consequential decisions, for example, admissions decisions, funding decisions, content

moderation, etc. AI must also be disclosed when it involves direct engagement with users, as in a conversational or chat bot context.

On the other hand, background or infrastructural uses of AI, such as recommender systems, spam filters, search ranking, or backend optimization, do not require a front-facing label or declaration. For example, it is not as important that Netflix users know that AI is helping to produce their recommended selections, but this methodology should be documented and accessible on request. While disclosure is generally encouraged, over-labelling low-stakes, well-understood uses risks desensitizing users to disclosure altogether, making the genuinely important labels less effective.

What factors make a disclosure of AI use meaningful and effective for the user?

Simple declarations such as “This content was generated by AI” may be sufficient in some content generation cases, for example, an AI generated image or video but ineffective in other use cases. In cases where the AI interaction is more complex or interactive, meaningful and effective disclosure must provide users with details. Users must be able to determine: what content was used to train the AI; if inputs will be retained and how they will be used; if there are any known limitations or risks associated with generated outputs, including if sensitive information is involved (e.g. medical or legal content). Furthermore, users must be able to identify who developed and deployed the AI, what the intended purpose of the AI is and at what point, if any human oversight is engaged. Users must also be able to seek additional information on an AI system and pursue accountability if issues arise.

Who in the AI value chain (developers, deployers, or others) should be responsible for providing transparency around AI interactions? Why?

Transparency requires shared responsibility and Canadians must be able to not only determine when they are interacting with AI but assess the implications of doing so. Disclosure requirements could be scaled based on type of AI interaction and AI tool. For example, while a simple notice indicating a customer service chatbot is not human and cannot effectively make human decisions may be sufficient in some cases, more robust notices and information would be required for relational chatbots or AI interactions that simulate human contact or involve decision making implications.

Developers must document and disclose the inclusion of AI in products and deployers must also be responsible for making AI disclosures available to users in a way that can be easily understood. Uniform or consistent language and notification standards would benefit transparency efforts.

Are existing market practices and legal frameworks sufficient to support transparency around AI interactions? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?

Canada should consider regulatory measures or standards with respect to the disclosure of AI interactions. In addition to engaging with Canadian experts and AI developers and deployers in Canada on the best way to achieve this, the government can monitor legislation in other jurisdictions, such as California and New York that have sought to regulate AI relational interactions. In addition to other provisions, both California's companion chatbot law, SB 243, and New York's AI Companion Models statute include user notification requirements and additional transparency protocol obligations. Canada has an opportunity to learn from this legislation, and it can inform the implementation of Canadian transparency requirements for AI interactions.

AI interaction disclosure regulation and transparency requirements should consider the type of interaction and implications of the interaction. Canadians should be provided with sufficient information to understand they are interacting with AI and the limitations of the information and interaction. Additional measures should be considered for AI interactions that simulate human experiences and engagement.

Literacy is essential in helping users understand the implications of AI interactions and the scale of significance varies greatly depending on the application and AI interaction. Academic libraries engage with campus constituents on these topics, but additional national literacy programming is essential.

Information about AI Systems

What kinds of information about AI would you like to know in order to help you make informed choices about when and how to use AI? When and for which types of AI products is this information most important to you?

Detailed information about how and why an AI system was developed and the intended application must be made available for all AI systems and products. In order to make informed decisions about the acquisition and adoption of AI systems, Canadians must be able to determine what type of AI the technology or system is reliant on; who created it; and what content was used to create it, including what was used to train and develop the system or model and if updates are needed to ensure the system remain relevant and up-to-date. Canadians must also be able to understand the intended use cases for the application; any known limitations to the system or product; and, be able to assess if the system complies with relevant Canadian laws, such as privacy and copyright law. This information is essential in allowing Canadians to acquire and adopt AI systems that meet their operations needs and mitigate AI related incidents and issues. For example, knowledge of any system limitations, such as inherent biases that may result in misinformation or known hallucinations, fabrications, or vulnerabilities that may pose security issues are essential and would only be known if made transparent. In addition, detailed information is required to ensure deployers can fulfill any disclosure and

transparency requirements associated with the platform and support literacy efforts of academic libraries and other organizations.

Developers must clearly document and publicly disclose this information. And while this should be required for all generative AI systems, it would be a general best practice or requirement for all AI systems and products.

What kind of information about AI do businesses need to adopt AI safely and responsibly? Do businesses currently have access to the information that they need? How can the Government best balance the need to increase transparency (including input transparency) against the need to protect confidential business information?

Businesses and organizations must have sufficient information to determine if an AI system will meet operational needs and to assess reliability, veracity and any risks both legal, social and environmental associated with the adoption and implementation of an AI system. In addition to the technological capabilities, companies must be able to determine if systems align with ethical practices and commitments, including social responsibility, environmental sustainability, fair labour and transparent governance. For example, academic libraries need sufficient information about an AI system to determine if it meets functional and pedagogical requisites and institutional procurement standards but also to assess if the system or product is trustworthy and aligns with institutional academic integrity, copyright and information management policies as well as applicable provincial and federal legislation.

Currently, the lack of transparency and disclosure requirements is exacerbated by rapid technological expansion, with development exceeding physical computing power and in some cases lack of human control and oversight of AI systems. Information about AI systems is also often locked down or behind a paywall, creating barriers to businesses and organizations to obtain the information they need.

Canadian businesses and all adopters of AI products would benefit from established transparency guidelines that establish what information needs to be made available and that it be publicly accessible. The government could play an invaluable role in stipulating what information must be made publicly available upon an AI system becoming available for acquisition and adoption in Canada.

Who in the AI value chain (developers, deployers, or others) is best placed to provide what information about AI systems, how they are created, and their capabilities and limitations?

All parties should be accountable in advancing transparency and providing appropriate AI disclosures but the responsibility should be distributed according to what information each stakeholder has.

Developers hold direct knowledge of how a system was built, including training data composition, architecture, and known limitations. They are best placed to provide this foundational information and have a responsibility to disclose it in a human-readable form, not just retain it internally.

Deployers are best places to explain how a system performs in the specific context it is actually used, since practical capabilities and limitations in a given application can differ from a model's general documentation behaviour. Deployers bridge the gap for end users.

Both need reliable access to current AI research to meet these responsibilities. The government should support access rather than solely leaving it to the market. By backing open access to research on AI capabilities, limitations, and environmental impacts, and by channeling applicable research to developers and deployers, a robust and sustainable environment of disclosure can be fostered.

Are existing disclosure practices and legal frameworks sufficient to provide the information that users and other stakeholders need? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?

While the Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems references transparency and calls for sufficient information to be published to allow consumers to make informed decisions and for experts to evaluate risk, more must be done to ensure Canadians have the knowledge they need to adopt and implement AI systems. Clear disclosure requirements would reinforce what is currently voluntary and ensure model cards that describe AI systems include all relevant and pertinent information, including data training disclosures, data provenance and how risk has been assessed and managed. Additionally, the creation of a required risk based framework, like the structure in the EU AI Act, and safety guardrails for the development of AI systems, particularly those that simulate or create human-like experiences, would aid transparency by establishing a structured method to ensure AI systems are understandable, auditable and traceable supporting safe adoption and encouraging public trust. Recent revelations from Anthropic and OpenAI CEOs regarding the significant risks associated with the rapid development of AI and their own admission that public fears are justified illustrates the need for strong transparency measures.

In addition to the EU AI Act, the South Korean AI Basic Act serves as another model for pre-market compliance and risk evaluation for advanced artificial intelligence systems. The legislation requires companies to address transparency, user protection and risk management prior to market entry and includes human oversight frameworks and requirements for certain sectors and checklists for others.

While any regulation, standards and frameworks must be nimble and able to address rapid development, changing landscape, and product capacity, they are essential to safe acquisition and adoption of AI systems and compliance with related Canadian laws. Our comments focus on regulated solutions, however, should the government consider and pursue legislative solutions, it is critical that current outstanding legal disputes are taken into account as the courts may provide valuable guidance that informs the drafting of new legislation or amendment of existing laws. Ideally any measures would be accompanied by literacy initiatives and guidance for all AI system stakeholders. There is a role for academic libraries in literacy campaigns as university campuses, including libraries are not only procurers and adopters of AI systems but also support faculty, students and staff in the use of these products and tools.

AI Agents

Are you concerned about the use of AI agents? In what contexts is it most important to know when an AI agent is in use? Why?

CARL has concerns about the use of AI agents. Generally, our concerns center on two related risks: uninformed use of agents by end users who do not understand what the agent is doing on their behalf, and the overuse of agentic systems in contexts where their autonomy outpaces institutional readiness to govern, support, or evaluate them.

The National AI Strategy's commitment to providing all post-secondary students with "access to trusted AI agents" raises immediate concerns. For example, how will the agent be selected and will institutions be free to choose their own agent. Without clarity on procurement and governance, institutions risk either being locked into a single vendor's ecosystem or face an uncoordinated patchwork of tools with inconsistent standards for privacy, data governance, and transparency.

Any governmental oversight per the use of AI agents on academic campuses must also take into account the rigorous standards associated with academic integrity and scholarship and ensure that AI agents rolled out for post-secondary students are highly transparent when it comes to content generation.

Who should be responsible for disclosing the use of AI agents? What kinds of information should they be recording regarding the activities of AI agents, and what kinds of disclosures should they be making, and to whom?

Transparency requires shared accountability and anyone deploying and using AI agents must disclose their use publicly. Detailed information, including prompts, answers and the number of iterations it took to obtain the desired response should be documented, retained and readily accessible to the public and developers should be required to provide the technical capacity and means to facilitate this. Disclosures could be as straightforward as acknowledgement statements or citations and best practices and providing guidelines could contribute to consistent declarations.

Additional measures and recording keep may be needed for more advanced and complex and multi-agent interactions. And specific controls and protocols should be in place for AI agents that collect and use personal information and data. The ability to identify malicious, malfunctioning and rogue agents and transparent disclosure of these events is necessary in order for there to be accountability and to assess liability and prevent ongoing or future harm. Transparency measures and disclosure requirements will be essential in mitigating risk and pursuing recourse and correction when situations call for it. The adoption of disclosure requirements should complement and uphold existing relevant Canadian laws, such as privacy, consumer rights and civil liberties.

Academic libraries are seeing increasing embedded AI agents within the collections that they purchase and in many cases, without the ability to control whether the AI agents are activated or not. Vendors need to make it clear to students and faculty that AI agents are being used while learning materials are being used, what the agents are collecting, and what the data they are collecting will be used for.

What kinds of information do businesses need in order to confidently and safely use AI agents? Do businesses currently have access to the information that they need about AI agents?

Academic libraries need to assess several factors in order to deploy an AI agent and agents with agentic capabilities require additional considerations and review. Much of the information needed to make these determinations is similar to that which would be considered when adopting an AI system such as, how the agent was trained and on what data, what information the agent collects and retains, how the collected data is reused and where it is stored, does it comply with provincial and federal laws (such as privacy), does it comply with institutional policies (procurement; security classification, copyright, academic integrity, etc), what are the contractual obligations if it is licensed or acquired from a third-party, what verifications, and what monitoring or oversight exists.

Are existing technical solutions, market practices, and legal frameworks sufficient to support transparency around agentic AI systems? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?

Our points made under transparency around AI interactions would also apply to agentic AI systems, and we support the regulation of AI developers and deployers and additional transparency measures.

We think the government should partner with trusted organizations like academic libraries in addition to AI research institutes to develop and deliver literacy initiatives.

Broader considerations

Given the areas discussed here, when is AI transparency most important, and are there any areas where action would be premature? Are there other areas that should be considered for action to enhance AI transparency?

Transparency is critical in every aspect of AI deployment and use. Everyone has a part to play in AI transparency but responsibility scales with control. Developers and deployers hold the core responsibility for ensuring safe, transparent use.

Transparency is most important where AI outputs could be mistaken for human judgement or human-authored content, where AI materially informs consequential decisions about a person, and where a user might otherwise reasonably assume they are interacting with a human rather than an AI bot. Canada already has a relevant reference point in the federal Directive on Automated Decision-Making, which establishes transparency and accountability requirements for automated decision making in the government.

Transparency needs will continue to evolve as AI itself evolves, and any regulation or policy response must be able to evolve with it to keep pace with consumer needs. For this reason, we caution against enacting rigid legislation at this stage. Requirements must be rigorous, but the mechanism for enacting and updating them needs to remain nimble, allowing standards to be revised as the technology, its applications, and its risks change, rather than locking in requirements that quickly become outdated or mismatched to new use cases.

How should any measures account for the circumstances of small and medium-sized enterprises?

AI transparency requirements should be able to be applied regardless of the size or circumstances of the enterprise but with greater expectations for larger operations and more consequential AI tools and platforms. AI transparency requirements should enable compliance with relevant Canadian legislation, including privacy, copyright and the Safe Social Media bill. And expectations and responsibility should be scaled in response to impact. For example, platforms that make content available and social media platforms should be held accountable for exposing users to content that is misleading, misinformed or malicious and subsequently AI systems that create or replicate human-like experiences must be responsible for transparency and protections for users.

How can Canada design measures to best account for the ongoing evolution in AI technologies and their use?

Canada has the benefit of monitoring AI legislation in other jurisdictions (e.g. European Union's AI Act) and can adopt elements of these Acts that have been effective or particularly successful in facilitating transparency and more generally addressing other AI issues and concerns. Continued engagement with a wide range of stakeholders through the development of legislation and policy will reduce the risk of AI enterprises driving legislation and will proactively create an inclusive environment in AI-driven strategies. Additionally, the government's approach to AI regulation, technology, and use should be supported by a robust independent research and development arm.

Any regulatory framework that is developed should include review mechanisms to regularly update technical standards that reflect developments in technologies and evidentiary outcomes. This would be done in tandem with organizations' regular review of processes, data, intended users, as operating environments change. This would also serve to stay abreast of international approaches/changes and would maintain oversight to ensure interoperability.

Ensuring human oversight and accountability would keep responsibility with identifiable people or organizations even when AI systems become more autonomous. The governance framework should remain clear about who is accountable for decisions and outcomes.

Keeping AI technologically neutral would ensure interoperability and portability and avoid pitfalls related to a particular vendor, model, technical architecture, or implementation and would support a more resilient governance model as organizations adopt new AI systems.

Having regulators collect information on actual system performance, incidents, unintended effects, and user experiences would support evidence-based decision making and enhance learning requirements which would lead to evidence-based decision making rather than assumptions about future technology.