



CRAiEDL

The Canadian Robotics and Artificial Intelligence Ethical Design Lab

Submission to the ISED National Sprint to Shape Canada's Renewed AI Strategy

Dr. Jason Millar (Director, CRAiEDL; Canada Research Chair in the Ethical Engineering of
Robotics and AI)

Dr. Kelly Bronson (CRAiEDL; Canada Research Chair in Science and Society)

Talia Cameron (PhD Candidate, CRAiEDL)

Caitlin Heppner (PhD Candidate, CRAiEDL)

Mikeli Habash (MA, CRAiEDL)

Amal Hussein (MA, CRAiEDL)

October 31, 2025

Table of Contents

1. Director's Remarks	3
2. The Canadian Robotics and AI Ethical Design Lab (CRAiEDL)	4
3. Targeted AI Sprint Question Responses	4
3.1. How can Canada strengthen coordination across academia, industry, government and defence to accelerate impactful AI research?	4
3.2. What conditions are needed to ensure Canadian AI research remains globally competitive and ethically grounded?	5
3.3. What are the key barriers to AI adoption, and how can government and industry work together to accelerate responsible uptake?	6
3.4. How can Canada build public trust in AI technologies while addressing the risks they present? What are the most important things to do to build confidence?	6
3.5. What frameworks, standards, regulations and norms are needed to ensure AI products in Canada are trustworthy and responsibly deployed?	7
3.6. How can we enhance AI literacy in Canada, including awareness of AI's limitations and biases?	8
3.7. What can Canada do to ensure equitable access to AI literacy across regions, demographics and socioeconomic groups?	9

1. Director's Remarks

This submission is in response to Innovation, Science and Economic Development Canada's (ISED's) National 30-Day Sprint to Shape Canada's Renewed AI Strategy.

An important goal of this rapid public consultation is determining how to get Canadians to trust AI, so they will use it more. I worry that this sprint, if the Government of Canada plans to treat it as a one-off, will undermine that goal.

There is no doubt that AI is one of the most promising technologies of our time; AI stands to benefit Canadians immensely, if we build and deploy it responsibly. But recent headlines demonstrate the significant risks to which AI subjects Canadians, especially children and youth, when it is deployed frantically. Children are forming problematic emotional relationships with LLMs; Chatbots are exacerbating suicidal tendencies; Educators are deluged with pedagogical challenges due to AI-enabled cheating and deskilling; Democratic norms are being eroded through the rampant spread of misinformation and deepfakes; Studies suggest, despite anecdotal claims that AI is improving productivity, that AI is actually decreasing productivity in areas in which it is meant to excel, such as programming. All the while, AI-producing corporations continue to make unproven claims about the economic benefits of using their AI systems, and politicians, business leaders, and other decision-makers emphasize the need to rapidly adopt AI so Canadians don't fall behind. The technological FOMO is palpable.

These mixed public messages have a predictable effect: according to a recent KPMG report, Canadians don't really trust AI. But perhaps Canadians are justified in their reluctance to adopt AI. Their low trust of the technology seems appropriately calibrated—the current AI ecosystem appears untrustworthy.

We urge the Government of Canada to adopt the right goal: to increase the trustworthiness of AI systems that are deployed in Canada. Nobody should encourage unwarranted trust in AI. Increasing the trustworthiness of AI is at the core of our responses.

Because the sprint was rapid our approach at CRAiEDL was to select a subset of questions from ISED's longer list that we felt landed in our general pool of expertise. We convened a small half-day workshop to draft responses to those questions, and collaboratively edited our answers, which form the remainder of this submission.

I would urge the Government of Canada to treat this sprint as a first step in a series of meaningful and much-needed public consultations on much-needed (and overdue) sovereign AI policies. The risks we outline in this submission are backed by a growing body of publicly available evidence of harms to Canadians. Canadians, Canadian children and youth most urgently, deserve strong federal policy responses that protect them from the most egregious deployments of risky, often knowingly harmful, AI-based products. Strong, enforceable red lines prohibiting specific harms promise to help make AI more trustworthy. Only with those protections in place should Canadians begin to trust AI and consider its broader adoption. Only with those protections in place can the Government of Canada responsibly urge them to do so.

Dr. Jason Millar

BScE, BA, MA, PhD, P.Eng.

Canada Research Chair in the Ethical Engineering of Robotics and AI

Associate Professor, School of Engineering Design and Teaching Innovation

Cross-Appointed to the Department of Philosophy

Faculty of Engineering

University of Ottawa

2. The Canadian Robotics and AI Ethical Design Lab (CRAiEDL)

CRAiEDL is an interdisciplinary lab conducting research on the complex ethical and social issues raised by existing and emerging robotics and AI technologies and applications, to help inform the responsible development and deployment of robotics and AI technologies, applications and policy. CRAiEDL develops concepts, frameworks, and tools to support the responsible design, development, deployment and governance of robotics and AI. Our research outputs can be used by students, start-ups, private sector companies and policymakers to build applied ethics capacity and support the development and application of best practices for robotics and AI in Canada and around the world. CRAiEDL also works with stakeholders across Canada and internationally to help develop both general and industry-specific policy thinking to support the responsible governance of robotics and AI.

3. Targeted AI Sprint Question Responses

3.1. How can Canada strengthen coordination across academia, industry, government and defence to accelerate impactful AI research?

(i.e. mechanisms for cross-sector collaboration; integration of public and private research efforts; industry-sponsored research while preserving academic independence)

To achieve your vision of “Impactful AI” that protects human rights, serves the public good, inspires trust and equitably balances risks and benefits, we believe GOC needs to establish a new institute to supercharge the social, legal and ethical (i.e. sociotechnical) dimensions of Canada’s AI ecosystem—the Sociotechnical AI Research Institute (SAIRI). Canada has benefitted from its three fundamental AI research institutes—Amii, Mila, and the Vector Institute—but those institutes are not best-positioned to lead on sociotechnical issues. Thus, an institute focused on sociotechnical issues—SAIRI—is required to balance the scales. Determining how AI can protect human rights, serve the public good, and inspire trust are activities that do not fall under the domain of fundamental AI research—they are the proven domains of the social sciences and humanities.

SAIRI’s research would complement and strengthen Canada’s existing fundamental AI research and governance ecosystem by focusing on researching and developing reasonable policy recommendations (including regulations where necessary), ethical frameworks, standards, and methods and processes to ensure the ethical procurement, development and governance of AI. Further, this new institute could develop meaningful sociotechnical metrics to measure the

benefits and risks of AI, enabling the monitoring of AI research and deployment to ensure Canada's AI ecosystem delivers on ISED's vision of "Impactful AI".

Canada has a deep well of AI talent that has consistently demonstrated its ability to innovate in both fundamental AI research and in the sociotechnical aspects of AI and other emerging technology. Indeed, Canada once led the global charge for meaningful AI governance with its Directive on Automated Decision-Making. By establishing SAIRI, focused on the sociotechnical dimensions of AI, we can better leverage this talent and innovation to create more trustworthy, impactful AI aligned with ISED's vision of "Impactful AI".

3.2. What conditions are needed to ensure Canadian AI research remains globally competitive and ethically grounded?

(i.e. infrastructure, talent and governance enablers; ethical standards and risk mitigation; alignment of applied research with business and societal needs)

Canada has maintained global competitiveness in fundamental AI research through the establishment of its three AI research institutes. Canada has demonstrated global competitiveness in AI governance with the Directive on Automated Decision-Making, which grounds the public use of AI in ethical principles and administrative law. To ensure Canada's global competitiveness beyond the public sector, we suggest establishing SAIRI, focused on maintaining our leadership in the sociotechnical aspects of AI, to support the development and establishment of strong, sensible, regulations and standards for AI development and deployment grounded in trustworthy AI that serves the public good, and strengthens and protects human rights.

Ethically grounding AI requires the establishment of strong, sensible, regulations and enforceable standards restricting the development and deployment of AI systems that significantly harm Canadians, that is, systems that: induce emotional dependence; induce addiction; employ deceptive practices; target children; undermine democratic norms; or illegitimately concentrate power. A research institute (SAIRI) mandated to provide global knowledge leadership on AI governance in both the private and public spheres, would strongly support the development of such guardrails.

Canada's low AI adoption rates suggests that Canadians are sophisticated when it comes to thinking about AI—they understand the return on investment and/or safety of most AI products has not yet been proven. Canadians are not under-trusting AI technologies, they are trusting them appropriately, and do not currently have strong reasons to adopt AI. The solution lies not in making Canadians trust AI more, but in making AI more trustworthy. If we commit to researching and developing AI products that are ethically grounded and truly trustworthy, then we stand to outperform those who do not.

3.3. What are the key barriers to AI adoption, and how can government and industry work together to accelerate responsible uptake?

(i.e. sectoral vs. cross-sectoral challenges, such as liability and small to medium-sized enterprise constraints; potential government policies, incentives and ecosystem supports)

Recent indicators suggest that AI adoption is lagging in Canada. This is often interpreted as a problem of trust—if Canadians had more trust in AI, they would adopt it. The lagging AI adoption rates in Canada can also be explained if we consider that Canadians have an accurate read on the trustworthiness of current AI systems.

News headlines of late suggest that AI is making significant gains in some areas, but that it also poses significant unmitigated risks to Canadians, especially children and youth. Children are forming problematic emotional relationships with LLMs; Chatbots are exacerbating suicidal tendencies; Educators are deluged with pedagogical challenges due to AI-enabled cheating and deskilling; Democratic norms are being eroded through the rampant spread of misinformation and deepfakes; Studies suggest, despite anecdotal reports that AI is improving productivity, that AI is actually decreasing productivity in areas in which it is meant to excel, such as programming. All the while, AI-producing corporations continue to make unproven claims about the economic benefits of using their AI systems, and politicians emphasize the need to rapidly adopt the technology. These mixed messages have a predictable effect: Canadians are justified in their reluctance to adopt AI; their level of trust in AI seems appropriately calibrated—the current AI ecosystem is untrustworthy.

If this paints an accurate picture, the key barrier to AI adoption is a lack of trustworthy behaviour on the part of those building and promoting the use of AI. A potential solution to this problem is strong, sensible, regulations and enforceable standards restricting the development and deployment of AI systems that significantly harm Canadians, that is, systems that: induce emotional dependence; induce addiction; employ deceptive practices; target children; undermine democratic norms; or illegitimately concentrate power.

3.4. How can Canada build public trust in AI technologies while addressing the risks they present? What are the most important things to do to build confidence?

(i.e. risks posed by AI tools and services; drivers of public and business mistrust; educational and literacy strategies to foster informed confidence)

Canada can build public trust in AI by supporting AI-producing companies in their efforts to make AI more trustworthy. Canadians are warranted in their mistrust of AI—recent headlines are awash with examples of AI posing serious risks and causing significant harm. In the current

economic climate, in which companies are desperately trying to outpace each other in a race to capture billions of dollars in market share, it is simply unreasonable to expect AI producers to voluntarily improve their practices or to self-regulate in the public interest. There has never been more need for sober, informed, rational constraints to be placed on corporate actors.

Canada can build more trust in AI by implementing enforceable frameworks, standards, regulations, and norms targeting the development of trustworthy AI. Such mechanisms (e.g. regulations, risk reporting requirements) should draw clear red lines around certain known and demonstrably harmful AI system features that unfairly impact Canadians and erode trustworthiness. Those include but are not limited to AI system features that: induce emotional dependence; induce addiction; employ deceptive practices; target children; undermine democratic norms; or illegitimately concentrate power.

Enforceable governance mechanisms must also require meaningful and transparent risk analysis on the part of private actors, specifically for systems that pose risks to a significant number of Canadians or pose a high risk to some Canadians. This risk analysis should consider a broad range of sociotechnical risks, such as those mentioned above.

3.5. What frameworks, standards, regulations and norms are needed to ensure AI products in Canada are trustworthy and responsibly deployed?

(i.e. governance mechanisms for AI oversight; assurance of product integrity and ethical compliance; priority areas where trust issues are most acute)

This question implies that frameworks, standards, regulations, and norms are needed to ensure the trustworthiness and responsible deployment of AI products in Canada. We agree. However, we note that Canada's current publicly stated position is that we do NOT need AI-specific legislation, which would act as the cornerstone of enforceable frameworks, standards, regulations, and norms. The global AI industry is currently replete with companies adopting weak guardrails as ethics-washing tactics. It is folly to assume corporations will voluntarily adopt meaningful trustworthy AI practices. Therefore, we strongly advise that the Canadian AI strategy establish the need for strong, reasonable and enforceable frameworks, standards, regulations and norms.

The starting point for AI trustworthiness is compelling the public disclosure of known or anticipated risks that AI systems pose to Canadians. Enforceable governance mechanisms must require meaningful and transparent risk analysis on the part of private industry, specifically for systems that pose risks to a significant number of Canadians or pose a high risk to some Canadians. This risk analysis should consider a broad range of sociotechnical risks, including inducing emotional dependence; inducing addiction; employing deceptive practices; targeting children; undermining democratic norms; or illegitimately concentrating power.

Researching practical, enforceable governance mechanisms could be a mandate of SAIRI. This fourth AI research institute would employ interdisciplinary research involving experts in relevant fields (e.g. law, ethics, social sciences, engineers). This would result in robust proposed frameworks, standards, regulations, and norms that are co-developed by technical experts, social scientists, and humanities researchers.

3.6. How can we enhance AI literacy in Canada, including awareness of AI's limitations and biases?

(i.e. workplace training programs or credentials; targeted engagements and public awareness campaigns; international best practices)

The goal of AI literacy in Canada should be to properly align Canadians' trust of AI with the trustworthiness of these systems. Alignment involves giving people the tools and information they need to critically analyze the AI products that are being marketed or suddenly appearing in their existing apps. At the same time, adoption is not the necessary outcome of AI literacy programs. Rather, the goal should be to ensure Canadians can meaningfully assess the risks and benefits of adopting AI. This is central to ensuring Canadians' constitutional rights are respected.

We recommend standing up a research pillar at SAIRI, the AI research institute we have proposed elsewhere in our responses, focused on researching and developing AI literacy strategies. These strategies could include policies, a resource hub, curricula, and public awareness campaigns. All AI literacy strategies should educate Canadians about the risks and benefits of a wide range of AI systems, not just LLMs.

We recommend incentivizing the development of K-12 curricula that critically examines AI. Children are increasingly exposed to AI in social media applications, education technologies, and other tools, and are particularly vulnerable to its effects. We also recommend incentivizing the development of undergraduate courses on the sociotechnical aspects of AI at all public Canadian universities and colleges.

We recommend the establishment of workplace AI training programs to ensure that workers in all sectors are aware of AI-related risks and can critically engage with the deployment of AI tools to assist their work.

AI literacy will support responsible Canadian AI adoption and commercialization. As we have discussed in our other responses, greater awareness of the risks and benefits of AI systems will ensure that Canadians properly align their trust in AI with its trustworthiness. This will ensure that Canadians adopt trustworthy AI technologies and promote ethical AI innovation.

3.7. What can Canada do to ensure equitable access to AI literacy across regions, demographics and socioeconomic groups?

(i.e. collaboration with other levels of government; role of industry and private sector; educational and literacy strategies to foster informed confidence)

Understanding what “AI literacy” looks like in different social and cultural contexts is key to ensuring equitable access across regions, demographics, and socioeconomic groups. Our team cannot define “AI literacy” for all groups, but we recommend further research to this end.

We suggest that this research should be conducted at SAIRI, the fourth research institute we have discussed throughout our responses. This research would include a mandate to identify groups in need of AI education and the type of AI literacy that should be made accessible to each group. SAIRI could work with affected stakeholders to produce the required educational materials, tailor them to each audience, and ensure they are available in the language of each group’s choice through translation.