

Preparing for the AI Crisis

A Plan for Canada

Accelerating developments in AI are enabling important opportunities, but growing concerns about its impacts are early symptoms of a brewing global crisis. Leading AI labs such as OpenAI and Google DeepMind are developing smarter-than-human AI, and could succeed in as little as 1-3 years. Its scale of impact will likely bring widespread upheaval to the economy, geopolitics, and almost all aspects of Canadians' lives. Canada's only option is a global solution, and is well placed to lead one. This white paper outlines what federal political leaders need to do.

Authors

Wyatt Tessari L'Allié (Founder, AIGS Canada), with the support of the AIGS team & expert community

October 2025

Artificial Intelligence Governance & Safety Canada (AIGS) is a nonpartisan not-for-profit and a community of people across the country, working to ensure that advanced AI is safe and beneficial for all. Since 2022, AIGS has provided the federal government with forward-looking [public interest policy](#), and the top recommendation from its [2023](#) and [2024](#) white papers were adopted.

Where is AI headed?

While the first generations of AI-powered systems such as ChatGPT, drones, and digital companions have impressed audiences and reshaped lives, they are just the beginning. With human intelligence staying the same, and artificial intelligence getting better by the day, **we are heading into a world in which AI can outperform us in all domains.** This includes ones where we currently still hold a clear advantage, such as critical thinking, leading organisations, and caring for people. With advances in robotics, physical tasks will not be far behind. [Artificial General Intelligence](#) (AGI) is the industry term for the first systems that can match or outcompete human beings in the real world. They will likely be followed by vastly smarter ones (Artificial Superintelligence, or ASI) shortly thereafter.

AGI and ASI represent the full potential of AI, offering unparalleled opportunities in areas like healthcare and climate, but also unprecedented social upheaval and safety risks. Most notably, AGI represents a turning point in human-AI dynamics. As AI becomes more reliable than us at understanding the world, performing work, and making decisions, we will increasingly defer to it in our personal and societal affairs. This paper therefore focuses on the period before ASI – beyond that, human policy guidance will be less useful than what smarter-than-human systems provide.

There is a wide range of expert opinion as to when the first human-level systems will be built. Current AI systems [already surpass](#) us in areas like breadth of knowledge, [human persuasion](#), and speed of action, and it is difficult to know which capabilities will be added next. However, its economic potential and near-term feasibility have many of the leading tech companies – including OpenAI, Google DeepMind, and Meta stating [AGI is their goal](#) and that [1-3 year timelines are possible](#). While they could be wrong, recent [breakthroughs](#) and [trends](#) give credence to this scenario. Given AGI's outsized implications on human life and that there won't be a second chance to prepare, uncertain but potentially short timelines mean that **governments need to launch preparations now.**




CHART 1

High-level sketch of AI capabilities and impacts over time.

Key areas of uncertainty are U_1 : how many technical breakthroughs are needed to reach AGI, and U_2 : whether the current exponential rate of progress will continue. This determines U_3 : how much time governments have to prepare for smarter-than-human AI and its far-reaching implications.

The Elements of a Global Crisis




TABLE 1 Likely impacts as AI capabilities increase: Timing of the crisis points (coloured areas) are best estimates and will likely vary. Other AI impacts (e.g. misinformation) may also reach crisis levels.

With every aspect of human life shaped by intelligence – work, relationships, health, culture, politics, the environment, etc. – widespread disruption from **AGI development will likely cause overlapping crises.** We outline here some of the most prominent, but others – such as misinformation and deepfakes disrupting public communication – may also reach national crisis levels:

- ⊗ **Major accidents and attacks:** By definition, AGI systems will be able to outperform human beings in the real world. If poor design or accident lead an AGI system to cause harm and refuse to get shut off, there is no guarantee governments and labs can ever regain control. This is why [leading scientists warn of catastrophic risks](#). Moreover, AGI-powered terrorist attacks could be devastating.
- ⊗ **Geopolitical conflict:** AGI could severely [disrupt the current military balance of power](#). There is a potentially decisive first-mover advantage to whoever can build AGI first and harness it to stop rival AI development. States will face enormous pressure to pre-emptively impede rivals from racing ahead.
- ⊗ **Jobs displacement:** AGI and advanced robotics will be [able to perform all economically valuable work better than humans](#) can. Systemic jobs impacts are unavoidable and [some sectors are already impacted](#). Whether AGI leads to a post-work world, or how fast changes will happen, is uncertain.
- ⊗ **Economic and fiscal shocks:** A rapid and sustained rise in unemployment will severely [strain federal income tax revenues](#) while creating major demand for government assistance. It could also cause a rise in mortgage defaults and major shifts in consumer spending, impacting finance and other sectors.
- ⊗ **Public unrest:** While AGI is an unfamiliar and complex global issue, when Canadians see their current jobs disappear and realise there may never be a replacement, or a major incident makes them realise their lives are at risk, public unrest could erupt and AI developers may need to be protected.

From Global Solutions to National Action

Should AGI be built, or should its development be paused?

While it would bring unprecedented benefits, smarter-than-human AI would fundamentally reshape human life, likely cause multiple crises, and risk permanent loss of control. It would be, to say the least, a very grave undertaking requiring careful deliberation. With every living being impacted, **the decision to pursue AGI should involve extensive public, scientific, and international consultations.**

If there is a collective desire to proceed, the world would also be wise to take its time getting the technology right. Building it is at least as technically complicated as building a spaceship, but in the case of AGI, all of humanity will be on the first test flight. Permanent loss of control is possible with the first capable systems, and major job and geopolitical impacts will likely happen before then. Moreover, the current "move fast and break things" culture in the tech sector, combined with well-documented instances of [megalomania](#), [poor cybersecurity](#), and "Feel the AGI" chants at corporate parties, suggests that the companies most likely to build AGI first are currently not trustworthy. However, AI development is a global phenomenon and there are strong military and economic incentives to race ahead. This means that countries that pause unilaterally will be at the mercy of those who don't, and economic sanctions alone are unlikely to stop rogue actors. A global pause would require levels of cooperation, verification, and enforcement never seen before in human history.

Broad approaches to governing AGI development

At its core, AGI governance is a challenge of **1)** human coordination, and **2)** technical competence. Can the world agree on what to do and organise itself accordingly, and can technologists safely deliver the desired systems? Despite the risk that AGI could be developed soon, most governance proposals are still early stage and experimental. We list here the broad strategies commonly discussed:

- ⦿ **Shape lab actions:** Nations directly regulate frontier labs within their jurisdictions, or buy equity stakes and Board seats on foreign ones, in order to better shape their actions in the public interest.
- ⦿ **Reshape the field:** Nations pool resources for a central entity ("[CERN](#) for AGI") that could buy up the talent and compute to lead the field, mitigate global competition, and share the benefits evenly.
- ⦿ **Distribute frontier AI and its benefits:** Give all nations simultaneous access to the latest models, to reduce the odds that any lab, country or misaligned AGI system gains a decisive advantage. Avoid wealth concentration by sharing AI windfalls and ensuring equal access to AI services.
- ⦿ **The first mover pivotal act:** The first lab to build AGI safely, tasks it with governing AGI. This can be an "AI cop" that polices other AI systems and protects against misaligned ones, or that polices human rogue actors (states, labs, individuals) that contravene a global treaty.

Research is urgently needed to better assess these options and propose alternatives. However, with all strategies one thing becomes clear: **Every responsible path forward on AI requires a global deal.** These can range from a statement of shared principles for labs and AGI systems to abide by, to establishing a central agency to build or govern it, to treaties on mitigating the AI arms race and sharing AI's benefits, to globally enforced restrictions on its development. Despite the current rise in geopolitical tensions, every party is at risk of losing control of any AGI system they build and faces the prospect of millions of unemployed. The intensity of the AI crisis could therefore open doors to collaboration that were assumed closed.

The toolkit for governments

Ultimately, [global solutions for AGI will need to be driven by nations.](#) While the influence of leading technology firms has been rising and could reach a tipping point with AGI, the authority to act on behalf of the public, and the legal and military means to do so, remains with national governments. Moreover, international bodies like the United Nations and its [AI initiatives](#) will only be as effective as member states make them. Finally, national governments are also clearly responsible for ensuring their country's overall preparedness for the impacts of AGI development within their borders.

To advance global solutions while building resilience at home, Canada has a number of practical options, such as:

- ⊗ **Capacity building:** Hire and upskill the talent needed for an effective government response, establish task forces and committees to study the issues and direct action, ensure ministries, parties, and jurisdictions are well informed and able to coordinate, harness existing AI.
- ⊗ **Foreign policy:** Host and fund global talks and institutions focused on AGI, present actionable treaty proposals, lobby other nations and build alliances, promote scientific cooperation, support international institutions, apply best practices gleaned from other global crises.
- ⊗ **Security policy:** Strengthen critical infrastructure cyberdefenses, grow AI and drone defense capacity, prepare security agencies for [AI-enabled CBRN](#) proliferation.
- ⊗ **Economic policy:** Ensure sovereign compute capacity, boost productivity via responsible adoption, require AI labelling and privacy protections, ban [unacceptable capabilities](#).
- ⊗ **Fiscal policy:** Prepare the tax base for disruption to employment revenue and increase in assistance need, explore a [Universal Basic Income](#) and AI for efficient public services.
- ⊗ **Investments in research:** Clarify the science, timelines, and impacts of AGI, and develop the technical and policy solutions needed to mitigate the risks and harness the benefits.
- ⊗ **Public engagement:** Clearly and regularly communicate the situation to Canadians along with the plan to address it, and consult extensively to ensure that policies reflect the public interest.

Steps taken so far

Like other countries, the Canadian government has not yet responded to AGI development. Its priority has been to advance AI research and ensure economic competitiveness. Initiatives in this vein included the 2017 [Pan-Canadian AI Strategy](#) (PCAIS), and \$2.4B in [budget investments](#) in 2024 to provide Canadian organisations access to AI compute and to accelerate AI adoption. [This has continued](#) with the new government's focus on harnessing AI to strengthen Canada's economy.

Governance efforts have been focused on pre-AGI systems, such as the 2019 [Directive on Automated Decision-Making](#), 2022 [Bill C-27 AI & Data Act](#) (did not pass), 2023 [guide for government use](#) and [voluntary code](#) for industry. Canada also co-founded [GPAI](#) and approved resolutions at the [G7](#), [UN](#), and on [military use](#). The most forward-looking steps taken were the signing of the [Bletchley Declaration](#), and earmarking \$0.1B of the 2024 AI budget investments for displaced workers and an AI Safety Institute ([CAISI](#)), which now has a [Safe and Secure AI Advisory Group](#).

More promisingly, the new federal government established a [Ministry of AI](#) and [AI Task Force](#), expanding its capacity to act and signaling a new seriousness about the technology. They also provided funding for the UK AISI's [AI Alignment Project](#) and have been [meeting with scientists](#), suggesting that safety concerns are being considered. In short, while Canada has not yet mounted a response to AGI, the new federal government is well placed to do so.

Canada's Opportunity to Lead

"We will have to do things that we haven't imagined before, at speeds we didn't think possible."

Prime Minister Mark Carney, 2025



As with the lead-up to previous crises, there is currently no adult in the room, no global monitoring and accountability system ensuring that AI development is safe and beneficial for all.

On the contrary, AI development is currently the Wild West, with the big US and Chinese tech firms leading [a global race](#) to release the next capability and reap the reward, while [lobbying governments to limit regulations](#). Blind faith in technology prevails, with few asking "Where is this headed?"

Governments and international bodies have for the most part responded with a proliferation of half-measures – laws that are hard to implement, volunteer codes easily ignored – focused on reacting to yesterday's AI. The few forward-looking efforts, such as the [UK AI Security Institute](#) audits of frontier models, the [EU AI Act's Chapter 5](#) requirements, are at best a partial solution. Meanwhile, this year's global [AI Safety Summits](#) series was renamed the [AI Action Summit](#) to focus on economic opportunity, relegating the [International AI Safety Report](#) to a side event. While concern about AGI and short timelines has been openly expressed by leaders such as [JD Vance](#) and [Ursula Von der Leyen](#), no one has presented a serious plan.

This opens a window of opportunity for Canada. Like other nations, our response so far hasn't addressed AGI. And like other nations, we cannot protect our citizens from its risks on our own, making a global solution our only option. But unlike many nations, Canada is unusually stable and educated, well suited to piloting public consultations. We have [top AI talent](#) able to advance the urgently needed technical and policy solutions. Our good international reputation, combined with the fact that we are unlikely to build AGI first, make us an ideal neutral broker. Our new government is ambitiously staffed with top talent including a first-ever Minister of AI, who is already taking steps to strengthen our national resilience. And our new Prime Minister, with extensive global relationships and experience at the helm navigating the financial, COVID and climate crises, is arguably the best placed national leader in the world to spearhead a global response to AGI.

And so we face a choice: we can imitate other nations in the global AI race towards crisis and assume it's not up to us to take responsibility for the whole. We can tell ourselves the US and China will never commit to any global deal and so not even try. We can wait for the crises to develop only to realise once AGI is built it's too late to meaningfully respond. Or we can act now and act fast. We can recognise that every country will need a solution for millions of unemployed workers, and risks losing control of any AGI system they build. And we can focus our attention and resources on getting right a technology that will reshape everything Canadians care about.

To overcome the AI crisis and ensure a future that benefits all, governments and labs will need to present an actionable plan that is big enough to overcome the global forces driving us towards the precipice. Doing so is essential to retain public confidence, AI worker safety, and geopolitical stability.

In this global crisis, **Canada's best contribution is leadership** – to raise our hand and provide a rallying point, to think ahead when others have their heads in the present, to set the example by building resilience at home, and, most importantly, to provide the vision and raw determination needed to overcome the odds and begin shaping a global narrative that Homo Sapiens isn't done yet. We may have less than 18 months to put global solutions in place. Let's get to work.

Preparing for the AI Crisis: A Plan for Canada

Accelerating developments in AI are enabling important opportunities, but the concerns about its impacts on jobs, education, relationships, culture, misinformation, privacy, cybercrime, and warfare are growing and are early symptoms of a brewing global crisis. Leading AI labs such as [OpenAI](#) and [Google DeepMind](#) are developing smarter-than-human AI ([AGI](#)), and could succeed [in as little as 1-3 years](#). Smarter-than-human AI would disrupt almost every aspect of Canadians' lives, and scientists warn of permanent [loss of control](#). Its development is likely to cause a range of global crises including major accidents, geopolitical conflict, systemic job losses, economic and fiscal shocks, and public unrest. Canada's only option is a global solution, and is well placed to lead one. To seize this opportunity and prepare Canada for the AI crisis, the federal government should focus on the following four actions by Q3 2026:

1 Pivot to meet the AI crisis

- ⊗ Establish a permanent task force on AGI chaired by the PM
- ⊗ Involve key Ministers, Opposition leaders, and Premiers
- ⊗ Reassess the priority of files and projects likely to be significantly disrupted by AGI

Rationale: Governments may have less than 18 months to shape a technology that will permanently reshape all other aspects of human life. Much like with COVID in 2020, there are times when the responsible thing for government to do is to focus top leadership on the crisis at hand and re-assess the priority of other files accordingly. Given its wide scope, success will require coordination across cabinet, parties, and jurisdictions.

2 Spearhead the global response

- ⊗ Jumpstart global talks on AGI and host a summit by early 2026
- ⊗ Ground discussions in latest scientific findings and public consultations
- ⊗ Invest in research to advance governance solutions

Rationale: International agreements will be essential to reduce conflict, increase safety, and ensure the benefits of AI are widely shared. Top issues should include whether to attempt a global pause or a "CERN for AGI" initiative and how to avoid winner-take-all scenarios. As a neutral broker with a new PM experienced in global crises, Canada is uniquely well placed to spearhead such talks. Doing so is the government's biggest opportunity to ensure AI is safe and beneficial for all.

3 Build Canada's resilience

- ⊗ **Economic:** Build sovereign compute capacity, diversify supply chains, harness existing AI
- ⊗ **Fiscal:** Prepare for income tax loss, assistance needs
- ⊗ **Security:** Strengthen critical infrastructure, defence capacity
- ⊗ **Societal:** Require AI labelling, ban unacceptable capabilities

Rationale: While domestic action alone cannot protect Canadians against AGI risks, plenty can and must be done to mitigate the secondary impacts - such as AI-powered cyberattacks on critical infrastructure, supply chain disruption from accidents or conflict, fiscal shocks from sustained job losses, and loss of access to foreign compute. By taking the initiative at home, Canada will be in a stronger position to navigate the AI crisis and negotiate from a position of strength. Moreover, the Minister of AI has already taken the lead on some of these measures.

4 Engage Canadians in a national conversation on AI

- ⊗ Nationwide public consultations to educate and consult with the public on core AGI decisions
- ⊗ Use findings to inform domestic and foreign policy decisions

Rationale: Even if the safety risks are managed and the benefits shared, AGI will fundamentally transform people's lives. Managing this transition and ensuring it reflects the will of Canadians will take time and needs to start early. With its relative stability, educated population, and strong AI ecosystem, Canada is in a good position to pilot such a conversation.