




JULY 2026

THINK TWICE

A citizen's guide to artificial intelligence in Canada

Hadrian Mertins-Kirkwood and Rachel Pettigrew



CCPA
CANADIAN CENTRE
for POLICY ALTERNATIVES
CENTRE CANADIEN
de POLITIQUES ALTERNATIVES

Think twice

A citizen's guide to artificial intelligence in Canada

Summary **3**

Introduction: What does artificial intelligence (AI) mean for Canada? **7**

Context: Understanding the AI moment **9**

Implications: AI in the real world **16**

Policy: How governments are responding to AI **31**

Principles: Foundations for public interest AI advocacy **39**

Conclusion: Toward an AI policy for Canada **45**

Appendix: AI policy database **47**

Notes **55**

Summary

Canadians are worried about artificial intelligence (AI). Poll after poll finds that workers, students, artists, parents and other people who live and work in Canada view the proliferation of modern AI tools as a threat to jobs, cognition, democracy, human rights, the environment and more. Yet, to date, these concerns have been given short shrift in the national discourse surrounding AI, which has been dominated by the technology industry with the enthusiastic support of the federal government.

In its new national artificial intelligence strategy, the government acknowledges public skepticism of AI but concludes that “for Canada to thrive in the era of AI, Canadians need to trust in its promise.”¹ That promise includes assumptions of productivity gains, higher-quality services and geopolitical security. It is grounded in the belief, articulated largely without evidence, that AI will necessarily and inevitably reshape society for the better.

This report presents a critical alternative narrative that invites Canadians to think twice about the promise and inevitability of artificial intelligence. It is intended as a foundational resource for concerned citizens and civil society organizations in Canada, which have, to date, managed only a piecemeal response to the AI moment. To mount a more coordinated and effective response to AI in Canada—and to push back on the AI industry in particular—citizens require a better understanding of what AI is, what its implications are, and what should be done about it.

The development and adoption of AI entail a dizzying array of risks and harms, many of which this report addresses in detail, including:

- The unproven productivity benefits of AI tools, which are fueling unprecedented financial speculation;
- The erosion of job quality through workplace surveillance and algorithmic management, as well as the longer-term risks of large-scale job losses;
- The ballooning energy demands of AI data centres and their impacts on the environment and community well-being;
- The erosion of public health care through commercial AI systems;
- The stunting of cognitive development associated with AI use and its consequences for the education system and workforce development;
- The explosion of culture and knowledge production (i.e., “slop”) produced by AI systems that are trained on the works of real creators without consent or compensation;
- The rush to adopt AI in public services and public institutions without assurances of public control over those systems;
- The breakdown of information integrity and the rise of AI-fueled political machines, which threaten the democratic process itself;
- The disregard for human rights in AI applications spanning policing, immigration, warfare, hiring and other sensitive sectors; and
- The attack on Indigenous knowledge systems and data sovereignty by inherently colonial AI systems.

Drawing from these disparate concerns, the report identifies four key themes that encapsulate the current AI moment in Canada. Addressing AI at the personal, institutional and political level requires us to grapple with:

- The indiscriminate adoption of AI systems across all domains of society;
- The systematic devaluation of human thought, judgment and expression in favour of algorithmic output;
- The acceptance of real present harms in exchange for speculative future benefits; and
- The exacerbation—both incidental and deliberate—of existing structural inequalities.

Canada is not the only country grappling with these themes. This paper considers and compares the public policy responses of three key players in the field of AI governance: the United States, China and the European Union. Each is taking a drastically different regulatory approach, and each offers lessons for Canada. However, Canada is clearly lagging all three jurisdictions in terms of AI governance. While the release of a national AI strategy was a necessary first step, Canada still lacks a meaningful legislative or regulatory framework for AI.

It is not too late to put the public interest first in the AI era, but it will require a different approach. Based on the preceding analysis, this report presents six principles for forward-looking, public interest AI advocacy in Canada:

1. Distinguish between AI as technology, industry and ideology.

The term “artificial intelligence” does not refer to anything in particular. By failing to define AI, we conflate the various technologies that fall under the AI umbrella; the multifarious researchers, developers and investors behind specific AI tools; and the foundational beliefs and political ideologies that tie them all together.

2. Recognize the value of original and diverse human thought.

Humanity has inherent value, both practically and morally. We can and must make social, cultural and political choices about where AI usage is appropriate and where it infringes on our humanity in unacceptable ways.

3. Subject the AI industry to democratic oversight and control.

Protecting the public interest in the AI era requires a comprehensive legislative and regulatory framework for AI governance that puts public well-being over private profit.

4. Apply the precautionary principle to AI harms. Some AI-driven harms are already here. Others loom on the horizon. In all cases, we must proceed more cautiously, which means slowing the development of AI systems and/or pausing their adoption until there is conclusive evidence supporting their safety and efficacy in a given context.

5. Prioritize augmentation over automation in workplace AI deployment. A pro-worker AI—one that supports workers rather than replacing or subjugating them—is possible, but only if workers have a say in how AI tools are deployed in their workplaces.

6. Develop AI systems that solve real problems in the public interest. Artificial intelligence, as a category of technology, has genuine potential for advancing the public interest. To realize those benefits will require not only stronger AI governance, but also a commitment to developing AI tools with greater participation from workers, governments and the broader public.

These six principles are intended, first and foremost, to serve as organizing principles for citizens and civil society organizations in Canada, but they also serve as policy principles for Canadian governments committed to the effective governance of AI systems and the AI industry.

According to the federal government, Canada must pursue the widespread adoption of artificial intelligence systems “with urgency.”² This report urges Canadians to think twice. To the extent that the rise of modern AI offers potential benefits, it also entails serious costs. To ensure the AI era serves the public interest, we must move slower, more carefully and with a clearer sense of the technological future we aspire to.

Introduction

What does artificial intelligence (AI) mean for Canada?

As AI tools proliferate in our workplaces, schools and homes, it is a question that is increasingly top of mind for concerned citizens. On the one hand, we are told by technology companies that the deep integration of artificial intelligence into private enterprise, public services and our daily lives will unlock profound benefits. From health care to education to business productivity and beyond, the AI industry promises to deliver broad and unprecedented economic and social welfare.³ In response, our governments are now racing to “drive the adoption of artificial intelligence across Canada’s economy and society” as quickly and as comprehensively as possible.⁴

On the other hand, the rapid rise of AI is driving public fear and mistrust. Many of the productivity benefits associated with AI are still speculative, but they are premised on ultimately automating entire domains of human labour. Moreover, the proliferation of unregulated commercial AI systems poses potentially serious risks to privacy, sovereignty, human rights, cognitive development, the democratic process and other vital public priorities.⁵ All the while, the financial and environmental bill associated with AI training and processing continues to grow to eye-watering heights.

Compounding the multifarious impacts of artificial intelligence today is a deep and fundamental uncertainty about the future. Even among AI experts and within the tech industry, there is significant disagreement

about whether these systems work as advertised and how they will evolve over the coming years.⁶ How, and whether, consumers and institutions adopt AI adds another layer of uncertainty. The as-yet-unrealized potential for government regulation to re-shape the current AI moment is merely the cherry on top of a very shaky cake.

Nevertheless, for all the uncertainty, coming to terms with artificial intelligence is absolutely essential. Workers, students and other concerned citizens in Canada need to be equipped to confront this moment as pro-AI pressure from industry and governments mounts. There may be productive, pro-social applications of the various tools that fall under the AI banner, but we must approach AI development and adoption with a clearer sense of what problems these tools actually solve and what risks they entail for individuals and society alike.

This paper offers a crash course in artificial intelligence for concerned citizens in Canada. It begins with a high-level overview of the current AI moment, including the main technologies behind modern AI systems and the current shape of the AI industry. We then break down the current and potential implications of AI technologies across a wide variety of domains in the Canadian context, including labour, the environment, education, health care, culture and democratic institutions. We evaluate the regulatory approaches that have been taken so far in Canada and around the world to confront these implications, for better or worse, from which we identify best practices and notable gaps. Finally, drawing from the preceding analysis, we arrive at a set of six forward-looking principles for the responsible development and adoption of AI in Canada.

The goal of this paper is to establish a common conceptual framework for public engagement and advocacy on AI issues. To move forward collectively, we require a shared understanding of what artificial intelligence means, what principles should guide future AI development, and what role our governments and public institutions have to play. For concerned citizens and civil society groups, this paper offers a stable place to start in the midst of a deeply unstable moment. Although this paper does not make specific policy recommendations, the appendix outlines more than 60 policies that concerned citizens and organizations may consider in their advocacy efforts.

Context

Understanding the AI moment

Artificial intelligence refers generally to any computer system that takes data (input) and then produces either predictions or content (output) with some degree of autonomy from human oversight. In so doing, these systems perform tasks that would otherwise require human intelligence and decision-making.

In practice, this definition is extremely broad. Since the phrase “artificial intelligence” was first coined in the 1950s, it has been used to refer to a wide variety of different technologies, including everything from weather forecasting models to chess bots to spellchecking software. This imprecision is a key source of confusion today. When we hear the term “AI,” it is not always clear whether it refers to the general concept of (semi-)autonomous computing systems or to a particular technology that falls under the AI umbrella or to a class of consumer products being marketed by technology companies.

To make matters worse, “artificial intelligence” is itself a misleading phrase.⁷ Whereas “artificial” implies that these systems exist outside of human influence, they are fundamentally a product of human labour. These systems are designed by human engineers, trained on human data and refined by human users. Conversely, the term “intelligence” implies a human-like cognitive process, even though these systems—capable as they are of performing many human-like tasks—operate in a fundamentally different way than our own conscious brains. Researchers prefer to use terms like “deep learning” or “machine learning” to describe specific types of AI without the same connotations around intelligence.

Nevertheless, this paper uses the term “artificial intelligence” because it is broadly understood to refer to a particular set of technologies that have come to prominence in recent years with potentially transformative social and economic consequences. Where applicable, we use more precise language to refer to different categories of AI systems.

Categories of AI systems

Modern AI systems are sophisticated pieces of software that employ a handful of fundamental technical architectures.

Symbolic (or logic-based) AI involves a computer system that follows a set of formal rules, similar to traditional programming. Historically, these have also included so-called “expert systems” that draw inferences from a specific body of knowledge. For example, an AI system given the rules of chess can make deductions about optimal moves consistent with the rules of the game. The first few decades of AI development primarily used logic-based approaches and many AI systems today still do, especially in domains such as science or finance, where variables can be controlled and accuracy matters. However, these systems are often task-specific and thus struggle with messy real-world applications, such as interpreting or inferring natural language.

Machine learning (or pattern-based) AI involves the automated processing of vast amounts of data to identify patterns. For example, a machine learning system trained on works in the English language will come to recognize that a capital letter typically follows a period, even without specific instruction in the formal rules of English grammar. Machine learning has been the main area of focus for AI development for the past several decades and it is the basis for many modern AI applications, such as social media algorithms, facial recognition and chatbots. These systems can handle the messiness of the real world better than symbolic approaches, but machine learning algorithms are fundamentally illogical. These systems do not understand what they are predicting and are therefore prone to making errors. The capabilities of these systems are often described as emergent, which means they have strengths and weaknesses that cannot be predicted in advance.⁸

Deep learning (or multi-layer) AI is a subset of machine learning that attempts to replicate some of the complexity of the human brain. These systems use “neural networks” made up of layers of machine learning systems that can pass information between specialized nodes, allowing

for more sophisticated pattern recognition than a simple machine learning system alone. While deep learning has a longer theoretical history, the modern boom in deep learning only began in 2017. Many applications of machine learning today, including advanced chatbots, are products of deep learning developed in the past decade.

Neurosymbolic AI attempts to combine the logic of symbolic systems with the flexibility of machine learning systems to mitigate the weaknesses of each. For example, coding assistants increasingly combine a (pattern-based) chatbot interface with a (logic-based) back-end code base for producing new code. Some experts view neurosymbolic systems as the “third wave” of AI after symbolic and machine learning systems, but, to date, they account for a minority of applications.⁹

The preceding AI architectures are the foundation for a wide variety of AI applications, many of which will be discussed throughout this paper. To understand the current AI moment, however, we need to specifically address one particular application: generative artificial intelligence.

Generative AI systems are a type of deep learning system that can create new multimedia content, such as text, images, audio or video, when prompted by a user. These systems do so by looking for patterns in their training data that align with the prompt and then extrapolating forward. For example, a generative AI image system trained on the works of Salvador Dali can be prompted to “create a painting in the style of Dali.” The output will be an image that reflects the training data—the colours, textures, compositions and so on that appear most often in Dali’s body of work. While the output is fundamentally an amalgam of its data sources, it is still an “original” image to the extent that it is not a direct reproduction of any particular Dali painting. The same principles are at work in deepfakes—the false reproduction of real peoples’ voices and appearance based on past recordings.

Large language models (LLMs) are a subset of generative AI systems that specifically work with text. LLMs are particularly important because they are the technology behind the chatbot-based AI systems, such as ChatGPT, Copilot, Claude and Gemini, that have driven widespread AI awareness and adoption since 2022. These models have proven to be increasingly capable of completing complex linguistic, creative and problem solving tasks previously assumed to be impossible for machines.

The aptitude that LLMs have for natural language creates an especially convincing illusion of consciousness that makes these systems easy to anthropomorphize, but, like all machine learning systems, LLMs are inherently illogical and, by any human definition, inherently unintelligent. They do not understand language or the real world to which

that language refers—they merely predict the next word in a sentence in a sensible but fundamentally unintentional way. LLMs are prone to “confabulation” (also referred to as hallucination), which means they fabricate or misrepresent information and otherwise make basic factual or logical errors that a human never would. They struggle with reasoning, in part because they lack working memories.¹⁰ Generative AI systems like LLMs are also inherently backward-looking, which is to say they are limited by the data they were trained on and perform more poorly in novel circumstances and edge cases.

Yet, for all their limitations, LLMs and other generative AI systems have proven to be transformative. Even if they are not truly intelligent, their capacity for convincingly emulating human thought raises profound questions about the nature of human cognition and the value of human intellectual labour. As this paper shall discuss, the consequences are already widespread and compounding quickly.

As a final technical note, it is necessary to draw a distinction between **prompt-based** AI systems, such as chatbots, and **agentic** AI systems. A prompt-based system requires a user’s input at every step—you ask a question, and it provides an answer. Agentic systems, on the other hand, take a user’s general input but then make independent decisions about which specific steps are necessary to complete the task. An agentic system may be given control of a user’s email and calendar, for example, to proactively schedule meetings. Agentic systems are, at least in theory, considerably more powerful than prompt-based systems and are thus a growing focus for AI research and commercialization.

Dimensions of the AI industry

Similar to the concept of cyberspace—a virtual networked environment that exists outside the material world—artificial intelligence is often understood as an abstract or ephemeral phenomenon. When AI is discussed in news articles and policy reports, for example, it is often accompanied by images of robotic faces floating through abstract spaces, overlaid with math and geometry, rather than by photos of the human workers and data centres that build and sustain these systems.¹¹ However, AI today is very much a product of the real world, and understanding the AI industry and the physical infrastructure behind it is key to grappling with its implications.

The AI industry is made up of a complicated network of research institutions, technology companies, chip manufacturers, power utilities, financiers and other stakeholders. To the extent that it can be measured (independent of the rest of the technology industry), the AI industry was valued at an estimated US\$400 billion in 2025.¹² Most of that value is being captured by U.S.-based big tech companies, such as Microsoft, Alphabet (Google) and OpenAI (ChatGPT), as well as the manufacturers of the computer chips that power AI data centres, such as Nvidia. While Canadian researchers are responsible for many key breakthroughs in the field of AI, Canadian firms have not commercialized AI to a similar extent. For example, Canada's largest AI-focused firm Cohere was valued at just US\$7 billion at the time of writing compared to OpenAI's US\$852 billion valuation. The value of OpenAI, like other AI companies, is largely speculative, based on potential future revenues and inflated by circular financing.¹³ Nevertheless, private investors are banking on the AI industry being worth as much as US\$5 trillion within a decade—roughly equivalent to the entire economy of Germany or Japan.¹⁴

Modern AI systems work by processing enormous volumes of data, both during the development of new AI models (i.e., the “training” process) and for the operation of those models (i.e., the “inference” process). All of that processing occurs in dedicated data centres, which are large warehouses that contain thousands of specialized computer chips and consume large amounts of energy to process data quickly. The rapid expansion of AI services in the past few years has outstripped the growth of data centre capacity, leading to a gold rush of new chip manufacturing and data centre construction around the world. Total investment in new data centres in 2026 may reach half a trillion dollars.¹⁵ While the data centre boom has, to date, centred on the United States, it is increasingly spilling over into Canada and elsewhere as governments seek to attract investment in major infrastructure projects.

Even though AI is an increasingly capital-intensive industry, human labour plays a key role at every stage of development and operation. The data that AI models are trained on is collected from human outputs, including essentially all of the text and multimedia content ever posted on the internet (typically without the knowledge or consent of creators). Much of that data is then manually cleaned, sorted and tagged by precariously employed, low-wage human workers under poor working conditions, including prison labour.¹⁶ The AI models themselves are developed by human engineers based on the work of human researchers in academia and industry. The data centres that power the models are built and maintained by human workers. Finally, AI applications are

tested and refined by human users. However, technology companies are increasingly trying to automate these steps. For example, some AI models are now being co-developed with AI coders and models are being trained on AI-generated “synthetic” data.

From a business perspective, the AI industry is taking the spend-now-earn-later approach that is common to Silicon Valley startups. Despite attracting unprecedented levels of private investment and attracting hundreds of millions of users, most AI services are losing billions of dollars per year as the costs of developing and operating AI models outstrip revenues.¹⁷ The dire need to raise those revenues explains, in part, the rapid integration of AI tools into many existing digital services accompanied by the aggressive marketing of AI companies. To justify their massive valuations, these companies are trying to show investors that their products are being used and will eventually turn a profit. The need for operational funding also explains the rush for private AI companies, such as OpenAI and Anthropic, to go public.

Assumptions driving the AI moment

As we shall discuss in the next section, measurable improvements in productivity due to artificial intelligence systems remain elusive. The gold rush of investment into the AI ecosystem is simply not justified by its real-world benefits to date. Instead, four key assumptions are driving the current AI moment.

First, AI companies (and many governments and employers) assume that AI will eventually deliver massive productivity gains. For example, Anthropic, which is the company behind the Claude chatbot, claims that AI adoption will double the rate of productivity growth in the U.S. over the next decade.¹⁸ While independent analyses are less optimistic, across a range of forecasts, the use of generative AI is expected to deliver labour productivity gains of, on average, about 17 per cent.¹⁹ These gains, if realized at scale, would indeed add trillions of dollars to the global economy over the coming decades, but they remain largely speculative.

Second, AI proponents argue that these systems are neutral and objective and therefore inherently superior to human judgment in many domains, independent of economic efficiency. For example, AI may be seen as less biased in areas such as hiring and the delivery of public services.²⁰ In practice, true neutrality is a myth, and AI systems are inherently biased based on their training data and engineering.²¹ Yet the

allure of algorithmic objectivity remains a key argument for accelerated AI adoption.

Third, many AI proponents argue that the world is in a global AI race, and that the economies that develop and adopt AI the fastest will win this race.²² The AI race narrative is typically presented as a great power competition between the United States and China, which is a view that the U.S. government has adopted explicitly.²³ While the Chinese government has rejected the rhetoric of an AI race, its actual AI policy reflects a similar urgency to lead the world in AI development and adoption.²⁴ Other countries, including Canada, have adopted this framing, as have many businesses and public institutions that worry about being left behind by the technology.

In practice, it is unclear what “winning” or “losing” the AI race actually means. There is no automatic end point for technological development, nor is there a guaranteed commercial monopoly for any economy that first achieves particular AI capabilities. Indeed, the lack of a technological “moat” for generative AI systems means they are more vulnerable to competition than many previous disruptive technologies.

Fourth, and most fundamentally, the current AI moment is predicated on the assumption that technological innovation is inevitable and, moreover, that it necessarily leads to social progress. The myth of technological progress is not new, but with AI it has taken on a new life.²⁵ The framing of technological development as a moral good that all societies must embrace—or else be left behind—reproduces a colonial hierarchy in which the Global North sets the pace and direction of human progress and others are positioned as perpetually catching up.²⁶ This narrative is now fully centred on AI, and it is being perpetuated by a handful of mostly American corporations and a growing number of governments around the world. A fear of stifling innovation is commonly cited by AI proponents in industry and government as a reason to delay or weaken regulation of the sector, regardless of its harms.²⁷

As we turn toward an analysis of the real-world impacts of AI systems, we cannot lose sight of these pervasive assumptions, which are often invoked as justifications for the otherwise negative consequences of these technologies.

Implications

AI in the real world

Like electricity or the internet, many artificial intelligence systems are general purpose technologies, which means they have implications across all economic sectors and domains of human life. This grand scope presents a research challenge in itself as different domains have attracted different degrees of empirical attention. The labour and environmental consequences of AI, for example, are relatively better studied than the social, psychological and cultural implications.

To make matters worse, the novelty of many AI technologies means there are few longitudinal studies into their effects. In many contexts, it is simply too soon to determine, with confidence, what these systems are doing to the world. Future predictions are even more uncertain since there is no historical precedent to anchor them, and it is unclear how fast AI will improve, how businesses and consumers will respond, and what regulatory approaches governments will take in the years to come.

With these caveats in mind, this section provides a high-level overview of the current implications of modern AI systems across a wide variety of domains, in no particular order. It draws on emerging research and expert analysis to map the developing terrain while acknowledging that significant uncertainty remains and that there may be additional implications beyond those discussed here. The section concludes with an analysis of common themes across these domains.

Productivity and the economy

The most frequently cited argument for AI adoption is enhanced economic productivity. As noted above, increasing the rate of labour productivity growth by even one percentage point would add trillions of dollars to the global economy. In Canada specifically, research by the Vector Institute predicts that AI will add \$300 billion to the Canadian economy over the next 10 years.²⁸

Although that figure appears significant, it amounts to less than one per cent of GDP, which is hardly transformative. To make matters worse, empirical studies of AI's impact to date have not found that AI is delivering even these modest benefits. For example, a widely cited MIT study found that 95 per cent of enterprise adoption of AI systems is producing no measurable economic return.²⁹ Statistics Canada has similarly concluded that "there is no statistically significant direct association between AI adoption and productivity" among Canadian firms to date.³⁰

The fact that AI is not delivering measurable productivity benefits so far does not mean that it never will. As the technology improves and best practices are developed, businesses may yet realize those benefits. Nevertheless, fears of an AI-driven financial bubble loom over the industry.³¹ If AI-driven productivity fails to materialize at the scale and pace expected by investors, the financial crash of the AI industry would dwarf the dot-com bubble of the early 2000s. AI-connected firms account for a major share of stock indices, so a crash in the sector would have cascading effects across the economy, including for institutional investors, such as pension funds, and retail investors.

Labour and jobs

Despite popular fears of AI-driven job losses, the main impact of AI in the workplace to date has been a reorganization of how work is controlled, managed and valued rather than the broad displacement of workers. Statistics Canada reports that among firms using AI in 2025, 89 per cent saw no change in employment levels, while nearly half reported only small reductions in tasks.³² For their part, labour unions have expressed a greater immediate concern with the impacts of AI on job quality rather than on job quantity.³³

Underpinning these shifts is an expansion of workplace surveillance as AI systems now routinely collect data on worker location, pace,

communications and physical behaviour. Employers are deploying AI to increase output and reduce labour costs, significantly impacting migrant workers and new Canadians, who are concentrated in sectors where workplace surveillance is more intensive, including warehousing and platform work.

“Algorithmic management” is the growing trend of human managerial tasks being handed off to AI systems. For example, in transportation and logistics, AI systems are being used to track and evaluate how workers drive while inward-facing cameras monitor drivers’ eyes to ensure they are watching the road at all times.³⁴ In office environments, AI can monitor keystrokes, screen activity and email content, transforming routine data into automated performance evaluations that may influence promotion, discipline or termination.³⁵ These systems can encode bias and discrimination based on race, gender or immigration status, while also reducing worker agency, displacing human management and suppressing union organizing.³⁶

Algorithmic tools also shape workers’ access to employment and earnings before they ever set foot in a workplace. In recruitment, AI screening systems are being used to filter candidates based on keyword matching rather than substantive qualifications, often bypassing strong applicants before any human review occurs.³⁷ In gig and platform work, AI compensation algorithms can identify which workers are likely to accept lower pay and adjust earning opportunities accordingly.³⁸ The threat of automation and a tightening job market weakens workers’ bargaining power, with those in AI-exposed sectors more likely to accept lower wages or worse conditions rather than risk replacement.

Nevertheless, over the long term, structural risks to the labour market remain a serious concern. Task level analyses suggest clerical work, data analysis, operations monitoring and writing carry the highest automation risk, while the trades, care work and roles requiring social perceptiveness show the lowest risk.³⁹ In Canada, three quarters of public sector jobs in business, finance and administration are highly exposed to automation of at least some tasks.⁴⁰ Exposure to AI task automation does not necessarily equate to job losses, and researchers disagree on what share of jobs is likely to be displaced by AI in the coming years.⁴¹ However, even if the total number of jobs stays the same, AI may still lead to a significant redistribution of work and workers, with widespread social consequences.⁴²

Because many at-risk occupations are disproportionately held by women, new Canadians and recent graduates, these shifts risk narrowing career pathways and advancement opportunities for those already facing

barriers in the labour market. Statistics Canada data from early 2026 found that workers under 30 saw roughly half the employment growth of workers aged 30-49, with entry-level job postings declining at twice the rate of senior roles.⁴³ Although it is too early to determine to what extent AI is responsible for rising youth unemployment, it is almost certainly the case that junior roles, such as software developers, are being (or will be) automated before senior roles, such as software engineers. Young workers are thus the canary in the coal mine for AI labour market impacts.

Energy and environment

AI is driving one of the largest infrastructure expansions in modern history through the rapid construction of data centres. These facilities are highly energy intensive. In 2025, AI data centres consumed roughly two per cent of global electricity, a figure projected to more than double by 2030.⁴⁴ Although renewable generation capacity is expanding, data centre construction, often completed within two to three years, is consistently outpacing new clean energy development in North America.⁴⁵ As a result, much of the near-term electricity demand for AI infrastructure is being met by natural gas, coal and other carbon-intensive energy sources, in direct conflict with the necessity of industrial decarbonization to meet national and global climate targets.

A single hyperscale data centre can draw over 100 megawatts (MW) of power, equivalent to the annual electricity use of approximately 100,000 households.⁴⁶ In the U.S., communities in close proximity to data centres have experienced electricity price hikes of over 200 per cent.⁴⁷ Recent polling finds that two-thirds of Canadians expect AI data centres to drive up electricity prices here as well.⁴⁸ Rising consumer prices are due, in part, to utilities offering data centres subsidized industrial rates to attract investment while the infrastructure upgrades required to serve them, such as new transmission lines, substations and backup capacity, are financed through rate increases passed on to households. Data centres are rarely required to publicly disclose their energy consumption, making independent assessment of these costs difficult.

Proponents argue that AI could deliver environmental benefits through improved resource efficiency, climate modelling and precision agriculture.⁴⁹ However, like many other benefits of AI, these gains remain largely speculative to date. The AI applications most likely to deliver environmental benefits, such as targeted climate modelling tools,

are narrow and task-specific, bearing little resemblance to the large generative AI systems driving the current infrastructure build-out.

Beyond energy, the AI supply chain includes critical mineral extraction, hardware manufacturing and the disposal of toxic electronic waste. This contributes to water contamination, biodiversity loss and greenhouse gas emissions.⁵⁰ Data centres also require large volumes of water for cooling their servers, with global consumption projected to rise from roughly 560 billion litres annually to more than 1.2 trillion litres by 2030.⁵¹ Data centres consume far less water than they do electricity, but many of these facilities are being built in water-stressed regions, which increases risks of water scarcity.⁵² As with electricity consumption, there is little transparency around data centre water use.

In general, the environmental costs of AI are disproportionately borne by rural, Indigenous and Global South communities, deepening existing patterns of environmental injustice.

Public health

Public health care systems face rising costs, staff shortages and aging populations. AI is being positioned as a response to these pressures. For example, machine learning systems can analyze patient histories, genetics and lifestyle factors to recommend individualized treatments and detect early complications.⁵³ At the population level, predictive tools can forecast disease outbreaks and identify at-risk groups early, enabling faster interventions and more efficient use of resources.⁵⁴ During the COVID-19 pandemic, for example, AI was used to analyze population data to identify optimal locations for vaccine sites.⁵⁵ Since 2022, DeepMind's AlphaFold has predicted over 200 million protein structures, opening new pathways for treating cancer, Alzheimer's and neglected tropical diseases. Its database is publicly available, giving researchers worldwide access to tools that previously required years of expensive laboratory work.⁵⁶ In diagnostics, AI is improving consistency in lung cancer screening, where early detection remains one of the most difficult clinical challenges.⁵⁷

These promising applications tend to share a common feature: they apply AI to well-defined, replicable tasks, such as pattern recognition in medical imaging or protein structure prediction, where consistency and scale are advantages. There is less evidence to support AI use in areas involving complex human judgment, contextual interpretation or

emotional support, such as clinical settings or mental health care, even though users are increasingly turning to AI chatbots for health support.⁵⁸

In many health contexts, concerns about safety and equity are significant. Health data is among the most sensitive information people generate and health systems report hundreds of data breaches annually.⁵⁹ In many cases, that data is collected and sold to train AI systems without patients' knowledge or consent.⁶⁰ The deployment of AI-enabled care depends on cost and the digital divide, risking a two-tier system where lower-income populations receive low-cost automated services as a substitute for higher-quality, human care (with or without AI assistance).⁶¹ Diagnostic systems trained on non-representative data can produce less accurate results for women, racialized communities and people in the Global South, embedding into clinical decision-making the same patterns of bias documented across AI systems more broadly.⁶²

AI tools deployed in clinical settings have also proven unreliable. For example, a 2025 audit procured by Ontario Health found that AI scribe systems fabricated clinical information in nearly half of cases, 60 per cent recorded incorrect medications, and 85 per cent missed key details about patients' mental health.⁶³ The same concerns extend to consumers seeking medical advice from general-purpose commercial AI chatbots, which may make harmful recommendations leading to negative patient outcomes.⁶⁴ In mental health contexts specifically, the systems risk exploiting what researchers call the "Eliza effect", which is the tendency for users to form emotional attachments to AI systems and overestimate their therapeutic capacity.⁶⁵

AI development remains concentrated among a small number of firms, placing critical health infrastructure under private control while commercial pressures accelerate deployment before risks are fully understood. To the extent that AI presents opportunities for Canada's public health care system, the benefits will largely depend on how much oversight and control the public sector retains.

Education and cognitive development

Chronic underfunding and ballooning class sizes are straining public education in Canada. Governments and school boards are pressuring teachers to turn to AI tools to support assessment, lesson planning and differentiated instruction.⁶⁶ Meanwhile, students are turning to AI chatbots for homework, advice and emotional support, raising concerns

about safety and dependence.⁶⁷ This pattern echoes longstanding critiques of social media use, where platforms designed to maximize engagement have been linked to addiction, anxiety and declining academic achievement. Ontario school boards have launched multi-billion-dollar lawsuits against major social media companies on exactly these grounds.⁶⁸

Emerging research suggests that reliance on generative AI tools may produce measurable cognitive harm. Preliminary findings from a MIT study found that students who regularly used ChatGPT showed reduced neural engagement, memory recall and ownership over their own thinking, even when the output scored well.⁶⁹ A related concern is sycophancy. Because many AI systems are designed to be agreeable, they confirm what users already believe rather than challenge it.⁷⁰ One study found that participants relying on AI were nearly five times less likely to reach correct conclusions than those who did not, while growing more confident in their wrong answers.⁷¹ Overreliance on AI risks stunting the cognitive development that education exists to support, including critical thinking, independent problem solving and the capacity to reach conclusions without assistance. When AI displaces mentorship, productive struggle and collective sense-making, it can strip learning of what makes it transformative.

According to the OECD, the core problem is not AI in education as such, but the dominance of off-the-shelf chatbots that were never designed for learning.⁷² Purpose-built tools co-created with teachers, students and communities could expand educational access for under-resourced students and reduce barriers for those with disabilities or learning in a second language.⁷³ However, the Brookings Institution global task force on AI in education found these conditions remain an exception, while commercial products face no binding obligations to demonstrate improved learning outcomes before entering classrooms.⁷⁴

OpenAI CEO Sam Altman has described a future in which intelligence is a utility that people purchase by the meter, signalling an intent to monopolize access to knowledge itself.⁷⁵ The long-term risks are substantial, not only for individual well-being but also for the workforce and for our very democracy, which depends on a citizenry that can think critically and independently.

Culture and knowledge production

Artists and cultural workers in Canada are among the first to feel the direct impacts of generative AI.⁷⁶ Across federal copyright consultations, union submissions and a 2024 National Summit on AI and Culture, Canadian creators have converged on three core demands: consent over how their work is used to train AI systems, credit when it is, and compensation when it generates value.⁷⁷

Major technology companies face active lawsuits over their use of copyright-protected material, while Canada's copyright legislation has not kept pace. For example, in 2025, a coalition of Canadian media organizations, including the *Toronto Star*, *Globe and Mail*, Postmedia and CBC, commenced a copyright infringement action against OpenAI in Ontario, with the court confirming jurisdiction over a foreign AI company for the first time.⁷⁸ Because Canada still has no standalone copyright bill, creators are left to challenge infringement on a case-by-case basis.

The stakes extend beyond individual rights to the broader cultural ecosystem. Nearly one in four businesses in Canada's information and cultural industries already uses generative AI, with adoption highest among the largest firms.⁷⁹ AI-driven algorithms threaten the discoverability of Francophone, Indigenous and independent Canadian content, deepening existing asymmetries between platform owners and creators.⁸⁰

At the knowledge level, AI generates outputs through pattern recognition rather than comprehension. The result is a proliferation of content that can appear authoritative while reproducing inaccuracies, bias, mediocrity and what some call "AI slop."⁸¹ AI-generated material now comprises over half of new online content, making quality and credibility harder to locate amid a growing volume of AI-generated imitation.⁸²

Public services and public institutions

Canadian governments are rapidly integrating AI into the public service under the pretense of streamlining operations and modernizing service delivery. Delays in Employment Insurance payments, immigration processing, tax support and benefits delivery, for example, have material consequences for people's lives, and slower public service delivery undermines public confidence. Processing backlogs are a legitimate problem that well-designed AI tools could help address.⁸³

However, there are also structural limits to what AI can do. The federal government's own AI register documents over 400 instances of AI use across 42 agencies, with roughly one-in-five systems generating content that informs public decision-making.⁸⁴ Most of these systems were deployed without impact assessments, meaning their effects on service quality and accuracy remain largely unknown. Unlike human judgement, AI can only look backwards, bound by whatever patterns it was trained on and unable to adapt to circumstances it has never encountered.⁸⁵ Offloading core functions to AI risks hollowing out the expertise and institutional memory that an effective public administration requires over the long term.

Automated decision-making in public services is already causing demonstrable harm. The Canada Revenue Agency's (CRA) AI chatbot "Charlie" cost taxpayers over \$18 million and a 2025 report from the Auditor General found that it gave incorrect answers two-thirds of the time.⁸⁶ At Immigration, Refugees and Citizenship Canada (IRCC), the use of automation has coincided with a quadrupling of federal court cases since 2020, with lawyers documenting decisions that show no engagement with the evidence filed. In one case, IRCC acknowledged using AI to reject a permanent residence application based on job duties the algorithm had fabricated. In Quebec, an algorithmic overhaul of the social assistance system that eliminated assigned human agents is believed to have contributed to a patient's death after a man received incorrect eligibility information without a caseworker who knew his file.⁸⁷

These technical limits are compounded by unequal access. AI-enabled service delivery assumes universal internet access and digital literacy, conditions that do not hold across the population. Rural, low-income and aging Canadians continue to rely on in-person and phone-based services.⁸⁸ Human judgment and contact are also needed for accountability, dignity and recourse. When these interactions are automated, there is a risk of displacing the relational core of the public service.⁸⁹

Underneath these concerns sits a more fundamental question of sovereignty. Digital sovereignty requires governments to account for where public data goes and who controls it. Forty per cent of federal AI systems are developed by external vendors, the majority American-owned,⁹⁰ meaning sensitive public data including health, immigration and tax information may be accessed through privately owned software governed by foreign jurisdictions. Canada's *AI Strategy for the Public Service* deployed Microsoft Copilot for internal tasks,⁹¹ while the national

AI strategy consultation used three U.S.-based LLMs to analyze over 64,000 public responses without clear disclosure on data handling.⁹²

Canada's recent investments in sovereign AI infrastructure are nominally intended to address these concerns, but they will be insufficient if procurement continues to default to American-owned platforms while austerity erodes public sector capacity that institutions depend on to remain adaptive and legitimate. Canada's *AI Strategy for the Federal Public Service* commits to responsible adoption but stops short of the procurement standards needed to ensure digital sovereignty in practice. The European Public Service Union has called for frameworks that require bidders to disclose ownership structures, mandate impact assessments before deployment, prevent public data from being repurposed or used to train AI systems without authorization, and ensure public authorities retain control over the systems they procure.⁹³ Canada's procurement framework offers no such guarantees.

Democracy and governance

AI systems are contributing to new pressures on information integrity in ways that destabilize democratic participation. The problem begins with how information is curated. Algorithms owned by private companies can influence what millions of people know about an election, a policy or a geopolitical crisis—not based on accuracy or truth but on engagement.⁹⁴ These curated feeds narrow the range of perspectives people encounter and deepen political polarization.⁹⁵ Political campaigns exploit AI-driven micro-targeting techniques to deliver individualized messages calibrated for persuasion rather than addressing a shared public sphere.⁹⁶

Due to a 2024 British Columbia Supreme Court ruling, B.C. is the only province where provincial privacy law applies to federal political parties. Canadian parties elsewhere operate largely outside of privacy legislation, meaning there is limited insight into how parties collect voter data, use AI to target citizens or share that data with third-party technology firms. Exposure to AI-generated “deepfakes”—highly convincing fabrications of real people or real events—compounds the problem by making people more susceptible to believing misinformation or losing confidence in authentic evidence altogether.⁹⁷ Taken together, these conditions erode democratic life by producing a fragmented public, less able to engage in healthy debate, build shared values or hold power to account.

The governance of AI itself reflects a growing disconnect between governments and the public. Two thirds of Canadians express low trust in federal institutions,⁹⁸ 85 per cent want AI to be regulated,⁹⁹ and 79 per cent support taxing companies that replace workers with AI.¹⁰⁰ Yet policy-making in Canada is prioritizing competitiveness and alignment with industry interests, as we will discuss in the next section. Canada's fall 2025 AI task force spent only a month studying the issue and it lacked adequate representation from civil society groups, labour unions and citizens.¹⁰¹ Canadians cannot meaningfully consent to technologies they have no say in shaping, and until democratic participation is built into how AI is governed, regulated and procured, decisions will continue to be made by those with the most to gain from them.

Human rights

AI has been used in ways that can support human rights and humanitarian response. For example, satellite imagery analysis enabled organizations such as Amnesty International to document attacks on civilian infrastructure and identify potential safe corridors in conflict zones.¹⁰² AI-powered translation tools and chatbots are being developed to help refugees navigate complex asylum systems.¹⁰³

However, a 2026 Amnesty International briefing concludes that the most widely used generative AI systems are fundamentally incompatible with international human rights law, finding that their reliance on unlawful web scraping constitutes mass invasion of privacy by design.¹⁰⁴ An investigation of OpenAI's ChatGPT by the Privacy Commissioner of Canada similarly concluded that the company "did not respect Canadian privacy laws."¹⁰⁵

Across justice, policing, immigration and warfare, government deployment of AI is transforming decision-making in ways that undermine core human rights protections. In policing, for example, facial recognition technologies and algorithmic risk assessments demonstrate documented racial bias and inaccuracy, contributing to wrongful arrests and expanded systemic discrimination.¹⁰⁶ Automated decision-making has been embedded in Canadian immigration and refugee processing for over a decade and predictive analytics systems influence who is flagged for secondary screening, delayed, or denied entry at borders, often with limited transparency and constrained avenues for appeal.¹⁰⁷

Globally, the stakes are most stark in warfare. For example, Israel is using AI-assisted targeting systems to generate strike lists, with human oversight reduced to just 20 seconds per decision, despite acknowledged error rates and authorization of significant civilian casualties.¹⁰⁸ This exemplifies moral outsourcing, where life-and-death decisions are delegated to systems devoid of conscience and incapable of ethical reasoning. Across domains, AI is accelerating harm, expanding the scale and speed at which violence and exclusion can be administered, while eroding due process and democratic oversight.

These examples are among the most acute, but human rights concerns arising from AI extend across employment, child welfare, fraud detection and any context where automated systems make or inform consequential decisions about people's lives.

Indigenous data sovereignty

Dominant AI paradigms encode Western ways of knowing as universal standards—quantifiable, objective and algorithmically legible at the expense of knowledge systems that are relational, place-based, embodied and intergenerational.¹⁰⁹ Generative AI systems scrape and commodify knowledge at scale, assuming it can be abstracted from context, relationship or ownership. Absorbing Indigenous Knowledge Systems (IKS) into AI systems strips them of the ethical, spiritual and governance dimensions that give these systems life and meaning.¹¹⁰ Indigenous researchers call this “data colonialism”—the appropriation of human and ecological knowledge as a raw material, commodified for capitalist computation, replicating historical logics of dispossession and erasure into the digital realm.¹¹¹

In response, Indigenous data sovereignty movements are asserting the right to control how community knowledge is collected, used and governed. Rooted in Article 31 of the UN *Declaration on the Rights of Indigenous Peoples*, the “CARE Principles” for Indigenous Data Governance (Collective Benefit, Authority to Control, Responsibility and Ethics) offer a critical alternative to data frameworks that prioritize unfettered open access and interoperability over Indigenous rights and collective interests.¹¹²

Emerging work in Indigenous algorithmics is moving beyond reductive data models toward systems designed around reciprocity and land-based ethics. Some communities are deploying AI on their own terms.

For example, remote sensing enables the monitoring of illegal logging on Indigenous territories;¹¹³ the MaxEnt model helped the Heiltsuk Nation identify culturally significant plants;¹¹⁴ and Canada's NRC Indigenous Languages Technology Project is supporting language revitalization through community-driven design.¹¹⁵

Grounded in Indigenous governance and knowledge systems, AI holds the potential to become a tool for self-determination and more just relations between humans, technologies and land. The challenge is not to integrate IKS into existing AI systems, but to dismantle the knowledge hierarchies those systems are built upon. Without that structural change, AI deepens digital colonialism.

Common themes

The rise of the modern artificial intelligence era has different implications in each of these domains, but there are clear through lines to our analysis. We identify four core themes that help answer the question of what AI means for Canada today.

The first theme is the indiscriminate adoption of AI systems across all domains of society. It may be the case that AI applications are ultimately of net benefit in some contexts, such as medical research, but not in others, such as primary health care. Determining where AI has value requires a degree of experimentation, but it also requires a sense of what problems these technologies are supposed to solve and an awareness of what kinds of solutions have been tried before. In many contexts today, AI is, instead, a solution in search of a problem. AI companies and governments promote AI adoption for its own sake rather than offering clear evidence for added value. As a consequence, both personal and institutional AI use is rising, but it is not necessarily delivering benefits for those users, institutions or the people they serve.

Indeed, the exceptionally high failure rate for AI pilots in the private and public sector is not necessarily an indictment of the underlying technology but, rather, of an unwarranted enthusiasm for AI adoption in contexts where it is not needed or may be actively harmful. Premature AI adoption at scale also introduces risks of its own, such as compounding security vulnerabilities, dependence on private, foreign technology and the deskilling of the workforce due to widespread cognitive offloading. These are avoidable problems if AI is developed more slowly and adopted with greater human oversight.

The second theme is the comprehensive devaluation of human thought, judgment and expression in favour of algorithmic output. The dehumanization and commodification of labour (physical and intellectual) is not new. It has proceeded apace since the industrial revolution. However, AI accelerates the trend to new extremes. No human domain is now beyond the scope of automation. Many sectors previously assumed to be “safe” from AI, such as creative and care work, are increasingly exposed.

Behind the application of AI systems to these domains is the fundamental assumption that the output of labour has value but the labour itself does not. For example, it does not matter whether a human doctor or an AI tool prescribes a medication to a patient, provided the prescription itself is passably similar. The same can be said of a piece of legislation, a business plan, a novel or a citizenship decision. This logic both dehumanizes labour and serves to depoliticize it. Human workers can organize, bargain and strike, which has historically forced employers and states to negotiate the terms under which work is conducted. When labour is automated, that leverage disappears with it and so does any clear accountability for the outputs of AI systems.

In practice, AI tools are still inferior to human expertise across many domains today. Moreover, the assumption that outputs matter but process does not fails to recognize that the labour itself is often the point. For example, writing an essay, a report or an email forces productive engagement with a topic, regardless of the quality of the final document. However, the gap between AI tools and human experts is narrowing, and as AI systems become increasingly capable of producing outputs that meet or exceed human standards, the value of intellectual labour becomes increasingly invisible and/or economically irrelevant. As a consequence, AI systems facilitate the centralization of power in the hands of employers and technology companies at the expense of workers and communities. The devaluation of human thought and the attendant concentration of power matters in both practical and moral terms, but, so far, it has not prominently featured in debates around AI.

The third theme is the tension between the speculative future benefits of AI systems and their real present harms. Where AI does offer theoretical benefits, as in business productivity, those gains are neither guaranteed nor immediate. In many cases, those gains hinge on continued technological development and the widespread social and commercial adoption of AI systems. Even in a best-case scenario, it may take many years for the AI-powered utopia promised by proponents to truly arrive.

Yet AI is already having material impacts on people and communities, and many of those impacts are negative. The erosion of the digital information environment due to AI is already a crisis, as is the use of unregulated and unmoderated AI tools in the education system and many workplaces. The proliferation of data centres powered by fossil fuels is actively undermining decarbonization efforts. Workplace surveillance and algorithmic management is accelerating. Policing and citizenship decisions are increasingly taken without human oversight.

To be truly socially beneficial, AI systems must meet high thresholds of capability and trustworthiness. To cause harm, they need only be accessible. Right now, AI systems are far more accessible than they are capable or trustworthy.

The fourth, and final, theme is the exacerbation of structural inequalities that are being driven by AI development and adoption. For example, to the extent that commercial AI use can or will improve productivity, the benefits will mainly accrue to employers at the expense of workers whose labour is automated, intensified or appropriated. AI in the public service, including in health care, shifts public resources away from public institutions and toward private firms. AI in sensitive social sectors, such as immigration, housing and policing, further dehumanizes and disempowers already marginalized groups and amplifies existing patterns of inequity across race, gender and class.

Ultimately, the biggest beneficiaries of widespread AI adoption are the private, mostly American tech companies that own and operate these systems. As they become more powerful, they are also able to exercise greater influence over regulators, users, workers and the rest of the private sector. The developers of AI tools have implied that the automation of labour (and the broader diminishment of human labour power) is a key goal of theirs, which means AI adoption is as much a political project as it is a technological project.

These are not the inevitable consequences of artificial intelligence itself. Instead, they reflect an AI industry that is produced by, and thus reinforces, a political and economic system of global capitalist exploitation. The difference between AI that surveils workers or empowers workers, for example, is political, not technical. What AI could do under different conditions of ownership, governance and democratic control remains an open question.

In the next section, we turn from implications to responses. We ask how governments around the world are responding to the AI moment and assess whether and how those policy measures are grappling with the themes identified above.

Policy

How governments are responding to AI

The three largest economies in the world—the U.S., China and the European Union—are all grappling with the AI moment in starkly different ways, which provides a useful natural experiment as well as offering diverse points of comparison for Canada. In this section, we assess each jurisdiction in turn, asking how they have responded to AI at the regulatory level and what each governance model reveals about that jurisdiction’s AI strategy. We recognize that the regulatory environment is evolving quickly, so the pictures presented here are more of a snapshot than a definitive analysis of each jurisdiction’s approach to AI governance.

We conclude with a summary of AI regulation in Canada and an assessment of how Canadian AI governance differs (or aligns) with public policy elsewhere.

United States

The trajectory of U.S. AI governance shifted sharply under Donald Trump’s 2025 *AI Action Plan*, rolling back the Biden administration’s emphasis on safe, secure and trustworthy AI. This shift is often framed as deregulation, but it is better understood as aggressive state intervention on behalf of the American tech industry.

A recent report from Data & Society describes this as the emergence of a “Big AI State” that is operationalized through three pillars. First, derisking data centre and energy development through executive orders that allow developers to bypass environmental review, to exploit federal lands, and to access public financing, limiting community input while socializing the costs of private AI investment. Second, the U.S. is exporting its AI tech stack as geopolitical infrastructure, offering full-service packages (compute, data systems, models and applications) through government-backed partnerships that condition market access on regulatory alignment with U.S. priorities. Third, the federal government is acquiring direct equity stakes in strategic firms throughout the AI supply chain, such as American rare earth mineral producers.¹¹⁶

The U.S. remains a fragmented regulatory environment, with states and municipalities actively advancing AI governance measures in the absence of federal guardrails. For example, over 350 bills related to AI and work were introduced in the 2025-26 legislative session alone, covering algorithmic management, automation disclosure and bias protections.¹¹⁷ The Trump administration responded with a December 2025 executive order aimed at preempting state AI regulation.¹¹⁸ Communities near proposed data centre sites are also organizing, with resistance expanding into litigation, calls for moratoria and electoral organizing.¹¹⁹

What is emerging is not an innovation economy but a technology infrastructure empire—over 5,000 data centres, more than ten times any other country, and \$286 billion in private AI investment in 2025 alone.¹²⁰ The emphasis on rapid infrastructure buildout and market expansion reflects the broader pattern of indiscriminate AI adoption, where deployment is prioritized ahead of social value. The alignment between state policy and private capital accelerates the concentration of economic and political power, deepening the inequalities already associated with AI development. The *AI Action Plan* removes references to misinformation, equity and climate change from federal frameworks, purporting that AI should pursue “objective truth,” free of ideological bias.¹²¹ Yet framing bias mitigation as ideological overreach does not produce unbiased AI. Instead, it produces unaccountable AI. Rather than mitigating the risks of AI, the U.S. model is actively upscaling them. However, as resistance from elected officials, civil society and affected communities continues to mount, the foundations of the Big AI State are more contested than the elite consensus would suggest.

China

China's strategy for AI is threefold: first, to compete with the U.S. for technological leadership; second, to reduce dependence on Western digital infrastructure; and, third, to extend state AI capacity across governance, services and security.¹²² Since the release of its *New Generation AI Development Plan* in 2017, China has backed AI-related technologies with large-scale investment and coordinated action across government, industry and research.¹²³ The release of DeepSeek-R1 in 2025, a generative AI model developed at a fraction of Western costs, established China as a legitimate competitor in frontier AI development.

Chinese firms often release models as "open weight," meaning they can be downloaded and further developed. This accelerates global adoption of Chinese AI and challenges the U.S. strategy of controlling AI diffusion through export restrictions. China is embedding AI across society—from chatbots in local government services and education systems to factory robotics in industrial production—at a scale and pace unmatched elsewhere. The 2025 State Council Opinions on the "AI+" Initiative have set aggressive targets for deepening AI integration across energy, health, transportation and public services.¹²⁴

China regulates AI through a layered set of sector-specific rules oriented toward state priorities rather than individual rights. Measures include mandatory pre-deployment safety assessments, algorithmic registration requirements and content labelling rules requiring watermarks on AI-generated content.¹²⁵ AI systems are also prohibited from generating content that opposes the ruling party or the broader political system.¹²⁶ Enforcement is tightening through significant penalties and a revised cybersecurity law that brings AI directly under government oversight.¹²⁷ While these measures reflect strong state control, they also produce forms of protection that are often absent in less regulated markets. For example, China's 2026 *Interim Measures for Anthropomorphic AI* prohibit virtual intimate relationships with minors, restrict content that could encourage harmful emotional dependence and require clear labelling so users know they are interacting with a machine.¹²⁸ UNICEF has recognized these provisions as a significant step forward for child protection in AI regulation.¹²⁹ In another notable case, a Hangzhou court ruled that AI replacement does not automatically justify terminating a labour contract, establishing that the costs of technological transformation cannot be shifted onto workers.¹³⁰

China is advocating for global AI governance through a proposed "World Artificial Intelligence Cooperation Organization."¹³¹ Whereas the

United States has abandoned human-centred multilateral coordination, China is moving to set the agenda. However, the tension in China's approach is impossible to ignore. While simultaneously using AI to drive economic growth and technological leadership, China maintains tight control over information and data, and it operates the most extensive AI-enabled censorship and surveillance apparatus in the world. State coordination allows for more direct intervention in deployment and, in some cases, stronger legal protections for citizens. Yet it concentrates both governing authority and the informational advantages of AI within the state, subordinating human autonomy to state priorities and leaving little room for the public deliberation that democratic governance requires.

European Union

The European Union has pursued a dual strategy on AI: leading in rights-based regulation while simultaneously investing in rapid AI adoption and industrial capacity. These ambitions sometimes pull in opposite directions, and the consequences are visible in the architecture and rollout of the EU's groundbreaking *AI Act*.

The EU has built a substantive policy framework for regulating the digital economy. The *General Data Protection Regulation* (GDPR) establishes strict rules for how personal data is collected, processed and shared.¹³² The *Digital Markets Act* and *Digital Services Act* extend this foundation, targeting platform monopolies and user rights, respectively, including recent investigations into cloud providers like Microsoft Azure and Amazon Web Services, whose infrastructure underpins much AI development.¹³³ Whether a handful of American companies can act as gatekeepers to Europe's AI future is now an active regulatory question.

At the centre is the *EU AI Act*, the world's first comprehensive AI law. It introduces a tiered, risk-based framework that bans certain applications of AI systems outright, such as social scoring systems that rate behaviour to determine access to services or opportunities.¹³⁴ It also imposes strict requirements on high-risk systems in employment, health care and public services, including mandatory audits, technical documentation, public registration and penalties reaching up to €35 million, or seven per cent of a company's global turnover, for the most serious violations.¹³⁵ The act reflects a precautionary logic. Systems must demonstrate safety and accountability before entering the market.

Yet the *AI Act* contains structural exclusions that undermine its stated commitments. It allows a full exemption for military, defence and national security uses. In migrant contexts, for example, this enables the use of prohibited and high-risk tools like biometric identification, AI lie detectors and predictive scoring to be deployed without registration or public oversight.¹³⁶ Worker protections remain limited to notification and transparency, falling short of rights to appeal or refuse workplace AI deployment.¹³⁷ Environmental impacts and supply chains are also largely unaddressed.¹³⁸

The competitiveness pressure further constrains EU regulation. Through the Digital Omnibus, the council and parliament are negotiating compliance delays that critics argue create a window for risky systems to reach the market before oversight applies.¹³⁹ The same process is advancing proposals to narrow the GDPR's definition of personal data, responding to industry arguments that existing privacy protections limit training data availability and disadvantage EU developers relative to less-regulated competitors. The EU has also actively pursued industrial policy through the *AI Continent Action Plan*, targeting €20 billion in annual AI investment, signalling that the AI race narrative is actively shaping its priorities.¹⁴⁰

Despite these limitations, the EU policy approach remains the most explicit attempt to govern AI in line with public interest principles. The EU exerts outsized global influence through the "Brussels Effect", whereby global companies align their practices with EU legislation.¹⁴¹ When the GDPR passed, for example, major platforms restructured data privacy practices globally, rather than maintaining separate systems for EU users.¹⁴² The *AI Act* is expected to operate similarly, making regulatory authority the EU's primary instrument of global AI influence. However, while the EU seeks to mitigate the harms of AI deployment, its parallel push for rapid adoption reproduces the same speculative logic identified earlier. As a result, while the EU is attempting to moderate some AI risks, it is not fundamentally altering the conditions producing those risks.

Canada

Canada entered the AI era as a research leader, with institutions that helped shape the technology now transforming the global economy.¹⁴³ The Canadian Institute for Advanced Research (CIFAR), which was first launched in 1982, has played a central role in convening researchers,

industry and governments in service of AI development over the past four decades.¹⁴⁴ Canada was the first country to adopt a national AI strategy in 2017, which was administered by CIFAR, and a revised national AI strategy was released in 2026.¹⁴⁵ However, after years of false starts and unfulfilled promises, Canada still has no federal legislation in place to govern AI systems. Rather than develop a formal regulatory and governance structure for AI, the federal government has prioritized public-private partnerships with the AI industry. Surveys consistently rank Canadians among the most cautious and skeptical populace on AI adoption, but those concerns are not reflected in Canada's approach to AI governance.¹⁴⁶

The failure to regulate is, in part, a structural problem. AI governance in Canada is mainly the purview of Innovation, Science and Economic Development Canada (ISED), which encompasses the Minister of Artificial Intelligence. The AI minister and ISED share a mandate to promote economic growth and to regulate the marketplace.¹⁴⁷ This dual mandate produces a predictable outcome, with policies that prioritize industry access over public oversight. For example, of 84 federal and provincial AI initiatives identified by academic researchers, 50 focus on promoting industry and innovation, compared to just 19 addressing social and workforce impacts.¹⁴⁸ The result is a regulatory landscape where standards about how AI shapes work, health and public life are either set by industry or are entirely absent.

The failed *Artificial Intelligence and Data Act* (AIDA) was emblematic of the government's approach to AI regulation. The AIDA was drafted within ISED, with no public consultation before being tabled in 2022.¹⁴⁹ Of more than 300 stakeholder meetings held afterward, only nine included representation from civil society, who, along with researchers and labour unions, called for its withdrawal.¹⁵⁰ The AIDA was eventually abandoned when parliament dissolved in 2025.

The death of the AIDA and the election of a new federal government in spring 2025 created an opening for a different approach, but the new government has fallen into familiar patterns. In fall 2025, the government began consultations on a new AI strategy, which included both an expert task force, composed mainly of industry voices and industry-adjacent academics, and a public call for comments. The public portion of the consultation drew more than 11,000 participants—the largest in ISED's history—signalling strong public demand for involvement.¹⁵¹ However, ISED privileged the views of its industry-dominated task force in reaching the conclusion that the government should prioritize innovation and growth of the AI industry rather than lead with democratic governance.¹⁵²

The pro-industry orientation of the consultation was reflected in the government's new AI strategy, *AI for All*, which was released in June 2026. Of the plan's six pillars, five are either directly or indirectly focused on scaling up AI development and increasing AI adoption across Canadian society. Among other goals, the strategy seeks to increase business adoption of AI from 12 per cent today to 60 per cent within a decade. The government also hopes to reach millions of students with "AI literacy" initiatives designed to encourage greater AI use. The strategy includes billions in subsidies for the AI industry, mainly focused on providing increased domestic compute capacity, although the sums remain relatively small compared to other public industrial spending.

The sixth pillar of the 2026 strategy is focused on AI safety. Among other commitments, the government promises to establish a "fundamental right to privacy," to protect elections from misinformation, and to require the watermarking of AI-generated content. However, no details or timelines are provided. Moreover, not all of the AI-related concerns raised in the public consultations or by civil society experts are acknowledged or addressed in the strategy. For example, *AI for All* makes no mention of the environmental harms of AI infrastructure, of the risks to cognition and mental health of AI use, or of the theft of copyrighted works to train AI models.¹⁵³

Indeed, the new federal strategy is guilty of perpetuating all of the assumptions driving the AI moment that we discuss above. The strategy assumes AI will "unlock productivity gains" and "improve quality of work." It speaks of the "urgency" of AI adoption. It also appeals to the myth of technological progress in calling for Canadians to "trust in [AI's] promise." The strategy provides little evidence to support any of these assertions. Even if the government delivers on its vague promise of safety regulation, the vision for AI described in this document is a far cry from a comprehensive, public interest governance framework.

The strategy's aspirations of digital sovereignty are also difficult to reconcile with the government's current practices. The federal government is deeply reliant on Microsoft and other American technology companies for computing systems and software, including for AI tools. Real sovereignty requires technical and legal safeguards, including audit rights, breach notification requirements, control over who can access Canadian data and the ability to refuse foreign government requests hardwired into enforceable contracts.¹⁵⁴ Microsoft's own representatives have acknowledged under oath that they cannot guarantee data sovereignty to non-American customers, even when data is stored domestically.¹⁵⁵ By equating digital sovereignty with domestic data centre

construction, the federal government is leaving the software layer—where data is created, processed and accessed—largely ungoverned.¹⁵⁶

Canada's federal institutions have consistently chosen dependency and perceived economic competitiveness over accountability when the two have come into conflict. That pattern is unlikely to change while AI governance remains housed in a department whose mandate is to promote the industry it is meant to regulate. Canada is not without leverage. As a middle power, it has substantial procurement power, research capacity, existing legal frameworks and strong public demand for regulation. Yet governments have not yet used that leverage to advance a public interest approach to AI development and adoption.

In the absence of federal legislation, provinces have begun to fill in the gaps unevenly. For example, Manitoba has moved to address algorithmic price discrimination, Quebec requires transparency in workplace AI deployments, and Ontario requires disclosure around AI use in the hiring process. On the other hand, Alberta has exempted data centre proposals from provincial impact assessments, which has triggered legal action in the province.¹⁵⁷ For example, Sturgeon Lake Cree Nation issued a cease-and-desist to the provincial government over the Wonder Valley data centre project, stating they were never consulted and learned of the proposal through a press release, despite the project sitting on land shared under Treaty 8.¹⁵⁸ Unless and until the federal government provides legislative and regulatory clarity around AI issues, Canada's AI governance landscape will remain a contested patchwork.

Each of the four jurisdictions examined in this section has taken a different approach to AI governance. The U.S. is prioritizing market power and tech imperialism, China is prioritizing centralized state oversight and coordination, and the EU is advancing comprehensive safeguards with carve-outs for industrial development. Canada's AI governance approach, lacking a formal federal framework of its own, is largely shaped by these international dynamics rather than shaping them to reflect Canadian values.

Nevertheless, across all four jurisdictions, AI adoption is treated as an end in itself, with policies designed around industrial capacity and geopolitical competition rather than the needs of the people most affected by these systems. The public interest principles that should guide AI governance are largely absent from these approaches, which is precisely why articulating them matters.

Principles

Foundations for public interest AI advocacy

This paper sets out to establish a common conceptual framework for public interest advocacy around artificial intelligence issues in the Canadian context. For citizens, workers and civil society organizations to productively engage on AI issues, they require actionable principles that can serve as a foundation for specific recommendations and demands. Based on our analysis of the real-world implications of AI in Canada and informed by existing government responses to AI, we offer six core principles for consideration.

Note that these principles are deliberately high level. They are intended to be salient for citizens, civil society, employers and governments at all levels in Canada. We discuss next steps for specific policy recommendations in the concluding section to this report.

1. Distinguish between AI as technology, industry and ideology

The term “artificial intelligence” does not refer to any one phenomenon in particular. That ambiguity creates conceptual problems and practical risks for both AI advocacy and AI governance.

As a technology, AI refers to a category of computer systems that share a variety of underlying architectures. As an industry, AI refers to a network of researchers, developers, investors and other stakeholders that seek to build and/or bring AI technologies to market. As an ideology (or as a political project) AI refers to a belief in the inevitable transcendence of human capabilities (especially human labour) by automated computer systems.

These are overlapping conceptions, but they are distinct in important ways. When the Canadian federal government, for example, advances an “AI for all” agenda, it is unclear whether the ultimate goal is technological accessibility, industrial development or socio-economic transformation. That ambiguity may reflect a lack of clarity in the government’s own thinking, or it may be a deliberate obfuscation. For example, the federal government evokes the benefits of narrow, sector-specific AI applications (e.g., in agriculture and manufacturing) as justification for a society-wide project of generative AI adoption, including new data centres to power generative AI demand. This equivocation (i.e., bait-and-switch) is only possible because AI is not clearly defined in the strategy.

AI proponents must be clearer about whether they support AI as a technology (and which types of AI systems, specifically), AI as an industry, and/or AI as a matter of principle. For their part, AI critics must articulate whether their concerns stem from the AI technologies themselves, from the structure of the AI industry, or from AI as an ideological or political project.

2. Recognize the value of original and diverse human thought

To address AI as ideology, we must first recognize that original and diverse human thought has value in itself.

Many artificial intelligence proponents claim that by outsourcing your thinking, you reclaim your time. This devalues the process of human thinking by focusing on the outputs of intellectual labour alone. There are contexts in which the ends may justify the means, such as rote administrative tasks, but it is not a universal truth. As various commentators have remarked, we want technology to do our laundry so that we can spend more time creating art, but modern AI tools promise to create art so that we can spend more time doing the laundry.

A related consequence of a handful of AI tools being used by billions of people is that they are intellectually homogenizing. If ChatGPT or Claude or Gemini is viewed as “neutral”, then its voice and perspective—shaped by a small segment of mostly high-income, white, English-speaking, American men—becomes an unassailable default. The illusion of objectivity that these tools offer entrenches the epistemic hegemony of the tech industry and undermines the value of human diversity, including non-Western values and Indigenous knowledge systems.

There are practical reasons for believing in human diversity and originality—a raft of evidence across domains finds that incorporating multiple perspectives leads to better outcomes¹⁵⁹—but it is also a moral position. We can and must make social, cultural and political choices about where AI usage is appropriate and where it is not, and we cannot be afraid to reject AI systems outright where they infringe on our humanity in unacceptable ways.

3. Subject the AI industry to democratic oversight and control

To address AI as industry, we must establish a robust system of public interest AI governance.

Even if the hype surrounding the artificial intelligence industry today does not reflect the current capabilities of AI systems, there is no doubt that these systems are transformative and will be a destabilizing feature of societies and economies moving forward. Any technology this disruptive requires democratic oversight and control.

Oversight means an expectation of transparency from the AI industry. Among other issues, AI companies must be clear how data is collected to train an AI model, whose labour was used in that training and under what conditions, what assumptions the model has been programmed with, and what biases it exhibits in its outputs. It must also be clear how much energy and water these systems require to operate. All industry claims must be independently verifiable.

Democratic control means governing AI in the public interest with binding regulation. That begins by rejecting the primacy of private sector innovation, business investment and geopolitical competition. None of those priorities is more important than the safety and well-being of citizens. It also means ensuring democratic control over the data—both public and private—that is the lifeblood of the AI industry.

To govern effectively, the public sector must have the legislative authority and technical capacity to regulate the industry and to protect Canadians and their data. Preserving regulatory space is especially important given the changing nature of AI. Governments must be prepared to (re)regulate as new technologies, issues and evidence emerge.

4. Apply the precautionary principle to AI harms

To address AI as technology, we must stop AI systems where they are causing harm today and we must slow AI adoption where there is significant risk of harm in the future.

Some of the social and environmental harms that AI systems have created are already obvious and concrete, and those harms must be immediately prohibited. There is no defence for sexually exploitative AI deepfakes, for example. Nor should new physical infrastructure, such as data centres, be built in a manner that entrenches fossil fuel dependence in the face of an escalating climate crisis. Speculative future benefits do not excuse immediate material harms.

In many contexts, however, we cannot say with certainty what the long-term consequences of AI systems may be. For example, whether AI use by students will enhance learning or permanently weaken cognition—and under what circumstances—remains a vital open question. There is a real risk that future generations will lose the ability to perform complex cognitive and creative tasks unless they are deliberately shielded from AI use. Whether AI use in medicine will make doctors more accessible to marginalized groups or widen existing health inequities is similarly uncertain.

Until we have more evidence, especially in sensitive areas such as education, health care and social services, where the risks are significant and the consequences difficult to reverse, we should apply the precautionary principle. That means slowing or pausing the adoption of AI systems until there is conclusive evidence supporting their safety and efficacy.

5. Prioritize augmentation over automation in workplace AI deployment

To address AI as technology, we must also adopt a pro-worker orientation to AI deployment.

Productivity, in economic terms, is calculated simply as output divided by hours worked. Productivity can thus be increased either by increasing output with the same amount of labour or by producing the same amount of output with less labour input. In many cases, AI systems may be technically capable of doing both. Whether we augment workers with these systems to make them more productive, or whether we automate the tasks previously done by those workers, is a political choice.

When AI systems are deployed in workplaces, they should be oriented around the augmentation of the human workforce. However, realizing the productivity potential of an augmentation approach requires investment, training and negotiation, whereas automation promises a shortcut to easy productivity gains. To ensure employers do not default to the easy option when AI systems are deployed, workers must have power in the workplace. Collective bargaining is one of the greatest tools for ensuring workers exercise control over the introduction of AI systems in their workplaces. In non-unionized workplaces, stronger government regulation becomes essential for giving workers tools for managing and pushing back on AI deployments.

6. Develop AI systems that solve real problems in the public interest

Finally, to further address AI as industry, we must, to the extent possible, reorient the development of AI systems themselves toward the public interest.

In many contexts today, AI systems are a solution in search of a problem. AI-powered tools are marketed by the tech industry as productivity or well-being enhancing with little regard for the specific needs of users or of the unintended consequences of indiscriminate adoption. Governments are similarly guilty of pushing AI adoption based on the ideological assumption that AI is a net positive in all contexts—and with no acknowledgment that the adoption of AI systems may be counterproductive or create new problems.

Yet there are contexts in which AI systems offer solutions to real problems or where AI tools open new doors to human flourishing. AI systems in medical and scientific research can accelerate new discoveries, for example, while AI systems in the utilities sector can help the electricity grid operate more efficiently. To properly identify those contexts, and to ensure they deliver on their potential, we require greater involvement of workers and citizens in the development and deployment of new technologies. Better AI governance models will help, but that's not the only path forward. Building new AI tools from the ground up, with public participation, can also ensure they better serve the public interest.

We should not be afraid to experiment with sector-specific AI applications co-developed with workers and the public sector, provided deployment is consistent with the principles of precaution and augmentation described above.

Conclusion

Toward an AI policy for Canada

The rapid rise of novel artificial intelligence tools is undermining the very value of human intellectual labour, all the while driving a gold rush of investment into a booming technology sector. Governments in Canada and around the world are all in on AI, which they view as central to the future of the economy. Indeed, the development and deployment of AI systems (and the accompanying physical infrastructure) is being sold as so vitally important for future well-being that it cannot afford to be slowed down by regulation or oversight, regardless of the concerning and increasingly visible harms associated with these technologies.

This paper has attempted to contextualize this tenuous situation and unpack its implications for concerned citizens in Canada. It considered the different types of AI technologies in question, the current shape of the AI industry and the assumptions driving the development and adoption of AI systems. It evaluated the measurable and potential future impacts of AI systems in the real world across a wide variety of domains. It also surveyed the various measures that governments in Canada and around the world have already taken to regulate artificial intelligence.

Based on this analysis, we arrived at six core principles for the responsible development of AI in Canada. We argue that governments, industry and civil society must:

1. Distinguish between AI as technology, industry and ideology,
2. Recognize the value of original and diverse human thought,

3. Subject the AI industry to democratic oversight and control,
4. Apply the precautionary principle to AI harms,
5. Prioritize augmentation over automation in workplace AI deployment, and
6. Develop AI systems that solve real problems in the public interest.

These are not specific policy recommendations aimed at any one government or institution, but there are clear policy implications. Above all else, artificial intelligence must be regulated through some combination of government legislation, workplace negotiation and broader social and cultural activism. The real and immediate harms that AI is causing are inexcusable when the only justification is the speculative promise of long-term benefits. There is also no guarantee that those benefits, if and when they arrive, will serve the broader public interest over private profits in the absence of a strong governance framework.

Specific policy ideas are included for reference and consideration in the appendix to this report. A common theme in the approaches identified in our research is that AI must serve people and not the other way around. Like any tool, whether AI systems ultimately benefit or harm the public is not inevitable. It is a choice that must be made thoughtfully, collectively and freely from the pressures of a self-interested technology industry.

Appendix

AI policy database

Table 1 summarizes specific AI-related policy recommendations that we encountered during our research. They are presented here as a resource in no particular order.

Table 1 / Summary of AI-related policy recommendations

Title	Type	Description
<i>Worker protection</i>		
Workplace AI bill of rights	Regulation, Governance	Require worker consent and union involvement before AI deployment in any workplace. Definitions of harm caused by AI systems should be clearly articulated and developed with workers and their organizations. The bill should mandate transparent communication about monitoring, performance metrics and job impacts and give workers meaningful veto power over surveillance systems, with accessible appeal rights for all automated decisions affecting employment.
Mandatory pre-deployment testing and algorithmic transparency	Regulation, Governance	Require independent third-party audits for bias, accuracy, security vulnerabilities and labour impacts before any AI system deployment in public services and high-impact private sector applications. Mandate public disclosure of training data sources and the logic underlying automated decisions.
Ban algorithmic management	Regulation	Prohibit AI systems that automate or substantially determine discipline or performance evaluation without direct human review by a qualified decision-maker and accessible worker appeal rights.
Worker consent for workplace data collection	Regulation	Require employers to inform workers of all data collected through AI systems and obtain prior informed consent before use. Limit monitoring to occupational health and safety purposes only.
Worker voice in public AI research and development	Regulation, Governance	Require that workers and their unions be meaningfully involved in publicly funded AI research and development processes from the earliest stages of design through deployment. Workers as end-users have practical expertise that improves safety and effectiveness. Any federal or provincial funding for AI innovation should be conditional on demonstrated worker and union participation in governance and decision-making.
Mandatory AI impact negotiations	Regulation, Governance	Require employers to negotiate with unions before introducing AI or algorithmic management systems. Include redeployment, reskilling and social protection for displaced workers in collective agreements.
Productivity gains sharing requirement	Regulation	Mandate that efficiency and productivity gains from AI deployment be shared with workers through wage increases, improved working conditions or reduced working hours. Distribution of gains must be negotiated through collective bargaining.
Job transition fund	Spending	Create a dedicated public fund for workers displaced by automation. The fund should provide income support and community-led job creation initiatives as an alternative to corporate retraining programs.
Recognize AI labour	Regulation	Extend employment standards, labour protections and fair wages to data annotators, content moderators and other workers in AI supply chains. These protections should apply to all workers engaged by Canadian companies, including those located in the Global South.
Sectoral pro-worker AI standards	Regulation	Establish sector-specific standards for AI tools that demonstrably augment rather than replace worker expertise, beginning with health care and education. Public procurement criteria should be tied to evidence of worker benefit.

Title	Type	Description
Using grant-making to support pro-worker AI investments	Spending	Establish a federal grant program to direct public AI research funding toward labour-augmenting technologies that expand worker capabilities rather than automate expertise. Grant criteria should include evidence of worker benefit and legal protections against the commodification of workers' skills and creative output by AI systems.
Ban productivity-based AI evaluation	Regulation	Prohibit public sector AI deployment justifications based solely on efficiency or cost savings. Any proposed deployment must demonstrate improved quality, equity or public benefit before receiving approval.

Democratic governance

Public AI advisory council	Governance	Create an independent body with a mandatory consultative role in federal AI procurement and regulation. Seats should be reserved for unions, environmental groups, Indigenous communities, disability advocates and civil society organizations.
Ban high-risk AI applications	Regulation	Impose an immediate ban on AI applications in predictive policing, welfare fraud detection, border risk scoring, facial recognition in public spaces and automated hiring and firing decisions, pending independent bias auditing, legislative authorization and meaningful public consultation.
Needs assessment requirement	Governance	Mandate that any proposed public sector AI deployment demonstrate that the problem it purports to solve existed before the technology was considered. Public consultation must be conducted on whether AI is the appropriate solution, with evidence of community benefit required before approval.
Public data trusts	Governance	Establish publicly governed data trusts where citizens collectively control how their data is used, with opt-in requirements, community benefit agreements and independent oversight to prevent data from being captured by commercial interests.
Digital safety commission	Governance	Establish an independent federal Digital Safety Commission with enforcement powers over digital platforms and public-facing AI, including authority to conduct algorithmic audits, issue penalties and require corrective action. The commission should be supported by a Digital Safety Ombudsperson mandated to receive public complaints, advocate for users and report publicly on systemic harms.
Personal data protection standards	Regulation	Strengthen existing federal privacy law to guarantee all users of AI and platform services operating in Canada have the right to transfer or delete their personal data. Companies must not retain personal data beyond what is necessary for its original purpose, with sensitive data subject to strict time limits on storage. Companies must publicly disclose the terms under which personal data is stored, shared and retained.
Citizens' assemblies on AI governance	Governance	Create standing, demographically representative citizens' assemblies with structured roles in AI policy design and oversight. Assemblies should include dedicated representation for youth, racialized and Indigenous stakeholders and have binding input into regulatory processes.

Title	Type	Description
Canadian digital sovereignty	Regulation, Governance	Establish Canadian ownership and democratic control over the digital infrastructure and software underlying public services, including data storage, cloud computing and AI systems. Require vendors to disclose corporate control structures and extraterritorial legal exposure as weighted criteria in all public procurement. Mandate preference for open-source software that can be audited and governed by public institutions, consistent with approaches already adopted in France, Germany and the Netherlands.
Government AI expertise capacity	Governance, Spending	Fund dedicated AI expertise within federal and provincial governments to enable meaningful regulatory oversight, procurement standards and evidence-based policy-making independent of industry influence. In addition to technical expertise, effective governance requires human rights, labour, Indigenous rights and equity perspectives to be equally represented in decision-making roles.

Transparency and safety

Strict liability framework	Regulation	Hold AI developers and deployers legally and financially liable for harms caused by their systems, including companies headquartered outside Canada. There should be no corporate immunity shields, and clear standards for damages and redress must apply to all AI systems operating in the Canadian market.
Algorithmic justice audits	Regulation, Governance	Mandate regular independent third-party testing for racial, gender, disability and class bias in all public and high-impact private sector AI systems. Results must be publicly disclosed.
Authenticity requirements	Regulation	Mandate clear labelling of AI-generated content across all public-facing communications, with a particular enforcement focus on education, journalism, electoral content and public services. Penalties should apply where AI-generated content is used without disclosure.
Age-appropriate AI design standards	Regulation	Require AI products likely to be used by minors to meet independent child safety standards before deployment, including prohibitions on behavioural profiling, addictive design features and data collection beyond what is necessary.

Environmental justice

AI climate impact disclosure act	Regulation, Governance	Mandate public reporting of energy consumption, water use, carbon emissions and e-waste for all data centres and AI training operations, with standardized metrics and independent verification.
Cap data centre energy use	Regulation	Set hard limits on data centre energy consumption as a percentage of the grid. Renewable energy sourcing and energy efficiency standards should be required as conditions of operating approval.
Community consent for data centres	Regulation, Governance	Require binding community approval before any data centre development proceeds, with special protections for Indigenous territories and rural areas experiencing water stress. Affected community members must be meaningfully included in all environmental and social impact assessments, with community accounts formally incorporated at every stage. Municipalities must negotiate binding community benefit agreements before approving development, including energy impact assessments, local employment commitments and revenue sharing.

Title	Type	Description
Community capacity for data centre assessment	Governance, Spending	Fund dedicated technical and regulatory capacity within regional, municipal and Indigenous governments to evaluate data centre development proposals. Rural and Indigenous communities must be equipped to assess environmental and social impact assessments, negotiate meaningful conditions and understand the long-term implications of data centre siting decisions. Public funding should ensure communities can access independent technical advisors before any approval process begins.
Data centre moratorium	Regulation	Impose an immediate moratorium on new data centre construction until binding regulatory frameworks are in place to protect communities, workers and the environment from the harms of unchecked AI infrastructure development. The moratorium should apply at the federal, provincial and municipal levels, with exemptions only where communities have provided free, prior and informed consent and enforceable community benefit agreements are in place.
AI supply chain accountability	Regulation	Extend environmental and labour accountability requirements to the full AI hardware supply chain, including critical mineral extraction, chip manufacturing and electronics disposal. Enforceable standards should apply to all Canadian companies operating globally.
AI environmental harm registry	Governance	Establish a publicly accessible national registry tracking environmental impacts of AI infrastructure by site, including energy subsidies received, land use, water consumption, pollution and community complaints.
AI data centre reclamation fund	Regulation, Spending	Establish a publicly administered fund to cover the environmental reclamation and remediation costs of decommissioned AI data centres. Companies operating data centres in Canada must contribute to the fund as a condition of operating approval, ensuring that communities and taxpayers are not left to bear the costs of site cleanup, land restoration and infrastructure removal when data centres close or relocate.

Decolonization

Indigenous data sovereignty framework	Regulation, Governance	Recognize Indigenous communities' collective ownership and control over data about their lands, languages and cultures. Require free, prior and informed consent for all collection, use and sharing of Indigenous data.
Indigenous AI capacity and digital infrastructure	Governance, Spending	Fund dedicated digital infrastructure, data literacy programs and technical capacity within Indigenous communities to enable meaningful exercise of data sovereignty. Sovereignty over data requires the resources and tools for communities to collect, manage, govern and apply their own data on their own terms. Funding must flow directly to Indigenous-governed institutions and be designed in accordance with community priorities.

Public ownership

Public AI infrastructure crown corporation	Governance, Spending	Establish a federal entity to develop and maintain AI tools for public services. The entity will operate with transparent, democratically accountable governance and will prioritize public benefit over commercial return.
Ban privatization of public AI	Regulation	Prohibit the sale or licensing of AI systems developed with public funding to private companies. Publicly funded AI tools must be released as open-source software, with community governance over ongoing development and use.

Title	Type	Description
Community technology cooperatives	Spending	Fund worker and community-owned AI development cooperatives as alternatives to the venture capital model. Support should include grants, procurement preferences and technical assistance for non-profit and solidarity-based AI development.
<i>Economic justice</i>		
Progressive revenue redistribution (AI profit tax)	Regulation	Tax windfall profits generated through AI deployment to fund transition supports, public services, critical digital literacy and community-led alternative infrastructure. Revenue should be directed toward workers and communities most affected by AI-driven displacement.
AI reparations fund	Regulation, Spending	Require companies that have profited from AI systems to contribute to a dedicated reparations fund supporting impacted communities. The fund should include compensation for data labellers, writers and cultural workers whose labour or creative work was extracted to train AI systems without their consent or fair compensation.
Close automation tax loopholes	Regulation	Ensure federal tax policy limits structural incentives that favour AI-driven automation over employment. This includes reviewing tax treatment, capital investment incentives and public subsidy programs to ensure they do not systematically disadvantage hiring and training workers relative to investing in AI software and equipment. Any public financial support for AI adoption should be conditional on demonstrated labour and community benefit.
Antitrust enforcement for AI competition	Regulation, Governance	Direct antitrust enforcement toward concentrated AI development markets to break up monopolistic control over data, compute and distribution infrastructure, creating space for alternative business models more conducive to pro-worker technology design.
<i>Procurement</i>		
Procurement preference for worker and community-owned AI vendors	Regulation, Spending	Establish explicit procurement preferences for worker-owned, cooperative and non-profit AI vendors in federal and provincial government contracting.
Procurement power standards	Regulation	Require the federal government to refuse contracts with AI companies that have records of human rights abuses, union-busting, environmental destruction or failure to pay fair wages to data workers. Mandatory due diligence assessments and public reporting must accompany all AI procurement decisions.
<i>Governance and information integrity</i>		
Algorithmic transparency for political content	Regulation, Governance	Require platforms to publicly disclose how content is ranked, weighted and amplified in political and electoral contexts, with independent auditing and enforceable standards against algorithmic manipulation of public discourse.
Right to AI-free spaces	Regulation	Protect designated spaces where human thought, reasoning and deliberation develop without AI mediation or intervention. Educational environments, civic assemblies, therapeutic settings and democratic consultation processes should be protected from AI substitution by default, with opt-in rather than opt-out mechanisms for any AI tools introduced into these spaces. The capacity for independent human thought is a democratic and social good that requires active protection as AI systems become more pervasive.

Title	Type	Description
Protection of epistemic autonomy	Regulation	Require AI systems to be designed in ways that do not undermine human reasoning and independent judgment. Developers must demonstrate that their systems support informed decision-making rather than manipulate or replace it, with particular scrutiny applied to recommendation systems and generative AI tools used in news, education and civic contexts.
Democratic deliberation protection	Regulation	Prohibit the use of AI to simulate, substitute for or manipulate democratic deliberation including public consultations, citizens' assemblies, electoral debates and legislative proceedings. Submissions and participation must be human-authored and disclosed as such. Governments must not use AI to fulfill consultation obligations, misrepresent public input or rely on AI as the primary means of evaluating submissions when developing policy.
Education and youth		
Pedagogical standards for AI in schools	Regulation, Governance	Require that any AI tool deployed in classrooms meet evidence-based learning standards developed by educators, teachers' unions and independent researchers, with mandatory review of learning outcomes before procurement renewal.
Youth AI cognitive safeguards	Regulation, Governance	Restrict deployment of generative AI tools in K-12 settings without independent safety review and demonstrated evidence of learning benefit. Ongoing monitoring for cognitive and social development impacts must be required as a condition of continued deployment.
AI literacy as public education infrastructure	Spending	Fund a national AI literacy curriculum developed by educators and civil society with a particular focus on advancing equity and strengthening critical thinking.
Public health		
Health data sovereignty act	Regulation	Prohibit the use of patient data to train commercial AI models without explicit informed consent. Patients must have the right to request deletion of their data and all health data sharing agreements with AI companies must be publicly reported.
Representative training data requirements	Regulation	Mandate that AI diagnostic and clinical tools demonstrate representative and bias-tested training data before deployment in Canadian health systems, with particular attention paid to accuracy across race, gender, age and geography.
Preserve human-led clinical care	Regulation	Prohibit AI systems from substituting for human care in mental health, palliative care and patient-facing clinical contexts, ensuring AI functions only as a support tool with a qualified human practitioner accountable for care.
Human rights		
Mandatory human rights due diligence for AI	Regulation, Governance	Require AI developers and deployers to conduct ongoing human rights due diligence throughout the full AI lifecycle in line with the UN <i>Guiding Principles on Business and Human Rights</i> . Due diligence must include human rights impact assessments that identify, prevent and address potential harms. The results of all assessments must be publicly disclosed, including documentation of how identified impacts are being addressed.

Title	Type	Description
Facial recognition moratorium in law enforcement	Regulation	Ban use of facial recognition by police and security agencies pending independent bias auditing, legislative authorization and meaningful public consultation, with any future use subject to warrant requirements and mandatory public reporting.
Mandatory disclosure of policing AI systems	Regulation	Require all law enforcement agencies to publicly disclose what AI tools they use, for what purposes, under what oversight and with what error rates.
Prohibit AI in immigration and refugee decision-making	Regulation	Strictly limit automated scoring and risk-flagging in immigration and refugee processing. All consequential decisions must be made by a human decision-maker, with transparent criteria, meaningful review rights and accessible appeals processes for all affected individuals.
Protection of humanitarian workers from AI surveillance	Regulation	Prohibit the use of AI surveillance tools to monitor, track or collect data on humanitarian workers, legal aid providers, journalists and civil society organizations. Surveillance infrastructure that treats solidarity and humanitarian assistance as indicators of criminal activity undermines the legal right to provide aid.
Prohibition on biometric surveillance of migrants at borders	Regulation	Prohibit the bulk collection and processing of biometric data from migrants, asylum seekers and refugees at points of entry. Biometric data collection in border and immigration contexts constitutes a serious interference with privacy and data protection rights and disproportionately impacts racialized communities. Any biometric data collection must be strictly necessary, individually justified and subject to independent oversight, with data minimization principles applied at every stage.
Drone surveillance restrictions	Regulation	Establish clear legislative limits on state use of AI-equipped drone surveillance, including prohibition on use at protests and public gatherings and mandatory public reporting on deployments. These restrictions apply equally in national security and military contexts, where additional oversight and independent review must be required before deployment.
<i>Copyright and cultural workers</i>		
Training data consent and compensation framework	Regulation	Require consent from and compensation to creators whose work is used to train commercial AI models. Collective licensing mechanisms should be established to distribute revenue equitably to affected creators and cultural workers.
Public interest media fund	Spending	Fund independent journalism and human-authored public interest media as a democratic response to AI content proliferation.
Cultural heritage protection from AI extraction	Regulation	Prohibit scraping of Indigenous, community and public cultural archives for commercial AI training without explicit community consent and benefit-sharing agreements. Special protections must apply to materials held in public institutions, libraries and cultural repositories.

Notes

- 1 Innovation, Science and Economic Development Canada, [Canada's National Artificial Intelligence Strategy: AI for All](#), Government of Canada, June 2026.
- 2 *Ibid.*
- 3 See, for example: Peter Nicholson, ["Industrial Revolutionary: How Artificial Intelligence will Fuel Canadian Productivity and Prosperity,"](#) Public Policy Forum, December 2024; Shahed Al-Haque, Marie-Renée B-Lajoie, Erez Eizenman & Nick Milinkovich, ["The Potential Benefits of AI for Healthcare in Canada,"](#) McKinsey & Company, February 2024; and Andrea Willige, ["From Virtual Tutors to Accessible Textbooks: 5 Ways AI Is Transforming Education,"](#) World Economic Forum, May 2024.
- 4 Innovation, Science and Economic Development Canada, ["Pan-Canadian Artificial Intelligence Strategy,"](#) Government of Canada, 2017.
- 5 Mary Burns, Rebecca Winthrop, Natasha Luther, Emma Venetis & Rida Karim, [A New Direction for Students in an AI World: Prosper, Prepare, Protect,](#) Brookings Institution, January 2026.
- 6 Melanie Mitchell, Yuval Noah Harari, Carl Benedikt Frey, Gary Marcus, Nick Frosst, Ajeya Cotra, Aravind Srinivas & Helen Toner, ["Where Is A.I. Taking Us? Eight Leading Thinkers Share Their Visions,"](#) *The New York Times*, February 2, 2026.
- 7 Kate Crawford, [The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence,](#) Yale University Press, 2022.
- 8 Fabrizio Dell'Acqua, et al., ["Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality,"](#) *Organization Science*, vol. 37 (no. 2), March 2026.
- 9 See, for example: Artur d'Avila Garcez & Luís C. Lamb, ["Neurosymbolic AI: the 3rd wave,"](#) *Artificial Intelligence Review*, vol. 56, March 2023.
- 10 Dan Hendrycks, et al., ["A Definition of AGI,"](#) *ArXiv*, December 2025.
- 11 Michael Chui, et al., ["The Economic Potential of Generative AI: The Next Productivity Frontier,"](#) McKinsey & Company, June 2023.
- 12 UN Trade and Development, ["AI Market Projected to Hit \\$4.8 Trillion by 2033, Emerging as Dominant Frontier Technology,"](#) April 2025.
- 13 Cedric Sam, Rachael Dottle, Agnee Ghosh & Kyle Kim, ["A Guide to the Circular Deals Underpinning the AI Boom,"](#) *Bloomberg*, last updated May 6, 2026.
- 14 UNCTAD, "AI Market Projected to Hit \$4.8 Trillion by 2033."
- 15 Goldman Sachs, ["Why AI Companies May Invest More than \\$500 Billion in 2026,"](#) December 2025.

- 16 Ben Lee Taylor, "[Long hours and low wages: the human labour powering AI's development](#)," *The Conversation*, November 15, 2023; see also: Tech Equity, "[The Data Work Landscape](#)," 2026.
- 17 Dan Milmo, "[Can OpenAI Keep Pace with Industry's Soaring Costs?](#)," *The Guardian*, November 10, 2025.
- 18 Anthropic, "[Estimating AI Productivity Gains from Claude Conversations](#)," November 25, 2025.
- 19 Rosalie Wyonch, "[From Hype to Output: How AI Investment Translates to Real Productivity Gains](#)," C.D. Howe Institute, April 9, 2026.
- 20 See, for example: Justin Longo, "[The Transformative Potential of Artificial Intelligence for Public Sector Reform](#)," *Canadian Public Administration*, vol. 67 (no. 4), December 2024.
- 21 Jillian Fisher, et al., [Toward Political Neutrality in AI](#), Institute for Human-Centered AI (Stanford University), September 2025.
- 22 See, for example: Eric Schmidt, "[How the U.S. Can Win the AI Race](#)," *Time*, December 11, 2025.
- 23 Yi-Ling Liu, "[How Silicon Valley Sold Washington an AI Race](#)," *Transformer News*, May 8, 2026.
- 24 Laurie Chen & Eduardo Baptista, "[China's New Five-Year Plan Calls for AI Throughout Its Economy, Tech Breakthroughs](#)," *Reuters*, March 4, 2026.
- 25 Leo Marx, "[Technology: The Emergence of a Hazardous Concept](#)," *Technology and Culture*, vol. 51 (no. 3), July 2010.
- 26 The White House, [Winning the Race: America's AI Action Plan](#), Executive Office of the President, July 2025.
- 27 The Canadian Press, "[New AI Minister Says Canada Won't 'Over-Index' on AI Regulation](#)," *CTV News*, June 10, 2025.
- 28 Vector Institute, "[New Study Reveals AI's \\$100B Economic Impact Across Canada, with Ontario Leading the Charge](#)," November 25, 2025.
- 29 Aditya Challapally, Chris Pease, Ramesh Raskar & Pradyumna Chari, [The GenAI Divide: State of AI in Business 2025](#), MIT NANDA, July 2025.
- 30 Jiang Li & Huju Liu, "[Artificial Intelligence Adoption and Productivity in Canadian Firms](#)," Statistics Canada, April 22, 2026.
- 31 Bobby Allyn, "[Here's Why Concerns About an AI Bubble Are Bigger Than Ever](#)," *NPR*, November 23, 2025.
- 32 Valerie Bryan, Shivani Sood & Chris Johnston, "[Analysis on Artificial Intelligence Use by Businesses in Canada, Second Quarter of 2025](#)," Statistics Canada, June 2025.
- 33 See, for example: Canadian Union of Public Employees, [Understanding artificial intelligence: A guide for CUPE members](#), July 2024.
- 34 AI Now Institute, "[Biometric Surveillance Is Quietly Expanding: Bright-Line Rules Are Key](#)," April 11, 2023.
- 35 Emma Nelson, "[Redefining Productivity in the Age of Workplace Surveillance](#)," Human Rights Research Center, July 24, 2025.
- 36 International Trade Union Confederation, [Artificial Intelligence: What Are the Implications for Trade Unions?](#) October 2025.
- 37 Janine Berg & Hannah Johnston, [AI in Human Resource Management: The Limits of Empiricism](#), International Labour Organization, November 2025.
- 38 *Ibid.*
- 39 Matthias Oschinski & Ruhani Walia, [Harnessing Generative AI: Navigating the Transformative Impact on Canada's Labour Market](#), Institute for Research on Public Policy, May 2025.
- 40 *Ibid.*
- 41 Rossana Merola, et al., [Workers' exposure to AI: What indicators tell us—and what they don't](#), International Labour Organization, February 2026.

- 42 Sam Manning, Tomás Aguirre, Mark Muro & Shriya Methkuppally, [Measuring US workers' capacity to adapt to AI-driven job displacement](#), Brookings Institution, January 2026.
- 43 Tahsin Mehdi & Marc Frenette, "Canadian Employment Trends in the Era of Generative Artificial Intelligence: Early Evidence," *Statistics Canada Economic and Social Reports*, vol. 6 (no. 1), January 2026.
- 44 Harshit Agrawal, "Data Center Energy Consumption: How Much Energy Did/Do/Will They Eat?" Yale Clean Energy Forum, November 12, 2025.
- 45 Jordan Hulecki, et al., "Data Centres, AI, and Electrification—Legal and Corporate Approaches to Growing Power Demands in Canada," *Alberta Law Review*, vol. 63 (no. 2), November 2025.
- 46 Jo Nova, "New AI Data Centers Will Use the Same Electricity as 2 Million Homes," Iowa Climate Science Education, May 20, 2025.
- 47 Josh Saul, et al., "AI Data Centers Are Sending Power Bills Soaring," *Bloomberg Technology*, September 29, 2025.
- 48 David Coletto, "Canadians Split on AI Data Centres as Cost Concerns and Local Opposition Emerge," Abacus Data, March 24, 2026.
- 49 Institute for Environmental Research and Education, "How Can AI Help the Environment?" June 1, 2025.
- 50 Sibongiseni Hlabisa, "The Ecology of Artificial Intelligence: Energy, Water, Materials, and Land Limits of Digital Systems," *Carbon Neutral Systems*, vol. 1 (no. 19), December 2025.
- 51 Yoon Young Chung, Federico Darakdjian & Tom Leahy, "When AI Meets Water Scarcity: Data Centers in a Thirsty World," *MSCI*, December 9, 2025.
- 52 See, for example: Rory White & Natasha Bulowski, "Three Quarters of Data Centre Sites Planned in Alberta Are in High Water Stress Areas," *Canada's National Observer*, March 23, 2026.
- 53 Yosri A. Fahim, Ibrahim W. Hasani, Samer Kabba & Waleed Mahmoud Ragab, "Artificial Intelligence in Healthcare and Medicine: Clinical Applications, Therapeutic Advances, and Future Perspectives," *European Journal of Medical Research*, vol. 30 (no. 1), September 2025.
- 54 *Ibid.*
- 55 Dimitra Panteli, et al., "Artificial Intelligence in Public Health: Promises, Challenges, and an Agenda for Policy Makers and Public Health Institutions," *The Lancet Public Health*, vol. 10 (no. 5), May 2025.
- 56 Dolores R. Serrano, et al., "Artificial Intelligence (AI) Applications in Drug Discovery and Drug Delivery: Revolutionizing Personalized Medicine," *Pharmaceutics*, vol. 16 (no. 10), October 2024.
- 57 Wu Quanyang, et al., "Artificial Intelligence in Lung Cancer Screening: Detection, Classification, Prediction, and Prognosis," *Cancer Medicine*, vol. 13 (no. 7), April 2024.
- 58 Marc Zao-Sanders, "How People Are Really Using AI in 2026," *Harvard Business Review*, June 1, 2026.
- 59 Nathan Matias & Megan Price, "How Public Involvement Can Improve the Science of AI," *Proceedings of the National Academy of Sciences*, vol. 122 (no. 48), November 2025.
- 60 *Ibid.*
- 61 Deborah M. Li, Shruti Parikh & Ana Costa, "A Critical Look into Artificial Intelligence and Healthcare Disparities," *Frontiers in Artificial Intelligence*, 2025.
- 62 Leo Anthony Celi, et al., "Sources of Bias in Artificial Intelligence that Perpetuate Healthcare Disparities—A Global Review," *PLOS Digital Health*, March 2022.
- 63 Colin D'Mello & Isaac Callan, "AI Systems Used by Ontario Doctors Hallucinate, Auditor General Finds," *Global News*, May 12, 2026.
- 64 Nicholas B. Tiller, et al., "Generative artificial intelligence-driven chatbots and medical misinformation: an accuracy, referencing and readability audit," *BMJ Open*, vol 16 (no. 4), April 2026.

- 65 Shruti Karlekar, "[The Psychology of the Eliza Effect: Anthropomorphism, Social Presence, and Projection in Human-AI Interaction](#)," *International Journal for Multidisciplinary Research*, vol.7 (no.6), November 2025.
- 66 Katie Hyslop, "[AI Chatbots Are Coming to BC Classrooms](#)," *The Tyee*, May 28, 2026.
- 67 Burns, et al, *A New Direction for Students in an AI World*.
- 68 Kristin Rushowy, "[Five more Ontario school boards and two private schools file lawsuits against Snapchat, TikTok and Meta, alleging they're harming students](#)," *Toronto Star*, May 29, 2024.
- 69 Nataliya Kosmyna, et al., "[Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task](#)," *arXiv*, June 2025.
- 70 Rafael M. Batista & Thomas L. Griffiths, "[A Rational Analysis of the Effects of Sycophantic AI](#)," *arXiv*, February 2026.
- 71 *Ibid*.
- 72 Organisation for Economic Co-operation and Development, [OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education](#), OECD Publishing, January 2026.
- 73 Burns, et al, *A New Direction for Students in an AI World*.
- 74 Mary Burns, Rebecca Winthrop, Natasha Luther, Emma Venetis & Rida Karim, [A New Direction for Students in an AI World: Prosper, Prepare, Protect](#), Brookings Institution, January 2026.
- 75 Hope Nguyen, "[Sam Altman Sparks Backlash over Claims that Intelligence Will Be a Utility: 'A Future Where...People Buy It from Us on a Meter'](#)," *Yahoo Tech*, March 21, 2026.
- 76 Cultural Policy Hub at OCAD University, [Generative AI: Policy Pathways for the Arts and Cultural Industries](#), OCAD University, June 2024.
- 77 *Ibid*.
- 78 Monique Jilesen, Sana Halwani, Jim Lepore & Devon R. Kapoor, [Canadian Courts Have Jurisdiction to Address Misappropriation by AI Companies](#), Lenczner Slaght, December 2025.
- 79 Statistics Canada, "[Which Canadian Businesses Are Using Generative Artificial Intelligence and Why?](#)" March 18, 2024.
- 80 Cultural Policy Hub at OCAD University, *Generative AI*.
- 81 Marina Adami, "[AI-Generated Slop Is Quietly Conquering the Internet. Is It a Threat to Journalism or a Problem that Will Fix Itself?](#)" *Reuters Institute for the Study of Journalism*, November 26, 2024.
- 82 Frank Landymore, "[Over 50 Percent of the Internet Is Now AI Slop, New Data Finds](#)," *Futurism*, October 14, 2025.
- 83 Organisation for Economic Co-operation and Development, [Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions](#), OECD Publishing, September 2025.
- 84 Trevor Potts, "[How Is the Federal Government Using AI, and What Does It Mean for Science and Democracy?](#)" *Evidence for Democracy*, March 24, 2026.
- 85 Woodrow Hartzog & Jessica M. Silbey, "[How AI Destroys Institutions](#)," *UC Law Journal*, vol. 77, last updated June 2026.
- 86 Christopher Nardi, "[The CRA spent \\$18 million on Charlie, its new tax information chatbot](#)," *National Post*, December 12, 2025.
- 87 Natasha Tusikov & Blayne Haggart, "[Carney's AI push risks harm as Ottawa automates public services](#)," *Policy Options*, April 30, 2026.
- 88 See, for example: Tim Roberts & Associates Consulting, [Legal aid service delivery in rural and remote communities across Canada: issues and perspectives in the context of COVID-19](#), Department of Justice Canada, January 2023; see also: National Advisory Council on Poverty, [We can do better: it is not a safety net if the holes are this big—The 2025 Report of the National Advisory Council on Poverty](#), Employment and Social Development Canada, October 2025.

- 89 Hartzog & Silbey, "How AI Destroys Institutions."
- 90 Potts, "How Is the Federal Government Using AI?"
- 91 Treasury Board of Canada, [AI Strategy for the Federal Public Service 2025-2027](#), Government of Canada, February 2025.
- 92 Blair Attard-Frost, "[Canada's new AI strategy is off to a bad start](#)," *BetaKit*, February 9, 2026.
- 93 Christina Colclough & Hannah Johnston, [Promoting digital sovereignty and fundamental rights: Six principles for inclusion in public service procurement contracts](#), European Federation of Public Service Unions, May 2026.
- 94 See, for example: Raluca Csernatosi, "[Can Democracy Survive the Disruptive Power of AI?](#)" Carnegie Endowment for International Peace, December 18, 2024.
- 95 Manuel Goyanes, Porismita Borah & Homero Gil de Zúñiga, "[Social media filtering and democracy: Effects of social media news use and uncivil political discussions on social media unfriending](#)," *Computers in Human Behavior*, vol 120, March 2021.
- 96 Ashton Black, "[AI and Democratic Equality: How Surveillance Capitalism and Computational Propaganda Threaten Democracy](#)," *Lecture Notes in Computer Science*, vol. 14129, October 2024.
- 97 Matthew Miller, [Deepfakes: Real Threat](#), KPMG, 2023.
- 98 Statistics Canada, "[Do Canadians have confidence in their public institutions](#)," November 23, 2023.
- 99 Leger, [Views on Artificial Intelligence: A survey of Canadians](#), August 2025.
- 100 Angus Reid Institute, [Canadians call for heavy AI regulation, but three-quarters doubt any government can keep up with the technology](#), June 2026.
- 101 Innovation, Science and Economic Development Canada, [Engagements on Canada's next AI Strategy: Summary of inputs](#), Government of Canada, February 2026.
- 102 Amnesty International, "[Nowhere is safe for us': Unlawful attacks and mass displacement in north-west Syria](#)," May 2020.
- 103 Kevin Cole, [Navigating Humanitarian AI: Lessons Learned from Building a Chatbot Proof-of-Concept](#), Refugee Solidarity Network, May 2024.
- 104 Amnesty International, [Unlawful by design: Exposing the human rights costs of generative AI](#), May 2026.
- 105 Privacy Commissioner of Canada, "[Statement by the Privacy Commissioner of Canada regarding a joint investigation of OpenAI's ChatGPT](#)," Office of the Privacy Commissioner of Canada, May 6, 2026.
- 106 Marcus Smith & Monique Mann, "[Facial Recognition Technology and Potential for Bias and Discrimination](#)," in *The Cambridge Handbook of Facial Recognition in the Modern State*, eds. Rita Matulionyte & Monika Zalnieriute, Cambridge University Press, March 2024.
- 107 Petra Molnar & Lex Gill, [Bots at the Gate: A Human Rights Analysis of Automated Decision Making in Canada's Immigration and Refugee System](#), Citizens Lab, University of Toronto, September 2018.
- 108 Ramón Reichert, "[Autonomous Occupation: Israel's AI-Driven Drone Warfare and the Digital Architecture of Authoritarian Power](#)," *Dialogues on Digital Society*, vol. 1 (no. 3), October 2025.
- 109 Maneesha Perera, et al., "[Indigenous Peoples and Artificial Intelligence: A Systematic Review and Future Directions](#)," *Big Data & Society*, vol. 12 (no. 2), June 2025.
- 110 Mayadhar Sethy, "[Artificial Intelligence and Epistemic Justice: A Decolonial Turn through Indigenous Knowledge Systems](#)," *AI & Society*, vol. 41, February 2026.
- 111 James Muldoon & Boxi A. Wu, "[Artificial Intelligence in the Colonial Matrix of Power](#)," *Philosophy & Technology*, vol. 36, December 2023.

- 112 Research Data Alliance International Indigenous Data Sovereignty Interest Group, "[CARE Principles for Indigenous Data Governance](#)," The Global Indigenous Data Alliance, September 2019.
- 113 Perera, et al., "Indigenous Peoples and Artificial Intelligence."
- 114 *Ibid.*
- 115 National Research Council Canada, "[Indigenous languages technology program](#)," Government of Canada, last modified June 2026.
- 116 Brian J. Chen, [The Big AI State: How the Trump Administration Is Shaping US Industrial Policy Toward "Global Technological Dominance"](#), Data & Society, January 2026.
- 117 Mishal Khan & Annette Bernhardt, [The Current Landscape of Tech and Work Policy in the U.S.: A Guide to Key Laws, Bills, and Concepts](#), UC Berkeley Labor Center, December 2025.
- 118 The White House, "[Ensuring a National Policy Framework for Artificial Intelligence](#)," December 11, 2025.
- 119 Data Center Watch, "[Briefing 03/27/2026](#)," March 27, 2026.
- 120 Sha Sajadieh, et al., [The AI Index 2026 Annual Report](#), Institute for Human-Centered AI (Stanford University), April 2026.
- 121 Sorelle Friedler, "[What to make of the Trump administration's AI Action Plan](#)," Brookings Institution, July 31, 2025.
- 122 Ho Ting Hung, "[Exploring China's cyber sovereignty concept and artificial intelligence governance model: a machine learning approach](#)," *Journal of Computational Social Science*, vol. 8, January 2025.
- 123 Huw Roberts, et al., "[The Chinese approach to artificial intelligence: an analysis of policy, ethics, and regulation](#)," *AI & Society*, vol. 36, June 2020.
- 124 Center for Security and Emerging Technology, "[Opinions of the State Council on Deepening the Implementation of the "Artificial Intelligence+" Initiative](#)," September 24, 2025.
- 125 White & Case, "[AI Watch: Global regulatory tracker—China](#)," September 22, 2025.
- 126 *Ibid.*
- 127 Scott Singer & Matt Sheehan, "[China's AI Policy at the Crossroads: Balancing Development and Control in the DeepSeek Era](#)," Carnegie Endowment for International Peace, July 17, 2025.
- 128 UNICEF China, "[UNICEF welcomes China's groundbreaking regulations to protect children from AI-related risks](#)," UNICEF, April 24, 2026.
- 129 *Ibid.*
- 130 Amy Hawkins, "[Chinese court awards compensation to sacked worker replaced by AI](#)," *The Guardian*, May 13, 2026.
- 131 Nature Editorial Board, "[China is leading the world on AI governance: other countries must engage](#)," *Nature*, vol. 648, December 2025.
- 132 Ben Wolford, "[What is GDPR, the EU's new data protection law?](#)" GDPR.EU, accessed June 17, 2026.
- 133 European Commission, "[Commission launches market investigations on cloud computing services under the Digital Markets Act](#)," November 17, 2025.
- 134 Directorate-General for Communications Networks, Content and Technology, "[Article 5: Prohibited AI practices](#)," European Commission, June 2024.
- 135 Directorate-General for Communications Networks, Content and Technology, "[Article 99: Penalties](#)," European Commission, June 2024.
- 136 EDRI, "[#ProtectNot Surveil: The EU AI Act Fails Migrants and People on the Move](#)," March 13, 2024.

- 137 Alexandra Minotakis, "[Regulating AI in the workplace: A critique of the EU AI Act and the Platform Work Directive through a worker-centred lens](#)," *Platforms & Society*, vol. 2, November 2025.
- 138 noyb.eu, [Digital Omnibus: First Analysis of Select GDPR and ePrivacy Proposals by the Commission](#), European Center for Digital Rights, February 2026.
- 139 Aida Ponce Del Castillo, "[The Digital Omnibus: Eroding Worker Protection and Rights](#)," Social Europe, February 4, 2026.
- 140 European Commission, "[European Approach to Artificial Intelligence](#)," June 4, 2026.
- 141 Jon Chun, Christian Schroeder de Witt & Katherine Elkins, "[Comparative Global AI Regulation: Policy Perspectives from the EU, China, and the US](#)," *arXiv*, October 2024.
- 142 *Ibid.*
- 143 Adam McDowell, "[Artificial Intelligence Was Made in Canada. How Can We Be World Leaders Once Again?](#)" *The Hub*, February 14, 2025.
- 144 Robert Hunt & Théo Lepage-Richer, "AI and the Canadian Institute for Advanced Research," in [Northern Lights and Silicon Dreams: AI Governance in Canada \(2011-2022\)](#), eds. Fenwick McKelvey, Sophie Toupin & Jonathan Roberge, Concordia University, 2024.
- 145 Innovation, Science and Economic Development Canada, [Pan-Canadian Artificial Intelligence Strategy](#), Government of Canada, March 2017; see also: ISED, *AI for All*.
- 146 See, for example: Eddie Sheppard & David Coletto, "[Optimism Meets Uncertainty and Even Fear in Canada's AI Landscape](#)," Abacus Data, July 25, 2025; Nicole Gillespie, et al., [Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025](#), University of Melbourne & KPMG, 2025.
- 147 Innovation, Science and Economic Development Canada, "[Mandate](#)," August 28, 2018.
- 148 Blair Attard-Frost, Ana Brandusescu & Kelly Lyons, "[The Governance of Artificial Intelligence in Canada: Findings and Opportunities from a Review of 84 AI Governance Initiatives](#)," *Government Information Quarterly*, vol 41 (no. 2), June 2024.
- 149 Andrew Clement, "[AIDA's 'Consultation Theatre' Highlights Flaws in a So-Called Agile Approach to AI Governance](#)," CIGI, November 6, 2023.
- 150 Ana Brandusescu & Renée Sieber, "[Missed Opportunities in AI Regulation: Lessons from Canada's AI and Data Act](#)," *Data & Policy*, vol 7 (no.40), January 2025.
- 151 ISED, *Engagements on Canada's next AI Strategy*.
- 152 *Ibid.*
- 153 Hadrian Mertins-Kirkwood & Rachel Pettigrew, "[Magnifica Technologia: Seven key takeaways from Canada's new AI strategy](#)," Canadian Centre for Policy Alternatives, June 5, 2026.
- 154 See, for example: Lawrence Zhang, "[From Sovereignty to Control: A Clear-Eyed View of Canadian Cloud Policy](#)," ITIF, April 27, 2026; Joshua Van Es, "[Canada's Digital Sovereignty Debate Overlooks Software Risks](#)," *Policy Options*, April 21, 2026.
- 155 Paul Kunert, "[Microsoft Admits It 'Cannot Guarantee' Data Sovereignty](#)," *The Register*, July 25, 2025.
- 156 Joshua Van Es, "[Canada's Digital Sovereignty Debate Overlooks Software Risks](#)."
- 157 Rukhsar Ali, "[Impact Assessments Not Required for Olds, Mihta Askiy Data Centres. Expert Says Legislation Needs to Catch Up](#)," *CBC News*, April 26, 2026.
- 158 Jody MacPherson, "[First Nation Demands Alberta Halt O'Leary's \\$70B Data Centre Project](#)," *The Energy Mix*, January 21, 2025.
- 159 See, for example: Vivian Hunt, et al., "[Diversity Matters Even More: The Case for Holistic Impact](#)," McKinsey & Company, December 5, 2023.

This report is available free of charge at www.policyalternatives.ca. The Canadian Centre for Policy Alternatives (CCPA) is an independent policy research organization. This report has been subjected to peer review and meets the research standards of the centre. The opinions in this report, and any errors, are those of the author(s) and do not necessarily reflect the views of the CCPA or funders of the report.

We acknowledge the Anishinaabe Algonquin People whose traditional unceded, unsundered territory is where this report was produced.

Acknowledgements

This report is the product of many months of engagement with researchers and activists across Canada. Thank you to everyone who participated in the CCPA's AI Roundtable in November 2025 and the Council of Canadians' Civil Society Summit on the AI Industry in May 2026, as well as the many people we met and learned from at various AI conferences and dialogues over the past year.

Thank you especially to Simon Enoch, Fenwick McKelvey, Sarah Ryan and Luke Stark for their invaluable comments on earlier drafts of this report. Any errors or omissions are the authors' alone.

About the authors

Hadrian Mertins-Kirkwood (he/him) is a senior researcher and political economist at the Canadian Centre for Policy Alternatives. His work focuses on federal economic, social and environmental policy, especially in the areas of climate change, artificial intelligence and industrial strategy.

Rachel Pettigrew (she/they) is a research assistant at the Canadian Centre for Policy Alternatives. Drawing on a degree in sociology and a Master of Social Work, their work combines policy research and advocacy to advance the interests of workers, communities and the environment amid political, economic and technological change.

AI transparency statement

This report was researched, written, edited and designed entirely by humans without the use of artificial intelligence (AI) tools.





CCPA
CANADIAN CENTRE
for POLICY ALTERNATIVES
CENTRE CANADIEN
de POLITIQUES ALTERNATIVES