

INDUSTRIAL REVOLUTIONARY

AI, PRODUCTIVITY, AND PROSPERITY

by

Peter Nicholson

June 23, 2025

This essay updates a paper of the same title published by the Public Policy Forum in December, 2024 (<https://ppforum.ca/publications/industrial-revolutionary-ai-productivity-prosperity/>)

EXECUTIVE SUMMARY

Industrial Revolutionary makes the case that today's unprecedented developments in artificial intelligence (AI) will drive a new era of global innovation and economic growth, comparable to the transformative shifts brought by the steam engine, electricity, and the microchip. Drawing on breakthrough use cases—from driverless Waymo taxis navigating busy city streets to chatbots that simplify and accelerate everyday tasks—the paper argues that AI is not just another technological advance; it is a General Purpose Technology (GPT) with the power to profoundly reshape economies and societies.

This essay provides an overview of the current state and near-term prospects of artificial intelligence and is intended to inform strategies that will enable Canada to participate fully in the AI industrial revolution that is now upon us.

At the heart of the paper is an urgent message: in advanced economies like Canada, productivity and GDP growth per person have been slowing for decades, weighing on living standards and public revenues. While recent innovations—like information technology applied to business processes in the 1990s—offered a temporary boost to growth, they were insufficient to overcome long-term structural hurdles: weak productivity in services, a plateau in educational attainment, aging populations, regulatory burdens, and diminishing returns from innovations that have powered growth for the past century. AI promises to restore robust productivity growth through two mechanisms: **augmentation**, where AI enhances human capabilities, and **substitution**, where AI replaces human effort entirely, freeing workers for higher-value roles.

Today's "generative" AI systems—powered by incredibly fast microprocessors, deep neural networks, and transformer architectures—can learn from vast amounts of text, image, or sensor data to perform tasks flexibly and creatively in *unstructured* environments. Still, these systems are not perfect. They hallucinate, inherit biases, lack transparency, and can fail at reasoning tasks that humans perform with ease. But their continuing improvement and sheer processing speed make today's AI platforms increasingly valuable in an ever widening range of activities. The key point is that AI doesn't have to be flawless; it only needs to outperform human alternatives at scale.

Several application areas already demonstrate the transformative potential:

- **Information processing:** AI-driven summarization, pattern recognition, and content generation across media.
- **Software development:** Faster code generation and improved quality, raising output across industries.
- **Manufacturing:** Smarter robots, proactive maintenance, and optimized supply chains.
- **Marketing:** Tailored content and audience targeting using generative tools.
- **Translation:** Real-time multilingual interactions, enhancing communication and collaboration.
- **Finance:** Smarter risk analytics, document processing, high-frequency trading, and personalized financial advice.
- **Education:** Customized learning tools for students and workers, improving skills acquisition.
- **Healthcare:** Automated administration, diagnostic support, treatment planning, and system-level analytics.

- **Government:** Better policy design, streamlined services, and eligibility processing.
- **Robotics:** Advanced autonomous vehicles, drones, and physical agents navigating real-world environments.
- **Science and innovation:** Accelerated discovery through analyzing data, generating hypotheses, and narrowing research pathways.

Economic modelling suggests that AI could add 0.3 to 2.9 percentage points to U.S. productivity growth annually, with mid-range Canadian studies projecting a 0.5–0.7 percentage point boost—enough to meaningfully improve economic trajectories. But given so much uncertainty regarding future AI innovation, a much larger impact on productivity growth is possible. On the other hand, prominent economists like MIT’s Daron Acemoglu and Bank of Canada Governor Macklem caution that the economic benefits of AI, while likely to be very significant, will take longer to unfold than the AI optimists project.

AI is not risk-free. No significant technology ever is. The risks of AI must be accepted and managed because the evolving methods of AI cannot be un-learned, nor can the competitive energies released by its power and prospects be bottled up by any nation or company. Public anxiety so far tends to center on three main concerns: job loss, stifled competition, and the safety/regulation of AI products.

Regarding the job threat, history shows that technological shifts never increase long-term unemployment. Rather, they change which jobs are done—e.g., as farming and manufacturing gave way to services, new occupations emerged. So while AI will displace certain roles and tasks—especially those that are information-heavy or routine—it will create many others. The policy challenge is to manage this transition through skills training, education, and social protection to ensure that productivity gains are broadly shared.

AI development requires massive capital—e.g., costs of training AI models amounting to hundreds of millions of dollars—giving a potentially decisive advantage to dominant players like Microsoft, Google, Amazon, and Meta. This threat of stifled competition is not a new challenge and can be met with measures such as antitrust enforcement, non-discriminatory cloud access, open-source models, and public investment in infrastructure, like Canada’s AI Compute Access Fund, in order to create a more diversified ecosystem.

Regulation of AI-powered products and services should build on existing frameworks for consumer safety, liability, bias prevention, and transparency. Techniques like independent audits, stress testing, regulatory sandboxes, and certified trials (e.g., in autonomous vehicles or medical AI) can generate public trust. In fact, providers of AI services will have a strong incentive to earn and maintain public trust which should make the regulatory task easier. Meanwhile, international initiatives such as the G7 Hiroshima Process and the Global Partnership on AI, signal a global push toward coordinated governance.

Canada is exceptionally well positioned to capitalize on the AI revolution. The country boasts world-class research institutions (Mila, Vector, Amii) and pioneers like Yoshua Bengio and Geoffrey Hinton. It also counts many successful AI firms—including Cohere and sectoral leaders in finance and logistics. Government support is substantial: the Pan-Canadian AI Strategy and budget allocations exceeding C\$2 billion for AI infrastructure reinforce policy commitment.

On the other hand, AI deployment remains uneven. Fewer than 15% of Canadian companies currently use or plan to use generative AI tools, highlighting a gap between capability and uptake. Three broad policy strategies are proposed to address this gap:

1. **Global and national regulatory clarity:** Canada should harmonize AI regulations with international norms, build data-sharing frameworks, and balance innovation with safety. Provincial and federal collaboration is essential.
2. **Sectoral focus on high-impact areas:** Government should support AI in key domains that have high productivity leverage—like energy systems, health care, logistics, and public sector transformation.
3. **Government as early adopter:** By embedding AI in public administration, Canada can send a strong signal, support domestic suppliers, and learn firsthand how best to regulate. Strategic procurement, internal pilot projects, and transparency will prove AI's value and reduce public skepticism.

Prime Minister Carney and his G7 colleagues, meeting in Alberta in June 2025, strongly signalled a new optimism regarding the promise of AI “to unlock competitiveness and deliver unprecedented prosperity...” So, rather than fearing job losses or worse, AI should instead be seen as a catalyst for innovation and renewed growth, and the risks and uncertainties as challenges to be embraced.

Will AI power a new era of productivity growth and material prosperity for Canada? Yes, it will. Only the precise trajectory, and especially the timing, remain to be discovered.

INDUSTRIAL REVOLUTIONARY—AI, PRODUCTIVITY, AND PROSPERITY

The future of driving has become a fact of daily life in Austin, Texas, my seasonal home. Waymo’s robotaxis, with passengers but no drivers, are ubiquitous—smoothly gliding around the busiest downtown streets; squeezing into crowded parking lots; and stopping obediently for careless pedestrians. The sight of a steering wheel with no one behind it turning left in front of you is frankly spooky but we’re already getting used to it. The same way we’ve gotten used to digital “assistants” popping up to answer our questions when we try to book a flight or cancel our cable.

Robo-taxis and chatbots are early harbingers of a revolutionary artificial intelligence that is destined to transform economies and cultures, and in fact just about everything. AI amplifies, and sometimes replaces, the capability of the human mind to process and act on the information conveyed by our senses. This is made possible by combining the mind-boggling power of the latest computer hardware with awesome feats of software engineering that have harnessed that power. The result is AIs that can “learn” to function in unstructured and uncertain environments, whether it’s navigating busy streets; identifying a worrisome abnormality in an X-ray; providing precise advance warning of extreme weather events; or conjuring up in seconds a draft multimedia funding pitch for your latest great idea. The list goes on and we’re just getting started.

While the precise trajectory of AI is impossible to foresee, it’s certain there’s no going back because:

- AI’s potential benefits are incredibly compelling.
- The world cannot un-learn AI, and global competition among businesses will continue to drive the technology.
- It’s a borderless technology that will become increasingly available to anyone, anywhere.
- The major players (US, China, EU) see their economic prosperity and military security at stake and thus find themselves in an AI race without a finish line.

The potential power of AI and the uncertainty surrounding the course of its development have understandably caused widespread anxiety. Particularly in Canada and other advanced economies, the public discussion has focused far more on the potential risks of AI than on the enormous benefits it holds in store. The objective of this paper is to explore those benefits while also suggesting how certain foreseeable risks can be overcome.

Fundamentally, the paper makes the positive case for AI as a “General Purpose Technology”—analogous in impact to previous world-changing technologies like the steam engine, electrification, and the microchip—that promises to reinvigorate economic growth through sustained impact on productivity globally and in Canada. The paper also serves as a scene-setter for more detailed expositions that identify specific opportunities and challenges and that lead to policy recommendations that will enable Canada to maximize the net benefit of the AI revolution.

We begin by explaining why productivity and economic growth in Canada and other highly developed economies have been gradually declining for decades and how the advent of powerful

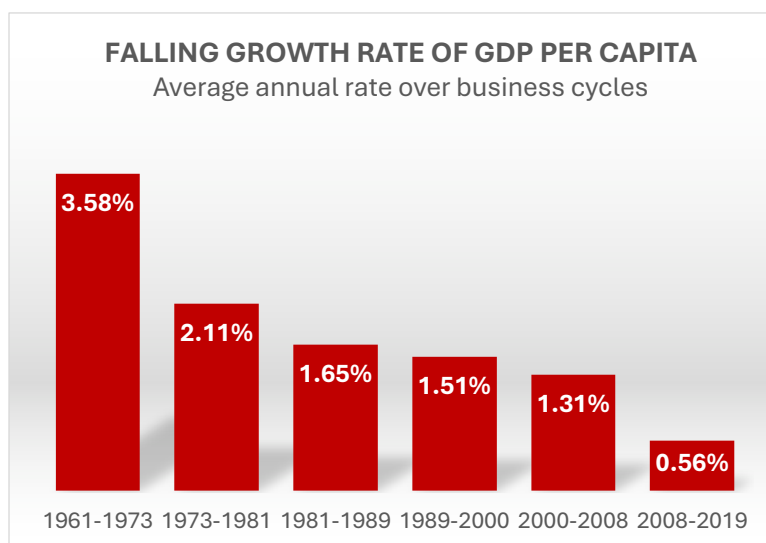
artificial intelligence will eventually reverse the trend. The discussion then turns to the need to manage three risks—AI’s impact on jobs, on competition, and on product regulation— that could delay or seriously mitigate AI’s benefit. The paper concludes with a high-level assessment of Canada’s readiness to participate in and benefit from the AI revolution. While it is outside the scope to make new policy recommendations, three theme areas are proposed to support an AI industrial strategy for Canada.

I. AI WILL POWER A NEW ERA OF ECONOMIC PROSPERITY

By now, it’s well known that the growth rate of Canada’s GDP per capita has been in a deep funk — the per person output of the economy is now barely greater than it was a decade ago.¹ Average living standards march in step with GDP per capita. Although this is not a perfect measure and does not adequately capture quality of life, GDP correlates positively, across nations and regions, with a great many social indicators including life expectancy, health status, and the incidence and consequences of poverty. Moreover, GDP defines the tax base and is therefore the ultimate source of funds for the social and other purposes of government. When the growth rate of per capita GDP declines, both average incomes and government resources are pinched. The sense of national opportunity can fade, and the public temper can sour.

What is less well known or appreciated is that this recent decline is not a new phenomenon. The rate of growth of Canada’s per capita GDP has been decreasing for many decades.² In fact, a [similar trend](#) is seen in virtually all the highly developed economies, the U.S. included. The robust economic growth rates of the post-war period, and the rising prosperity once taken for granted, now seem like a thing of the past.

A primary objective of this paper is to explain why this has occurred and to describe how the application of artificial intelligence (AI) throughout the global and Canadian economies promises to reverse the trend.



Just why has per capita economic growth been in long-run decline in virtually all the highly developed countries? The most important reason is that the rate of growth of productivity—the

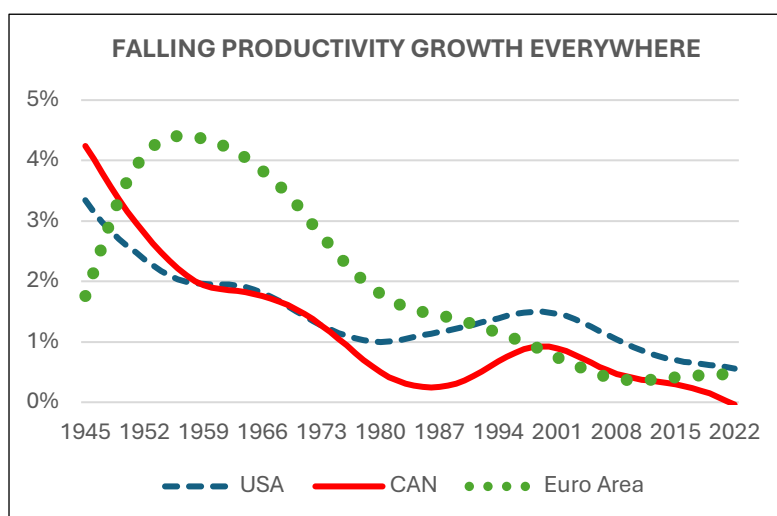
¹ Canada’s [per capita GDP](#) in 2023 was \$58,840 (in 2017 dollars), compared with \$56,610 in 2013, implying an annual average growth rate of merely 0.4%. In fact, GDP per capita was lower in 2023 than five years earlier in 2018.

² The chart is based on data from the [Centre for the Study of Living Standards](#) and the author’s calculations.

amount of GDP generated per hour of work, averaged across the economy—has been falling, amid the ups and downs of the business cycle, for most of the past 75 years.³ If short-run fluctuations are averaged out to focus on the trend illustrated in [the figure](#) below, productivity in Canada, the U.S., and western Europe grew at about four percent in the immediate aftermath of the Second World War. The steady decline thereafter was interrupted only by a decade of rising productivity growth in the U.S. and Canada in the 1990s driven by the application to business processes of computer and communications technology. Unfortunately, that mini-boom was not sustained and the rate of productivity growth throughout the West resumed its weakening trend. The pervasiveness of this phenomenon, both over time and across very different political systems, implies that the causes are deeply systemic. They have little to do with the political blame-shifting that dominates the news cycle and cannot be resolved by silver-bullet policy remedies.

Weak productivity has not been the only factor weighing on per capita GDP. The number of workers as a share of the population (the employment ratio) also matters. For the last several decades, the *growth* of that ratio slowed as women became more fully integrated into the paid workforce and the labour force participation rate plateaued.

Meanwhile, the populations in advanced economies have been growing older as both birth rates and death rates declined, which also causes the *growth* of the employment ratio to fall, intensifying the drag on the growth rate of per capita GDP.



Productivity growth is the most important factor going forward, and it has faced several structural headwinds:

- The increasing role in the economy of services—now constituting about 80 percent of GDP, and less amenable to the productivity-boosting effect of traditional automation and mass production than has been the case for goods production, particularly manufacturing.
- The declining growth impetus from human capital—as average educational attainment, at both the secondary and post-secondary levels, plateaued.

³ Per capita GDP is by definition “GDP/Worker X Workers/Population” or productivity multiplied by employment as a percent of the population (the “employment ratio”). The annual *growth rate* of GDP per capita is equal to the growth rate of productivity plus the growth rate of the employment ratio. As the population ages, the latter ratio tends to decline, and per capita GDP growth comes to depend entirely on productivity growth. That’s where we are today.

- The cumulative effect of regulation. While often justified to mitigate market failures and to achieve social and environmental objectives, a dense web of regulations has increasingly constrained decisions directed solely to growth maximization.
- Diminishing returns from the group of technologies that have powered productivity growth for more than a century—e.g., the electric motor, factory automation, industrial chemistry, the internal combustion engine, telecommunications.⁴

Take cars, or passenger aircraft, or even the internet. Each steadily got better and became more widely adopted for many years until technical improvement plateaued and their adoption rate saturated. This “S-curve”—a slow beginning followed by a period of rapid improvement before eventually tapering off—is characteristic of every technological innovation and is mirrored by a second S-curve of diffusion from the early adopters, to the majority in the middle, and finally to the stubborn holdouts. This process of innovation and diffusion is the fundamental engine of productivity and economic growth.

Simple arithmetic explains why the productivity growth impetus of every specific innovation eventually declines. It’s because sustaining an exponential growth rate—i.e. a constant percentage increase each year—becomes more and more challenging. The arithmetic of compounding means that the required absolute annual increment of growth keeps increasing while the returns on any specific technological innovation inevitably diminish. Fresh innovation is the only way to increase, or even sustain, a constant growth rate.⁵

Economic history demonstrates the extraordinary role in growth played by certain technologies that have pervasive effects on a very broad range of activities—e.g., the steam engine, electricity, the microchip, among others. Dubbed [General Purpose Technologies](#) (GPTs) they powered successive Industrial Revolutions by underpinning dense webs of complementary innovations in transportation, novel materials, construction, factory automation, retailing, and on and on. Every GPT is an innovation amplifier that stimulates a great deal more innovation and thus boosts the rate of growth of productivity. Collectively, these GPTs have created the modern world.

For a time, it was believed that applications of the microchip would power a plethora of innovations that would reignite productivity growth. Indeed, by the mid-1990s the production of information technology and communications (ITC) goods and the application of computers to business processes caused a productivity boom, but it [lasted](#) only through the early 2000s. The above noted structural headwinds proved to be more powerful than the ITC-driven productivity surge. Stronger medicine is needed.

⁴ Stanford economist Charles Jones and coauthors [have shown](#) that new ideas have been getting harder to find. Research efficiency—defined as productivity growth per researcher—has declined steadily even as R&D effort has increased rapidly. The economic historian, Robert Gordon, has also [argued](#) persuasively that the innovations associated with information technology have, so far, failed to drive productivity increases comparable to those of the past hundred years.

⁵ Imagine that a quantity like productivity (output per hour), starting at 100, is growing at 3%. The first year it grows by 3 units to 103. By the 25th year (1.03^{25}) it has grown to 209 units. To continue to grow 3% in the 26th year requires a new increment of 6.3 units—more than double the growth increment 25 years earlier. Eventually new and better ways need to be found to maintain a steady annual growth rate.

On the cusp of an Innovation Revolution

Artificial intelligence is the only technology, existing or on the horizon, with the potential to reverse the decades-long decline in the rate of productivity growth. The potentially unprecedented power of AI derives from its open-ended capacity to amplify and augment the human mind. This sets it utterly apart from any of the transformative technologies of the past (see [table](#) below). Fundamentally, AI increases labour productivity by:

- Augmenting human capabilities, thus enabling individuals to create value more quickly, and/or
- Substituting for humans in various tasks, thus creating value without requiring worker hours directly while freeing workers to create value in other ways.

The first process (“augmentation”) typically involves an AI taking on some of the tasks involved in an existing job—often the more routine ones or those that can take advantage of the AI’s information processing speed (see examples below). In such cases the AI functions as a subordinate [co-worker](#).

In other cases (“substitution”), the AI may be sufficiently sophisticated to replace an existing job entirely—e.g., as a routine customer service respondent—with the worker then freed up to perform some other job. In this case total output is increased—i.e. the AI replaces at least the worker’s original output while the replaced worker, if re-employed, produces new additional output—but with usually no net increase in human hours worked. Thus, labour productivity overall increases.⁶ Whether the AI functions as a co-worker or as a substitute will obviously depend on the nature of the job/task and on the sophistication of the AI itself. The distinction is familiar from the history of technology and automation. AI is simply the latest manifestation, but now with unprecedented capacities.

What is revolutionary this time is the ability of trained AIs to observe, process, and analyze digitally encoded information in all modalities—text, image, sound and even, via sensors, physical touch (Box I). Moreover, they can do this in unstructured environments and then generate appropriate responses via familiar interfaces with humans. Previous automatons—e.g., the assembly robots in car factories—were programmed to function in highly constrained and predictable environments, much like the early

Comparing AI to selected previous General Purpose Technologies

	Steam engine and electricity	Computers and internet	Artificial Intelligence
<i>Main output</i>	Energy	Calculations and information exchange	Advanced analytics (predictions, optimisation) and content generation
<i>Nature of tasks primarily affected</i>	Physical	Cognitive routine and communication	Broad range of cognitive
<i>Autonomy? (operate independently from humans)</i>	No	Limited	Potentially advanced
<i>Capacity for self-improvement?</i>	No	No	Yes
<i>A method of invention?</i>	No	Yes	Yes

⁶ AI is not restricted to augmenting or substituting for existing human work. It will also increasingly create entirely new sources of value that are wholly beyond human capability; familiar examples being GPS navigation and web search engines.

AIs that played a good game of chess but never had the “creativity” to beat the world’s top players. Today’s “[generative AIs](#)” are able through digital training processes to develop an internal model of some significant aspect of the real-world environment, such as human language, that’s rich enough to infer elements of that environment that were not explicitly present in the training data. That allows such AIs to cope flexibly with novel stimuli and to generate appropriate responses, a capability that is entirely unprecedented.

That said, today’s most advanced AIs are still far from infallible: they sometimes confidently assert things that are not true (“hallucinate”); they express biases implicit in their training data; they fail at certain tasks that humans can accomplish without even thinking. But they keep [improving](#) with better internal software, more powerful hardware, and better training data. Fundamental limits may exist, but they’re not yet in sight. Moreover, an AI does not have to be perfect: it only needs to perform reliably better than a human or other alternative in any given situation.

BOX I Peeking Under the Hood of AI

Machine intelligence of a kind has a long history of commercial application—in Jacquard’s punched card loom at the beginning of the 19th century, for example, and in the progressive automation of manufacturing ever since. What is so very different today is the phenomenal power of computer technology to enable [artificial simulation](#) of behaviour that is increasingly indistinguishable from its human counterpart.

For example, today’s graphic processing units (GPUs), the workhorse of Chat GPT and other generative AIs, [perform](#) an unimaginable 300 trillion basic arithmetic operations per second. To put this in some perspective, the GPU can do in one second what it would take a human—tapping in one number per second on a keyboard—ten million years! Such processing power is necessary but still not sufficient to perform the magic of today’s leading-edge AIs. First, data—whether text, image, sound, or virtually any other kind—must be digitally encoded into mathematical objects on which GPUs can operate. Then algorithmic procedures need to be designed that enable the AI to “learn” a model of a particular target domain like natural language (the training phase) after which it’s able to compute and regurgitate responses to queries relevant to that domain (the inference phase), whether it’s human language in the case of Large Language Models (LLMs) like GPT-4, or images in the case of “diffusion models” like DALL-E, and so forth.

Today’s leading-edge AIs employ a computational architecture called a “[deep neural network](#)”—based loosely on a highly simplified model of the brain—that processes digitally encoded input data via a series of computational parameters that number between 500 billion to more than one trillion in the latest models. By comparison, the human brain contains fewer than 100 billion neurons, although the interconnections among them are vastly more complex than those in contemporary artificial neural networks. On the other hand, the artificial networks process

information at incomparably greater speed. During the training process, the billions of parameters of the network are tuned, via a mathematical process of error minimization, until the network is able to form a good internal representation of the training data. Despite the phenomenal processing speed of GPUs, training the latest LLMs still requires 2-3 months and approximately 3 trillion-trillion (3×10^{24}) individual computations, a number that is vastly beyond human comprehension.

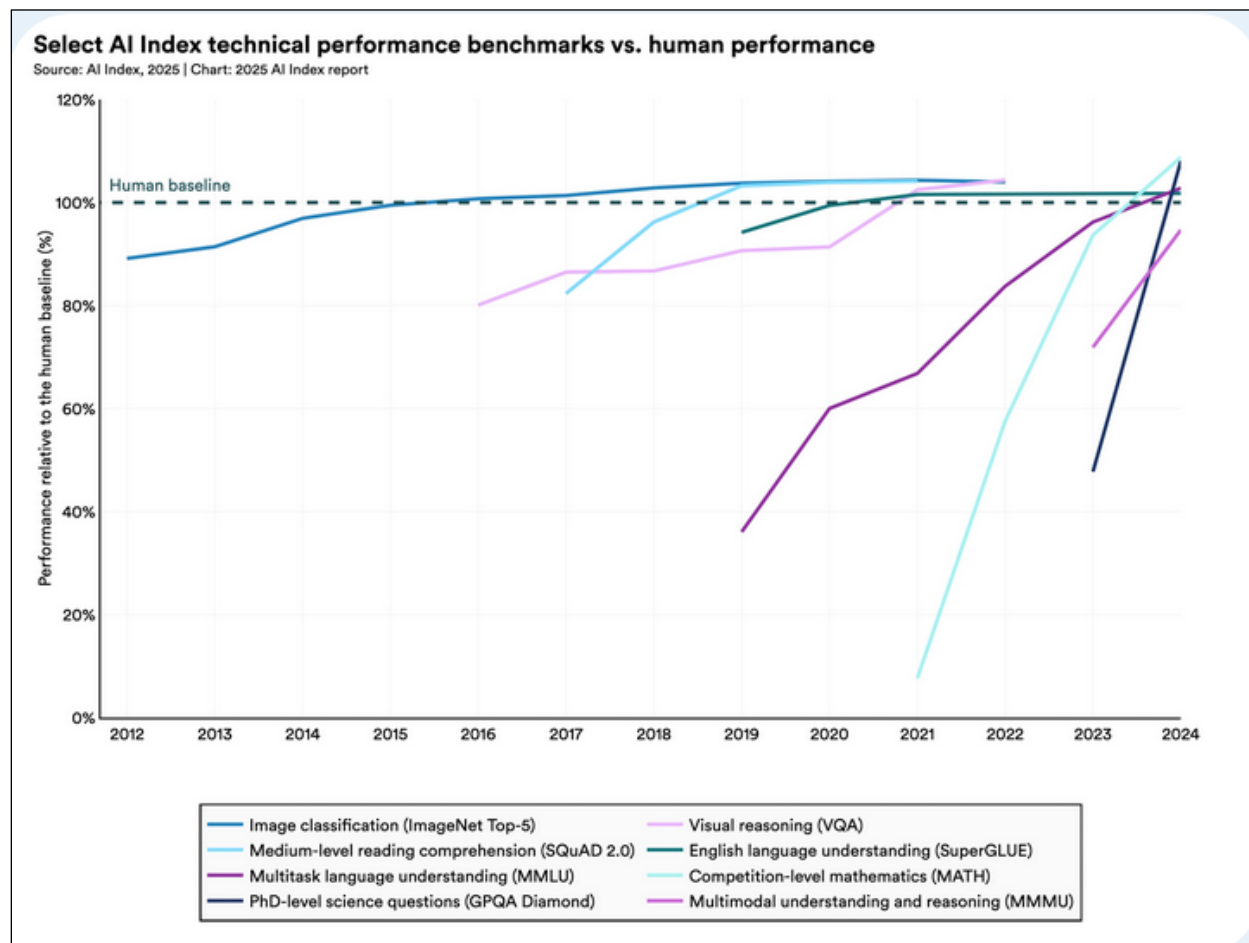
Such brute force is still not enough to “solve” a domain as complex as human language. That depended on an innovation, dubbed the transformer, developed by a team at Google in 2017. The transformer—the “T” in GPT (Generative Pre-trained Transformer)—is a computational architecture that enables LLMs and other leading AIs to recognize subtle contextual features of input data and to do so in a way that scales very efficiently to larger and larger models. As a result, AIs evolved from being expert at identifying and classifying objects of all kinds to being able to generate original content in text, image, and sound.

Chat GPT was thus born in late 2022 and attracted a million global users within five days. The public mind became fixated—both fascinated and alarmed by what seemed to be not just an artificial intelligence but actually an alien intelligence. Although computer scientists designed the software that animates today’s leading-edge AIs, and in that sense understand them, no one can see inside the “black box” to follow the countless billions of computational steps from input to output. Sheer complexity can cause surprising and creative behaviour to emerge, as anyone who interacts with a chatbot or image generator soon discovers. But the lack of transparency as to the step by step “reasoning” processes of a generative AI can stand in the way of trusting the output in respect of important decisions—e.g., in medical diagnosis, or conducting financial transactions. Lacking transparency as to process, trust can only be established through repeated testing under a broad range of real-world circumstances, as has occurred, for example, with self-driving vehicles (Box II).

Today, the ability of generative AIs to match or exceed human capability on an ever-expanding range of tasks has increased in rough proportion to the computational power being invested (see figure below). At present, the performance limits, if any, appear to be related to the availability of much greater volumes of high-quality training data that is specific to various application areas. Fundamental conceptual or technological roadblocks may eventually emerge to stymie progress. But to date AIs keep exceeding what was thought to be possible and there is no compelling evidence that such advances cannot continue, although not likely at quite the pace that many AI optimists project.

Significant electrical energy is required to train and operate today’s “foundation models” like the GPT series. Training such a model is estimated to require roughly 2,000 megawatt-hours, or enough to power about 180 Canadian homes for a year. The expansion of AI at very large scale will eventually require new electrical generation capacity but that will not be a show-stopper. Meanwhile, the much-publicized problems with the first generation of LLMs—the occasional

nonsense answers, biases in training data, and a relatively weak ability to cope with logical reasoning problems, and even simple arithmetic—are being overcome.



While AIs may now be super-human when it comes to passing college exams and conjuring up amazing feats of original text and image generation, the monumental private investment in their development—totaling an estimated US\$150 billion globally in 2024, of which about \$34 billion was allocated to generative AI—will not continue indefinitely without a commensurate financial return. That will depend on developing applications that can be trusted to outperform existing methods in terms both of cost and reliability. Although many AIs, using earlier generations of technology, already meet that test on relatively simple tasks, there are still few widely deployed use cases that exploit the full power of the latest generative models. Today, the really big money makers in the AI ecosystem are the infrastructure giants that supply the powerful computer chips and the “hyperscalers” like Amazon, Google, Microsoft, and Meta that provide the cloud resources on which everything rides. Now, having built it, the applications will come.

How AI will drive productivity growth

It's difficult to predict where AI will have its greatest impact, but several significant use cases are already being implemented or appear to be on the near horizon. With anticipated improvement and increasingly broad adoption, each has the potential to generate significant productivity growth.⁷ The examples include:

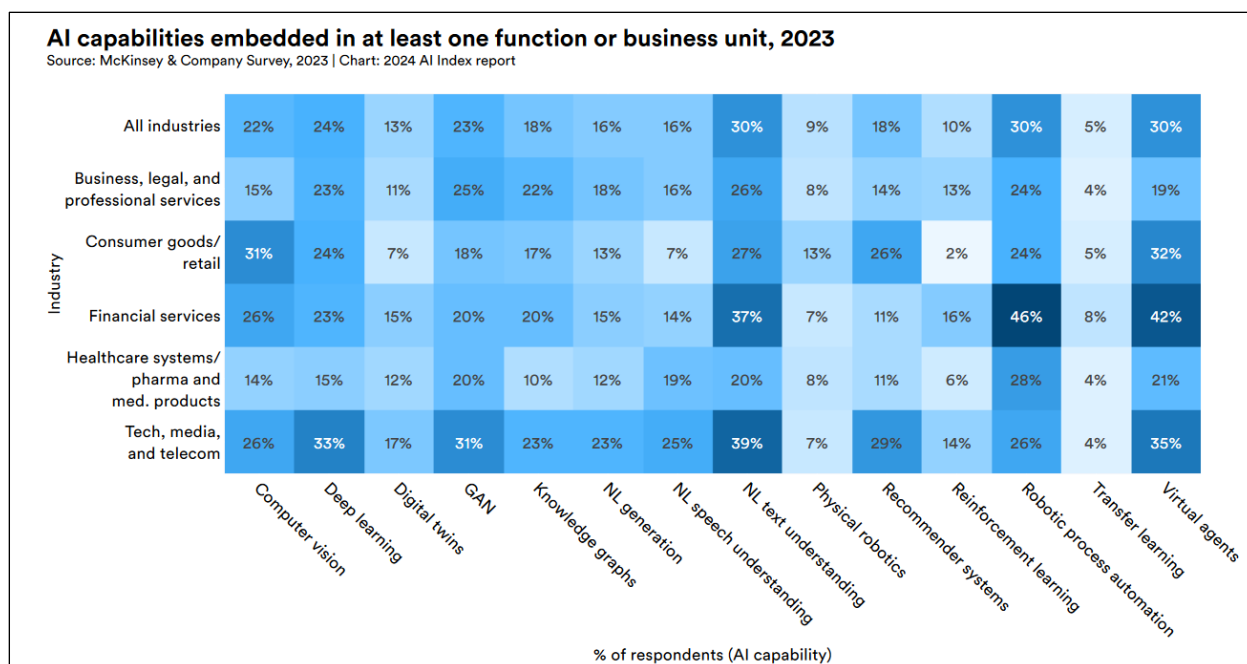
- Processing data and information of virtually any kind. This is a generic capability from which applications abound: summarizing enormous volumes of text; pattern recognition to enhance medical diagnosis; text, image, and sound creation in virtually any field including the creative arts. The AIs are typically co-workers here, amplifying the productivity and creativity of their human superiors.
- Augmenting the productivity of software engineers by significantly increasing output volume without sacrifice of quality. The observed [improvement](#) has typically been greater for weaker performers. Boosting software productivity and quality has the knock-on effect of improving productivity in virtually every other application area—an AI multiplier effect.
- Boosting [productivity in manufacturing](#) and in goods production generally—e.g., more flexible and intelligent robots; supply chain optimization; better fault detection leading to proactive maintenance and reduced downtime; generation of design options that optimize manufacturability, regulatory compliance, and cost.
- Improving [marketing strategies](#): e.g., preparation of promotional materials (text, image, video); micro-targeted customer identification and inducements. Such applications, while often controversial, underlie the business models of social media platforms and appear to be the most commercially valuable AI applications to date. Indirectly, they provide funding for the development of the leading-edge AI systems by, for example, Amazon, Microsoft, Alphabet, Meta.
- Enabling widespread, high-quality [language translation](#) implemented in real time with voice synthesis (including via a smart phone app), thereby enhancing productive collaboration and cross-cultural understanding in many everyday situations. Soon every tourist will be multilingual!
- Improved services and capabilities in the [financial sector](#): e.g., risk analysis; complex document processing associated with lending agreements and regulatory compliance; implementing high-frequency trading strategies; mass personalization of retail financial products.
- Enriching and personalizing [education](#) from the early years through adult learning and job training. The potential payoff in terms of productivity will be to build human capital far more efficiently than ever before, including re-training/up-skilling workers that AI augments or replaces.
- Boosting productivity across the [healthcare system](#): (i) automating much of the administrative burden using AI to process and provide management recommendations based on information in a variety of formats; (ii) assisting diagnosis and suggesting treatment

⁷ The productivity growth impact of any innovation depends both on its “unit impact” and on the pace and extent of uptake. While many AIs have remarkable capabilities, their effect on economic growth will remain limited until they are widely deployed in commercial applications.

options; (iii) providing deep analytical support for health system planning to optimize the allocation of scarce resources. Because healthcare is such a large and growing sector, and a domain where significant efficiency improvements have been hard to achieve, there are few if any areas with as much potential for AI-generated productivity growth. Of course, AI solutions will have to be thoroughly tested and proven in practice before wide acceptance and adoption.

- Improving the quality and efficiency of [government services](#). Because most government activity is information-intensive—e.g., processing payments/receipts that are subject to increasingly complex legal and policy criteria (like tax collection and EI payments); application of regulations; design of policy solutions—there are a great many tasks ideally suited for AI-based innovation. The potential for significant economy-wide productivity improvement is directly amplified by the sheer scale of modern government, and indirectly by the potential for more effective design and targeting of policy and regulation.
- Equipping various kinds of robots with the [capability to perform flexibly](#) in unstructured environments—e.g., self-drive vehicles and applications of drones already demonstrate AIs with agency in the physical world (Box II).
- [Amplifying innovation](#) itself. AI has literally a super-human capacity to: (i) absorb, process, summarize and interpret information of all kinds including from sensors in the physical environment and from the entirety of the world’s research literature; (ii) recognize subtle patterns in unstructured data, including insights that cross conventional disciplinary boundaries; (iii) perform increasingly complex logical inference (until recently this was a major weakness of LLMs); (iv) generate hypotheses based on all of the foregoing and engage in dialog with human researchers. These capacities to augment human research skills have the potential to radically increase the productivity of the discovery process itself. This will ultimately be the most transformative impact of AI since it promises to increase the rate of productivity growth of the economy as a whole.

The following table, based on data from McKinsey & Co., illustrates the uptake of various AI capabilities across a broad range of sectors. Although still early days, it’s already evident that AI has the salient characteristics of a GPT with the potential to accelerate productivity across much of the economy. In particular, AI promises to have a revolutionary impact on productivity in services, much as earlier generations of technology had in goods production, notably in agriculture and manufacturing.

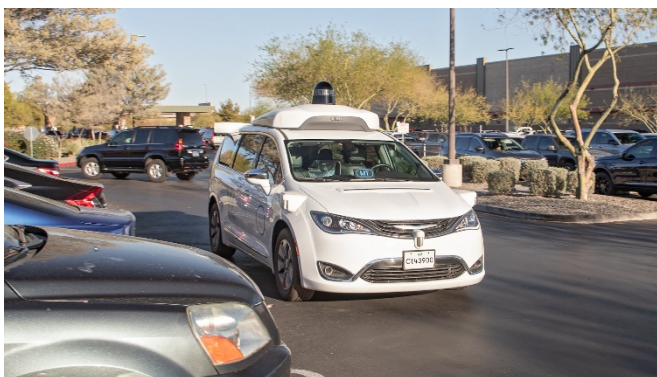


The hallmark of contemporary AI is the ability to function in relatively unstructured environments with the flexibility to respond to the unexpected. For an AI to be effective in specific application areas, in addition to generic capability it needs to be specifically tailored—e.g., trained on high-quality data relevant to the application area. This will increase the efficiency and especially the reliability essential to building the trust required for mission-critical applications and profitable business models. While artificial intelligence that is equal or superior to human intelligence in every respect—referred to as [artificial general intelligence](#) or AGI—may someday be [achieved](#) what seems most likely in the near to mid-term is the evolution of specialized AI “modules” optimized for specific domains.

BOX II Autonomous Vehicles—AIs as Agents in the Physical World

The Google spin-off company, Waymo, is [currently carrying](#) more than 250,000 paid riders a week in its fleet of wholly autonomous taxis—no driver at all—along the downtown streets of San Francisco, Phoenix, Austin, and Los Angeles with service in several more cities planned for 2025-26. Once seen as a threat to public safety, they’ve proven so safe and reliable that residents rarely give them a second glance. Waymo’s amazing achievement is the culmination of more than 35 million kilometers of on-road training dating from 2015. The driverless vehicles need to operate safely in an unpredictable urban roadway environment, drawing on sensor data that creates a picture of that environment as it evolves in real time. Waymo taxis have recently been [authorized to operate](#) on some California and Arizona freeways where speeds create an even higher bar for safety and rapid response. Yet the biggest challenge ahead is to turn Waymo’s miraculous technology into a viable business, a goal that is still on the horizon.

Waymo is simply the farthest along among a great [many start-ups](#) in the field. Some like Tesla and the UK company, Wayve, have focused on advanced driver-assistance technology. Others—notably the Canadian start-up [Waabi](#)—have targeted long-haul trucking. Meanwhile, at the leading edge of the technology, Waabi, Waymo, and a few others are seeking to combine the [transformer architecture](#) employed by LLMs with the standard sensors and AI software that already operate autonomous vehicles. The idea is to train the autonomous system to be even more adept at predicting the real-time behaviour of objects in the driving environment, much as LLMs are able to predict the most appropriate next word in a text response.

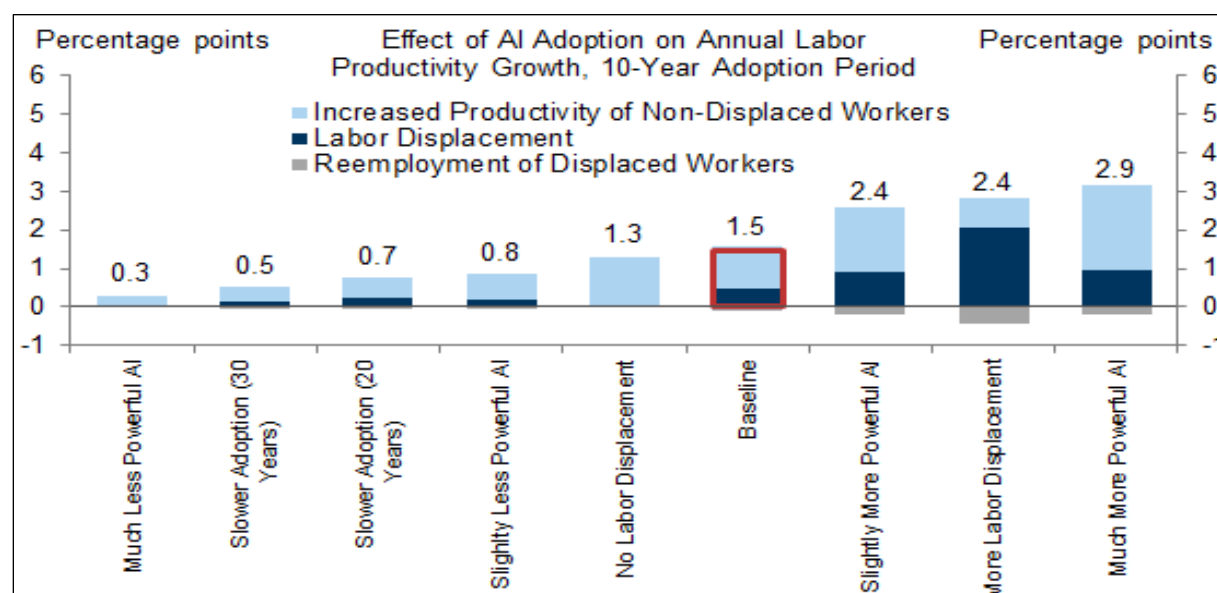


The autonomous vehicle—whether a taxi, a freight truck, or an aerial drone of the sort deployed in the Ukraine conflict—is essentially a new species of AI operating as an intelligent agent in the physical world. This represents a significant step beyond large language models. Why? It’s because apart from what is implicitly encoded in an LLM’s training data, those models have no understanding of how the physical world works and thus lack “common sense”. This compromises the ability of such models to resolve certain ambiguities that are intuitive to humans. For example, if we are told that “the gift could not fit in the suitcase because it was too big” we immediately understand that it is the gift, and not the suitcase, that is too big. An AI trained just on word patterns may not be able to make that inference. The self-drive vehicle, on the other hand, is out in the physical world, processing and acting on real-time information streamed from multiple sensors. By adding a transformer element, analogous to what has given the LLM its uncanny ability understand text, autonomous vehicles may eventually develop an equally uncanny ability to “understand” how the physical world works. In this way, self-drive cars, drones, and mobile robots of various kinds (such as those being developed by Vancouver-based [Sanctuary AI](#)) will learn about the tangible world from direct experience, much as human infants do. The implications for productivity, and much else, are profound.

Quantifying AI’s near-term productivity impact

There have been a number of attempts to [estimate](#) the increment of productivity and GDP growth that can be expected from AI applications that are either in place and/or reasonably foreseeable. Goldman-Sachs, in a widely cited [analysis](#) in 2023, projected that AI would cause the rate of growth of U.S. productivity to increase by an annual average of almost 1.5 percentage points over a 10-year period following widespread adoption of AI applications. That may seem a small number but it represents a more than doubling of the current trend rate, with much greater increases in sectors where AI applications show the earliest potential. The projection illustrated

below was subject to large uncertainty—ranging from a minimal increase of 0.3 percentage points to a transformative 2.9 points—depending on the difficulty of the tasks AI will be able to perform and the number of jobs that would be affected.



Source: Goldman-Sachs, 26 March, 2023

A recent analysis by TD Economics [estimated](#) that ramping up AI adoption could raise Canada's GDP by 5%-8% in 10 years relative to its baseline projection. This would translate to an AI-induced increase in the rate of productivity growth of 0.5-0.7 percentage points. While that is at the lower end of projections by Goldman-Sachs and most international projections, it's still significant by comparison with the dismal growth rate of Canada's productivity in recent years. And the possibility exists for a much larger productivity boost if the pace and breadth of AI adoption by Canadian businesses exceeds the assumptions in the TD projection.

Along with the optimists there are also some well-informed skeptics, notably the respected MIT economics professor, Daron Acemoglu, who in a June 2024 [commentary](#) concluded: "Given the focus and architecture of generative AI technology today, truly transformative changes won't happen quickly and few if any will likely occur within the next 10 years. The largest impacts of the technology in the coming years will most likely revolve around pure mental tasks, which are non-trivial in number and size, but not huge." Professor Acemoglu is nevertheless optimistic that generative AI "has the potential to fundamentally change the process of scientific discovery." His skepticism relates primarily to timing. Bank of Canada Governor Macklem expressed similar sentiments in a September 2024 [speech](#): "In the long run, we can expect AI to boost productivity... AI has all the hallmarks of a general-purpose technology, or GPT. But how large and how wide-ranging are hard to predict. We know from history that it takes years for a GPT to diffuse through the economy. We also know that the first applications are typically less transformative than the new businesses and new business models that eventually emerge. This all suggests that we won't see the full effects of this wave of AI anytime soon."

It's fair to say that anyone who hazards a forecast of the impact and timing of AI on productivity and GDP growth is peering into a dense fog. As with every major technological innovation, the early forecasts tend to extrapolate the recent past and miss entirely the most significant ultimate impacts—e.g., who could have predicted in 1900 that the automobile (dubbed the horseless carriage) would utterly transform the landscape, the economy, and the culture; or who foresaw in the mid-1980s, when the Internet was largely used for communication among small cadres of academics, that it would universalize access to information, spawn social media, and upend entire sectors of the economy; or that the plain old telephone would morph into a powerful computer in everybody's pocket.

That said, we can infer from the earlier summary list of examples how AI is likely to overcome the various “headwinds” that have caused the rate of productivity growth in the advanced economies to trend down since the early 1970s.

- AI clearly has the potential to transform productivity in most aspects of the service sector, just as earlier generations of machinery and factory automation did in agriculture and manufacturing. A services-focused economy will eventually no longer be a brake on robust productivity growth.
- Through its potentially transformative impact on education and training, AI promises to generate a new era of growth in average skills and competencies (human capital), possibly analogous in productivity impact to universal grade-school and widespread post-secondary education.
- The greatest productivity impetus of AI would come from its impact on the rate of innovation itself through augmentation of the human processes of discovery, understanding, and invention. Here the crystal ball is especially hazy, but the potential implications are most profound.

This vision is unabashedly optimistic. In purely technological terms it is plausible although the path forward is sure to be strewn with conceptual and engineering obstacles. Meanwhile, the reality of the AI transformation will be tempered by culture and politics and by the social habits and vested interests that constitute the status quo. Change is never easy. Change of the magnitude projected from the examples above will be disruptive, as has been the case with every major technological revolution. But opting out is not an option.

II. MANAGING THE RISKS AND BUILDING TRUST

Every major new technology carries society on a voyage into the unknown, buoyed by imagined benefits but always bringing risks. This inevitably creates a tension between innovation and restraint. Government finds itself on both sides seeking to maximize the opportunities while minimizing the risks. Artificial intelligence, more than any preceding technology, presents the greatest challenge of managing the opportunity-risk tension because of the scope and scale of its impact, the sheer pace of innovation, the largely opaque nature of the technology itself, and its inherently borderless characteristics.

There has been a great deal of discussion of the foreseeable as well as the potential risks of AI, from the relatively mundane to the [existential](#). The table shows there is widespread public

concern regarding a range of anticipated impacts of AI. Canadians appear to be among the most worried, ranking at or above a 21-country average on almost every impact area. This is to be expected given the widely cited concerns [expressed](#) by some prominent Canadian AI experts together with a

Global concerns about the impact of AI over the next few years (2023)						
	Global Ave	Canada	USA	China	India	Japan
Use for nefarious purpose	49%	52%	43%	39%	34%	52%
Impact on jobs	49%	49%	43%	39%	48%	36%
Violation of privacy	45%	48%	43%	39%	35%	33%
Dehumanization of services	41%	49%	39%	22%	28%	38%
Non-transparency re decisions	34%	38%	32%	27%	36%	24%
Impact on education	33%	34%	27%	26%	41%	25%
Ethical implications	30%	37%	32%	25%	32%	24%
Accuracy of results	28%	32%	30%	24%	39%	29%
Unequal access	26%	24%	19%	23%	29%	27%
Potential for bias	24%	30%	26%	20%	29%	18%
Own ability to use AI	22%	18%	19%	18%	36%	16%
Source: 2024 AI Index Report						

natural wariness of the unknown, particularly given experience with the downsides of social media and rampant disinformation on-line. Moreover, most of the potentially positive and compelling applications of AI have yet to appear.

The focus of this paper is on those positive applications and specifically on the economic significance of AI. But unless public skepticism is credibly addressed by governments and by businesses, many beneficial applications of AI will be delayed or prevented. And so too would be the positive impact on Canada's productivity and living standards. That's why the positive case for AI needs to be complemented by compelling evidence that the understandable public concerns can be effectively managed.⁸

While it is beyond the scope of this discussion to address all aspects of the public anxiety regarding AI, there are three issues directly associated with the economic implications that have seized the attention of policy makers everywhere. They are (i) the impact of AI on employment, (ii) assuring healthy competition in the AI marketplace, and (iii) regulating the role of AI in provision of goods and services. Experience from past technological revolutions suggests that each can be managed so that AI's transformative benefits are achieved while preserving human values and purpose.

Managing the impact on employment

As described earlier, AI (or any other productive technology) increases labour productivity by automating certain tasks, and/or by augmenting the capabilities of a human worker. In either case more output is produced per human hour worked. The extra value created shows up as increased compensation for workers and/or owners of AI capital. This increases overall demand in the

⁸ A recent [study](#) for the Canadian Chamber of Commerce underscored the point noting that "The factor of trust will be important for future adoption, with public interest and acceptance of AI likely being positively correlated with countries' business adoption rates."

economy, a portion of which generates new work for humans, including potentially for those initially replaced by AI.

The ultimate impacts on employment and on the shares of new income going to labour and to capital are complex and hard to foresee precisely. But the history of technological change demonstrates unequivocally that technology does not kill jobs overall. The long-term rate of unemployment is roughly constant. Technology simply changes what jobs get done. For example, in the first decade of the 20th century a million Canadians—about 35% of the entire workforce—were [employed in farming](#). In 1970 the number had fallen to 480,000 or 5.6% of total employment. By 2023 agricultural employment was down to 256,000 thousand, a mere 1.3% of Canadian jobs. But total farm output was many times greater than 50 or 100 years earlier. That’s the payoff from productivity growth. Meanwhile, the farm employment displaced by agricultural machinery and crop science was replaced by new jobs in rapidly expanding manufacturing and service industries.

By 1975 manufacturing accounted for about 20% of [Canadian employment](#) and services for 65%. Then, as technology enabled rapid productivity growth in manufacturing, that sector’s employment share fell from 20% to 9% currently, while employment shifted into an expanding range of services that now account for 80% of Canadian jobs. And within the broad ambit of services there continues to be dynamic birth and death of job categories—e.g., very few clerk-typists but lots of marketing managers.

People nevertheless focus on the particular job that is lost. It’s tangible and attached to a human face and to a community. The offsetting job that will eventually be created is an abstraction and may or may not be there for any particular individual. But new, unimagined jobs always do appear. A recent [study](#) in the US showed that 60% of employment in 2018 was in job titles that did not exist in 1940.

So, in terms of impact on total employment AI will be no different from past technological changes. Some jobs will disappear; new jobs will be created; and the resulting productivity growth will cause society’s material standard of living to increase. But the impact will vary according to sector of the economy, the specific job functions (“tasks”) that AI either automates or augments, and the characteristics of the impacted workforce—e.g., education, age, gender, income. A recent [study](#) published by the Canadian Chamber of Commerce shows which industries are most and least likely to be affected by AI in the near term (see below). Not surprisingly, the sectors most exposed to generative AI applications are those that mainly produce and use information, and the least exposed are those that engage heavily with the physical world or provide services that employ the human touch, like healthcare and social assistance.

Nevertheless, within every sector there are always specific tasks that can be automated or augmented by AI—e.g., administrative and decision-making support; document preparation; data analysis; or any task that requires subtle pattern recognition. In any event, AI is most likely to

augment higher skill jobs and automate those with lower skill requirements. That effect would be to increase income inequality. On the other hand, in the case of AI augmentation, the benefit appears to be [greatest for less experienced employees](#), presumably because the present generation of AI provides a smaller advantage to the most expert.

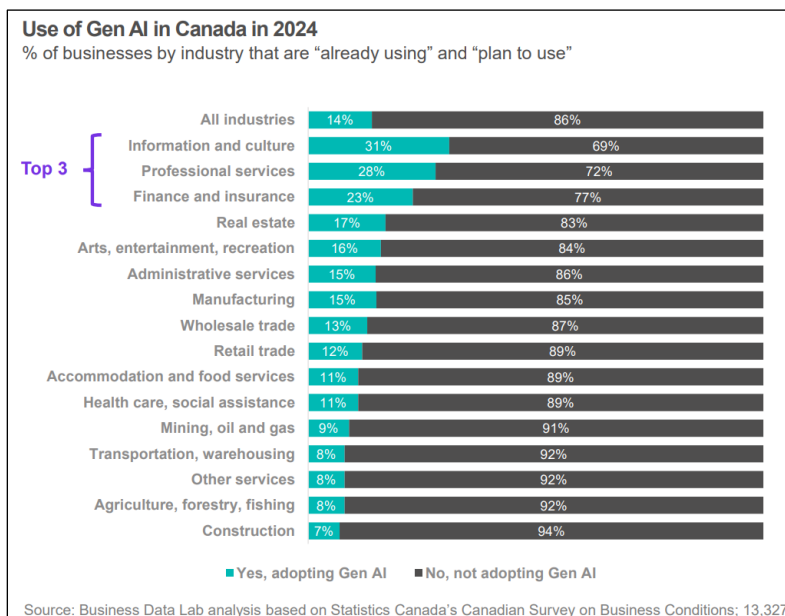
What can be said with confidence is that even as AI capability increases there will still be things for humans to do, and that the productivity growth enabled by AI will create a larger economic pie to be shared. The [policy issue](#) is therefore to manage the transition and to ensure that the new value generated by AI is shared fairly. These are not new challenges. In the context of past technological revolutions, they have been met with innovations in public policy—e.g., universal

education, progressive taxation, various worker protections, re-training, and the broad range of social programs that constitute the welfare state. Looking forward, we can build directly on that experience and innovate to address what will be unprecedented about the AI transition. The lesson is that it's society's *choice* as to how the employment impact of AI will be managed.

Ensuring healthy competition

Creating a state-of-the-art AI is hugely expensive due to the top-end processors (GPUs) and data center resources required—e.g., training Google's "Gemini Ultra" model in 2023 is [estimated](#) to have cost almost US\$200 million.⁹ A small number of transnational giants dominate the field—led by Microsoft, Google, Amazon and Meta—that possess the human talent, cloud infrastructure, access to massive proprietary data sets and customer channels, as well as the financial muscle to stay in the game. These companies can afford the "entry fee" after which provision of the resulting AI services is comparatively low-cost. A very significant barrier faces prospective competitors.

Vigorous competition is obviously beneficial for users and also promotes the innovation that drives an emerging field like AI. Moreover, excessive concentration of private influence over a strategically vital technology like AI can threaten the public interest, including national security.



⁹ Global private investment in AI in 2024 is estimated by the Stanford University based [AI Index 2025](#) to be more than US\$130 billion. The U.S. dominates with 80-85% of the total, far ahead of China's 7.2% and Canada's 2.3%.

Nevertheless, financial scale and private sector creativity are essential at this stage of the evolution of AI. So a delicate regulatory balance must be struck.

Fortunately, policies that promote fair and competitive markets have a long history that can be readily applied to the AI domain—e.g., review of merger and acquisition deals, and anti-trust regulation. Beyond that, recent experience in respect of the information and communications technology sectors will be of direct relevance in promoting a competitive environment—for example:

- Non-discriminatory access to cloud resources and other data infrastructure will be essential. Rules in an AI context could draw on experience in ensuring fair access to telecommunications networks.
- Choke points resulting from proprietary control of the most powerful AI models can be mitigated by support of open-source platforms like [Hugging Face](#) and [EleutherAI](#). There is a very large global community of researchers, developers, and financiers prepared to volunteer skills, time, and money to ensure a rich ecosystem of AI models analogous to those that developed open-source computer operating systems like Linux.
- Competitive AI development depends on affordable access to enormous computing power. Governments can contribute, as they have in the past, to this infrastructure—e.g., the federal government has [earmarked](#) \$2 billion over five years to launch a new AI Compute Access Fund and a Canadian AI Sovereign Compute Strategy.

These examples illustrate the rich body of already established regulation and practice that can promote a competitive AI ecosystem. What is nevertheless a unique challenge is to ensure fair access to the enormous volumes of data on which AI models are trained. Data is the fuel that powers generative AI and access to it has [emerged](#) as a contentious and potentially limiting constraint on future development. Policy creativity will be required to establish protocols for responsible data sharing that allow smaller companies to access high-quality datasets while maintaining user privacy, and to ensure that users and businesses have control over their data and can port it between platforms, enabling them to move from one AI service to another.

Regulating AI-based products

The sale and use of products—whether tangible goods or services—is already subject to extensive regulation regarding safety, liability, transparency, privacy, non-discrimination, consumer protection, among other things. How might this existing framework be adapted to the inclusion of AI? In many cases, it's relatively straightforward to apply existing regulations/standards. For example, product liability laws can be extended to AI-enhanced goods (e.g., medical devices), as can consumer protection regulations (e.g., to guard against deceptive practices), and anti-discrimination laws (e.g., to remove bias in AI-enhanced credit scoring or job screening). Professional licensing standards can be adapted to cover situations where AI is “included in the loop” through proof of competency—e.g., use of an [AI in radiology](#) should first have to conclusively demonstrate competency at least equal to the human standard.

While a great deal of product regulation can be simply “ported” into an AI environment with minimal adjustment, there are unique features of the technology that call for regulatory creativity. For example, generative AI is a dynamic, evolving learning mechanism that is in effect a “black box” with an opaque reasoning process that may inherit potential biases lurking in its training data or provide false answers. These characteristics create unique challenges in establishing the trustworthiness of AI products, particularly in important applications such as healthcare, legal contexts, autonomous driving, financial decision-making. It’s obviously in the provider’s interest to prove to users that its products are safe and perform as promised. That’s why a lot of technical effort is being made to minimize hallucinations, to improve the explainability of AI’s decisions, and generally to improve the amount and quality of training data for particular application domains. That motivation coming from the market can be amplified by application of existing regulations regarding liability and consumer protection as noted above. Crucially important in the latter regard is to require explicit notification when AI is a significant component of a product.¹⁰

Ultimately, the way to prove performance and establish trust in AI products is through disciplined demonstration of safety and efficacy. The required rigour would vary depending on the importance of the application. A lot of AI products should be regulated simply by market acceptance or rejection. But where more is at stake, rigorous testing protocols need to be developed and enforced. A number of methods have been applied or proposed including, for example: “red teaming” which involves stress tests on AI systems to identify vulnerabilities and weaknesses; independent audits of an AI system’s performance, safety, and compliance with ethical standards; “[regulatory sandboxes](#)” in which an AI product would be tested in a tightly-limited user environment and subject to light regulation; and controlled product testing (analogous to clinical trials for drug approval) as part of a certification process—e.g., the millions of kilometres driven before autonomous vehicles are certified.

Meanwhile, a great deal of work is underway at both the national and international levels to develop principles and codes of conduct to govern the development and use of advanced AI systems.¹¹ Given the inherently global nature of the AI phenomenon it’s obviously important to achieve as much international commonality as possible to minimize opportunities for regulatory arbitrage and inconsistent compliance obligations on transnational companies. Notable in this regard is the [Hiroshima AI Process](#) launched by the G7 in May 2023 to establish a framework for the trustworthy governance of AI systems. Although initiated by the G7, the framework invites international cooperation involving developing countries, private entities, and academic institutions. At the June 2025 G7 summit in Alberta leaders placed new emphasis on the economic potential of AI relative to the prevailing focus on regulation, [noting that](#): “We intend to double down on AI adoption efforts that connect research to practical applications, helping

¹⁰ A requirement for AI transparency is needed to minimize deceptive practices such as the use of “deep fakes”, AIs masquerading as humans, personal targeting based on facial recognition, etc. Effective control of these practices will depend on technological countermeasures combined with legal sanctions proportionate to the risk of harm.

¹¹ Significant initiatives include: OECD AI Principles; G7 Hiroshima AI Process; UNESCO Recommendation on the Ethics of AI; Global Partnership on AI (involving governments, industry and academia); ISO/IEC AI Standards; EU Artificial Intelligence Act; the U.S. AI Risk Management Framework (developed in 2023 by the National Institute of Standards and Technology).

businesses—especially SMEs—integrate AI technologies that drive productivity, growth and competitiveness.”

The national approaches to AI regulation in the U.S., EU, China, and Canada are summarized in Box III. Together with transnational initiatives to achieve harmonization, there is underway a concerted global regulatory effort to earn the trust of an often skeptical public. Inevitably, this will be an evolving learning process. But the history of managing past technological revolutions provides both the confidence that AI can be managed for human benefit and lessons as to how that can be accomplished.

BOX III National Approaches to AI Regulation

United States: The U.S. approach to AI policy under President Biden focused on promoting innovation while addressing potential risks through a combination of voluntary frameworks and sector-specific regulation. The White House 2023 [Executive Order](#) on AI was intended to coordinate significant federal AI efforts including mandatory safety testing for the most advanced systems. Donald Trump rescinded the Biden Order, replacing it with [his own](#) that shifts the emphasis decisively toward innovation, growth and competitiveness while loosening regulatory guardrails.

European Union: The EU has been at the forefront of AI regulation. The [AI Act](#) is the world’s most ambitious regulatory framework, classifying AI systems into risk categories with strict obligations on the highest risk systems. The Act also bans certain AI practices entirely—e.g., deploying subliminal techniques to materially distort opinions in a democratic society. The Act is expected to come into force in 2026 after a two-year grace period. But in common with the global shift in emphasis toward AI’s role as a growth engine, the [EU has indicated](#) some easing of certain regulations.

China: China has implemented several AI-related regulations—e.g., [Interim Measures](#) for the Administration of Generative AI Services that regulate AI-generated content and emphasize the need for safety and control over misinformation. But China has not yet passed a comprehensive law. Its approach is state-driven and employs strict regulations, particularly regarding the ethical use of AI, the protection of state interests, data privacy, and particularly recently, [technological self-reliance](#). China is committed to striking a balance between technological leadership and regulation to ensure that AI development remains consistent with the values and objectives of the Chinese Communist Party under President Xi.

Canada: The Trudeau government’s pending [Artificial Intelligence and Data Act](#) failed to pass before Parliament prorogued in January 2025 and the Carney government has yet to indicate when or if it would be re-introduced and, if so, in what form. The appointment of a first ever Minister of Artificial Intelligence and Data Innovation clearly signals the priority that Prime Minister Carney attaches to AI as a growth driver for the economy and the key way to enable more efficiently delivered public services. Meanwhile, several initiatives from the Trudeau government remain in place, at least for the time being: (a) the [Directive on](#)

[Automated Decision-Making](#) is in place for federal government AI use; (b) the Voluntary [Code of Conduct](#) on the Responsible Development and Management of Advanced Generative AI Systems, launched in September 2023, and signed by major tech companies; and (c) the 2017 [Pan-Canadian Artificial Intelligence Strategy](#), renewed in 2022. At the same time, several provincial governments have become [proactive](#) in the AI, data use, and privacy regulatory areas.

Common features: Across all four regions, there is now increased emphasis on the potential of AI to re-ignite stronger productivity growth, including applying AI to increase public sector efficiency. The shift in emphasis, relative to the prior overriding concern with the potential downsides of AI, was evident in the recent [G7 Leaders’ Statement on AI for Prosperity](#) (June 17, 2025).

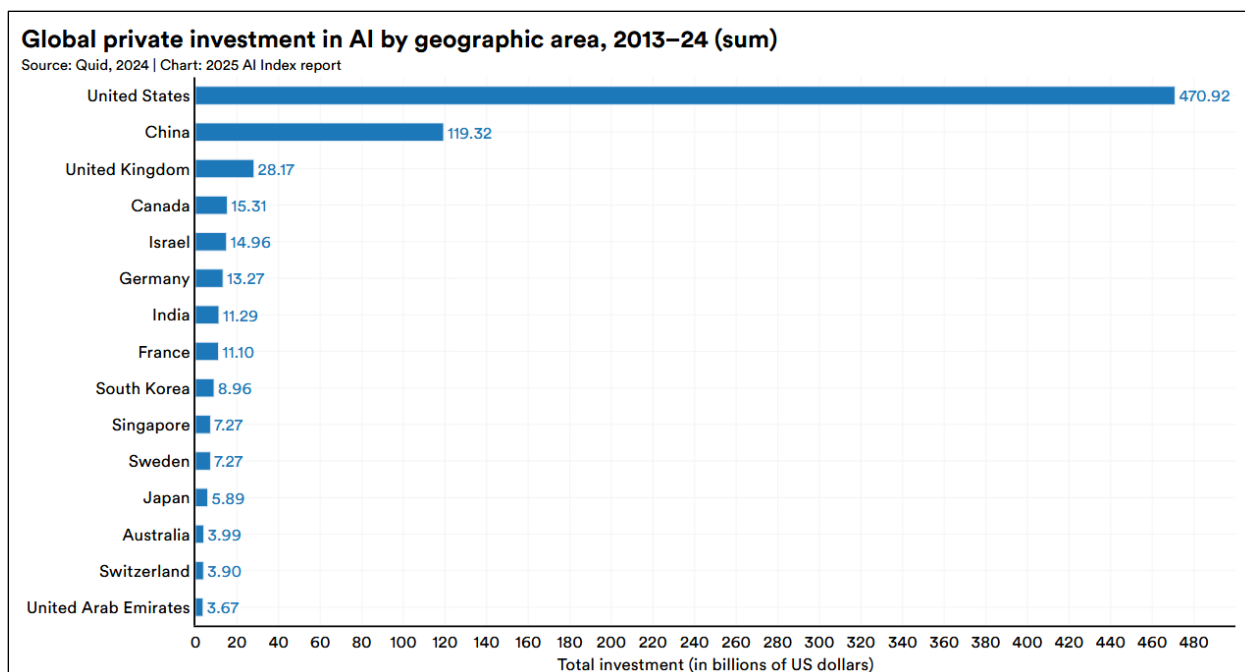
Divergent features: The EU and China are more proactive in setting comprehensive regulations, while the U.S. and Canada clearly favour more flexible, innovation-driven policies. Canada and the EU are actively involved in international forums like Global Partnership on AI (GPAI) and the Hiroshima AI Process, aiming for global consensus on AI governance, whereas the U.S. tends to emphasize international leadership through technological dominance. China is more insular in its AI development, prioritizing national objectives while engaging in some global AI dialogues.

III. PREPARING CANADA FOR THE AI ECONOMIC REVOLUTION

Is Canada ready, willing, and able to seize the AI opportunities outlined in this paper? The answer, at least in the author’s opinion, is a qualified “Yes”. We come to the AI revolution with several significant advantages:

- **Research excellence:** Canada is home to genuinely world-class AI research capability with (a) global leaders in the field like Yoshua Bengio (co-recipient of the Turing Award, the most prestigious recognition in computer science), Geoffrey Hinton (Turing Award co-recipient and Nobel Laureate), and Richard Sutton (internationally renowned for his path-breaking work on “reinforcement learning”), among many others; and (b) three top-ranked national AI research organizations —Mila in Montreal, Vector Institute in Toronto, and Amii in Edmonton, as well as vibrant regional hubs from coast to coast. This outstanding intellectual capital has branded Canada as a global leader while supporting a steady flow of superbly trained talent, thus making the country a compelling destination for investment. In the 10 years through 2023 Canada attracted almost US\$11 billion of private investment in AI, the 5th largest total in the world, although only 7 percent as much as the U.S., by far the global leader (Fig. 9).
- **A base of globally competitive AI companies:** Canada has already established a solid position in the still-emerging AI industry with pioneering companies like Kinaxis, Coveo, Element AI (since acquired by ServiceNow), BlueDot, Mindbridge, Benchsci, among many others. [Cohere](#)—co-founded in 2019 by CEO Aidan Gomez, a member of the Google team that developed the transformer architecture—is recognized as a world leader in the

integration of LLMs in various enterprise applications. Canada's major banks are among the global leading developers and users of AI in finance, with RBC currently ranked 3rd and all of the Big Five ranked in the global Top 25 according to the [Evident AI Index](#), the gold standard rating institution for AI in banking. Canada ranks fourth globally in respect of cumulative private sector investment in AI over the past 11 years (US\$15.3 billion) although investment amounts in the US and China respectively were 30 and eight times larger.



- **Supportive government:** The federal government, early on, made AI a focus of support—beginning with the 2017 Pan-Canadian AI Strategy, managed by the Canadian Foundation for Advanced Research ([CIFAR](#)), and so far funded with \$557 million.¹² In addition, the 2024 federal budget included \$2.4 billion for several AI support initiatives, headlined by \$2 billion for computing power needed to train and operate Canadian-based AI models. Regarding AI regulation and governance, the government has been active in various international forums, notably the G7 Hiroshima Process and the GPAI. Provincial governments have also been developing AI industrial strategies, most notably Quebec which has pledged a \$217 million [investment in the AI sector](#) (2022-27). Prime Minister Carney has signalled, with the appointment of Evan Solomon as Canada's first Minister of Artificial Intelligence and Data Innovation, that the federal government intends to accord even greater priority to AI.

Despite these impressive advantages, Canada continues to be challenged by the [long-standing difficulty](#) of converting its leading-edge knowledge into commercial innovation. This is the legacy of Canada's industrial structure—weighted toward resource extraction, construction,

¹² The initial \$125 million funding of the Strategy was managed by CIFAR and led to the creation of three national institutes (Mila, Vector, Amii). Further funding of \$442 million was announced in 2022 and will be allocated among CIFR (\$208 M), the national institutes (\$60 M to support commercialization), the Global Innovation Clusters (\$125 M to support AI development and use by SMEs), the Standards Council of Canada (\$8.6 M), and \$40 M for AI-dedicated computer resources.

finance, and U.S. branch plant investment—that reflects the country’s traditional position within a tightly-integrated North American economy. The AI revolution creates the opportunity to turn the page. But this will depend on the willingness of businesses, large and small and in virtually every sector of the economy, to step up to the opportunity.

So far, the response has been mixed. While Canada is well represented, relative to the size of its economy, by companies that are creating advanced AI services, uptake of applications has so far been limited.¹³ Polling of a representative sample of more than 13,000 businesses by Statistics Canada [revealed](#) that only 14 percent have used generative AI tools or are imminently planning to do so, while almost three-quarters are not even considering the option.¹⁴ They give a variety of reasons — e.g., no business case has been identified, they lack the skills, they’re concerned about data privacy, cost and financing. While it’s still early days, the technology is moving very fast and users that embark early on the learning curve are much more likely to end up with a durable competitive advantage. Moreover, the impact of AI on Canada’s rate of productivity growth depends almost entirely on the extent and speed of uptake of applications by businesses and public sector entities.

Three themes for an AI industrial strategy

It’s beyond the scope of this scene-setting paper to propose further specific policy measures to promote and accelerate the application of AI in Canada’s economy. That job will be the subject of future work by the Public Policy Forum. Following from the big picture analysis in this paper, three theme areas should be the key elements of an AI industrial strategy for Canada:

- **Regulatory development and harmonization:** A lack of regulatory certainty is well-known to discourage investment and AI will be no exception. Given the extremely dynamic nature of the field, AI regulation will inevitably be in flux, but every effort needs to be made to formulate and adhere to basic principles and to achieve consistency across jurisdictions. Canada acting alone has little influence on the course of AI governance and regulation, and therefore must continue to play a leading role in global forums like the Hiroshima AI Process and standards-setting bodies. Domestically, it is essential to promote AI regulatory harmonization in areas of shared jurisdiction and among provinces in their areas of exclusive jurisdiction, such as delivery of healthcare, grade-school education, energy and resource development. Rules in respect of data and privacy are central to development and use of AI models and thus require priority attention. Because AI creates many unprecedented opportunities and issues, domestic collaboration and harmonization may be easier to achieve than has been the case for established regulatory domains.

¹³ For example, Cohere [reports](#) that only 1 to 2 percent of its customers are in Canada.

¹⁴ According to Patrick Gill, the lead author of an [analysis](#) of the uptake of generative AI by Canadian business “Gen AI is a generational opportunity to boost Canadian productivity at a time when our performance is steadily headed in the wrong direction. Canadian businesses must innovate or die, and that means embracing Gen AI. While adoption has begun in every industry, it’s likely not fast enough for Canada to be competitive on the global stage.”

- **Supporting high pay-off sectors:** An AI industrial “strategy” requires that choices be made. While this is tough to do in a geographically and culturally vast federation like Canada, our resources are limited and impact depends on their focus. For example, the world is undergoing a multi-decade transformation of the energy system to renewably-generated electricity. AI can play a [major role](#) in this historic transition, but market forces alone in the heavily regulated electricity sector may not be sufficient to seize the opportunity in a timely way. Well-designed policy and program interventions can tip the balance. Beyond that, there is a limited number of areas that have particular potential to boost productivity through application of AI due either to their scale and importance, or their strategic position in the economy—e.g., healthcare systems and supply-chain logistics. Government support should be directed preferentially to such high-impact areas.
- **Government as a model user of AI:** The ancient proverb—“physician, heal thyself”—applies. There are at least three ways by which the application of AI to the government’s own operations can make a major contribution to an AI-based industrial strategy:
 - AI applied in the administrative and decision-support functions of government promises eventually to improve the efficiency and quality of service. If well-delivered, this will increase public confidence in AI applications, without which AI’s potential to improve productivity will be greatly diminished.
 - A government committed to AI can be a lead customer for Canadian suppliers through strategic procurement. The government, as early buyer, can help suppliers ascend the learning curve and then provide the validation that promotes market expansion. But this approach cannot work without explicit acknowledgement from the top that it’s an element of industrial strategy and consequently requires extra time and cost.
 - Finally, a government’s committed use of AI will provide a wealth of practical experience from a user’s perspective to inform development of wise policy and regulation in this novel area.

While artificial intelligence is still in its infancy, the child is already precocious. With the sudden advent of generative AI, the revolutionary potential of AI for human betterment has become evident. Because AI amplifies, and in many ways mimics, the human mind it has no precedent in the history of technology. With great power also comes great risk. But the risk has to be accepted and managed because AI cannot be “unlearned”, nor can its development be terminated for the simple reason that the prospective benefits are so compelling and there’s no jurisdiction on earth that has the authority and power to call a halt.

Will AI power a new era of productivity growth and material prosperity in Canada? Yes, it will. Only the precise trajectory, and especially the timing, remain to be discovered.