

AI in Canada

Meeting the opportunity and
governing AI in Canada



The promise of AI



Chris Barry,
President,
Microsoft Canada

Foreword by Chris Barry

The era of AI is here, ushering in a transformative wave with potential to touch every facet of our lives and enhance our experiences in unprecedented ways. It is not just a technological advancement; it is a societal shift that is propelling us into a future where innovation takes centre stage.

Yet, amid this excitement, a critical mandate remains; as we delve into this AI-powered future, it is imperative that we innovate responsibly, and any progress aligns with human values. With this exciting technological revolution, the role of human oversight cannot be overstated. AI is a tool harnessed by human ingenuity. Our ability to guide, govern, and ensure responsible deployment of AI is paramount. Regulation is crucial to helping ensure that, while we explore AI's vast potential, we do so responsibly.

Canada stands at the forefront of championing responsible AI governance.

The recent unveiling of a Canadian code of conduct for advanced AI systems and contributions to the historic G7 Hiroshima Process position Canada as a global leader, highlighting Canada's commitment to responsible AI

deployment. Codes of conduct serve as invaluable tools, complementing the ongoing efforts of governments worldwide, as they craft legislative and regulatory frameworks for AI. Codes provide the cornerstone upon which a globally interoperable framework can be built among G7 countries and other allies. Concurrently, the Government of Canada's proposed Bill C-27, *Artificial Intelligence and Data Act* (AIDA) signifies a pivotal moment and embodies a crucial step in responsible regulation.

At Microsoft, we champion the need for robust legislation that navigates the complexities of AI, ensuring we safeguard privacy, civil liberties, and ethical considerations, while nurturing the potential of this revolutionary technology. And in the heart of this transformation, we remain steadfast in our commitment to empower groundbreaking initiatives that demonstrate our nation's leadership in AI.

And we have already started to see the breakthroughs transpire.

During the pandemic, Canada's healthcare sector mitigated the spread of infections, enhanced patient safety, and charted a path toward a future where AI solutions helped to revolutionize healthcare and improve patient outcomes. By harnessing the power of AI technologies, Canadian healthcare organizations, such as [Toronto's University Health Network \(UHN\)](#) and [BC Cancer](#), are also unravelling hidden insights within vast datasets, leading to breakthroughs in cancer treatments.

[The City of Kelowna](#) is developing AI tools to speed up the building permit process by 20-30%, enhancing the overall experience of citizens and government workers alike: citizens who receive frictionless service and government workers whose wellbeing improves from the result of lower churn from data overload and repetitive work.

In partnership with [Ontario Power Generation \(OPG\)](#), we are working to shape a net-zero and environmentally resilient future where AI solutions are instrumental in powering new energy insights, matching local carbon-free sources, reducing emissions, and shaping a sustainable energy landscape.

Across the country, [AltaML is partnering with the Government of Alberta and Alberta Wildfire](#) to build Microsoft AI powered solutions to predict the likelihood of fire in any region, helping to determine how best to deploy scarce resources resulting in \$2–\$5 million in annual savings and helping protect lives.

It makes me particularly proud that [Microsoft brought Inuktitut](#), the language of more than 70% of Nunavut’s population, to Microsoft’s AI-powered Translator. Our work to help preserve cultural heritage with the Government of Nunavut gives tangible expression to Canadian truth and reconciliation.

Ensuring that AI benefits everyone will take collaboration across a diverse set of stakeholders, including policy makers, academics, civil society, and industry to ensure many voices and community interests are considered. Government and industry have a shared responsibility to make certain that Canada’s workforce is equipped with the skills, knowledge, and opportunity to gain jobs and increase livelihoods in an AI economy.

Microsoft is committed to helping people and communities learn how to harness the power of AI through our new AI Skills Initiative, and to continuing to invest in local talent capacity, including by building initiatives such as the [Canadian Tech Talent Accelerator](#), launched in partnership Digital Supercluster and NPower Canada.

The paper that follows shares our suggestions for the strong regulatory guardrails we believe are needed to see AI shape more prosperous, inclusive, and sustainable communities while prioritizing the safety, security and trust of citizens, governments, and businesses across Canada.

As Microsoft Canada celebrates our part in Canada’s AI-transformative journey, we are also mindful of the road ahead. AI’s potential is vast, and with responsible regulation, it remains a force for good, enriching lives and empowering Canada.

The collaborative partnership between Canada and Microsoft exemplifies the fusion of innovation with ethics, shaping a future where AI serves humanity, enriches the lives of Canadians, and propels our great nation toward a future defined by boundless possibilities. Microsoft is deeply committed to this journey where Canada continues to lead, inspire, and thrive in the digital age.



Chris Barry
President, Microsoft Canada



AI
Explained

AI is changing how we work, live, and play. Microsoft AI provides billions of intelligent experiences every day to people using our business software and services, gaming on our Xbox platform, and in keeping organizations secure. Our AI tools and technologies are designed to benefit everyone at every level in every organization. They are used in workplaces, home offices, academic institutions, research labs, hospitals, manufacturing facilities, and in emergency preparedness around the world. They are helping everyone from scientists and salespeople, to farmers, software developers, front line healthcare workers, and security practitioners.

What exactly is AI?

Using math and logic, an artificially intelligent computer system simulates the reasoning that humans use to learn from data and make decisions.

Algorithms are at the core of every AI system. They are sets of instructions or rules that the system follows to process and analyze data. Different algorithms are used for different tasks, such as processing images or language, for example.

Using algorithms, an AI model is trained on a dataset representative of the task or problem it is being developed to solve. Data could be in the form of unstructured data, such as text, images, audio, or structured data that has been stored in a standardized format.

The AI model uses the patterns and relationships it discovers in the data to adjust and optimize its performance resulting in a learned representation of the training data, which can be used to interpret and make inferences from new input or data.

One of the techniques used to train AI is machine learning. It is akin to teaching a computer program to recognize patterns by showing it examples. Just as we learn from experience, machine learning models are tuned from data. They look for similarities or patterns in the data to perform tasks such as making predictions, recommendations, classifications, and more.

Imagine you want to train a computer program to recognize cats. You would feed a machine learning algorithm a large number of pictures, each labelled as to whether the image contains a cat or not. By analyzing these labelled examples, the algorithm can identify patterns such as shapes,

colours, and textures associated with cats, and create a model based on this learning. Once trained, the model can then be used to analyze new, unseen images to provide a prediction of whether they contain a cat or not. The more examples the algorithm sees, the better the model becomes at recognizing cats accurately.

A subset of machine learning under the broader umbrella of AI is reinforcement learning. It involves directing an AI system to make a sequence of decisions to maximize a reward. The AI system learns through trial and error, receiving feedback in the form of rewards—or penalties—based on its choices. By learning which actions lead to higher rewards, the AI system improves its decision-making over time.

AI development has also been inspired by the structure and functioning of the human brain. Neural networks are computational models that consist of interconnected nodes called artificial neurons, organized in layers. Each artificial neuron receives input, performs a computation, and passes the result to the next layer. Neural networks excel at pattern recognition and train their models to demonstrate complex relationships in data.

As you can see, AI is a wide umbrella under which sits a range of different but complementary computer science fields and sub-fields—many of which have been in research for decades. But whichever approach is used to develop an AI

system, having access to data and compute power to train your AI model are fundamental build blocks.

AI-powered possibility

AI enables machines to carry out a wide range of tasks, many of which were not possible before its advent. Natural Language Processing (NLP), for example, extracts meaning from text or speech data. It enables computers to process languages and perform tasks like sentiment analysis, language translation, and text generation. AI-powered Computer Vision similarly enables machines to ‘see’ and decode the visual world around us. It involves techniques such as object detection and image recognition and has a wide range of applications from medical imaging to autonomous vehicles.

These capabilities and more are transforming how we use, benefit from, and interact with computer technology. AI can serve as a ‘copilot’ or virtual assistant to augment our human ingenuity and creativity through automating tasks and providing new insight.


Over the last decade, the AI field has made considerable progress on image, sound, and language tasks. With recent advances in the form of generative AI—a class of AI models that can generate new content such as text, images, code, and more. Generative AI is underpinned by a class of large-scale models known as *foundation models*.

Foundation models are trained on massive amounts of data and can perform a wide range of tasks. With a simple prompt like “describe a scene of the sun rising over the beach,” generative AI models can output a detailed description or produce an image, which can then be animated or even turned into video.

Many recent generative AI models, such as language models, are not only good at generating text but also generating, explaining, and debugging code. [GitHub Copilot](#), for example, leverages OpenAI’s Codex model to assist developers in writing code.

Another example of these AI systems in action is the new AI-powered Bing search engine. The web search experience often involves the time-consuming task of reviewing and synthesizing information from a variety of sources identified from different search queries. Now, Bing can do the heavy lifting for you, working behind the scenes to make the necessary queries, collect results, synthesize the information, and present a single complete answer.

Time taken in months to reach 100 million monthly average users



Source: Eyerys



AI at work – Q&A with Kevin Scott

Kevin Scott is Executive Vice President of Technology & Research, and the Chief Technology Officer of Microsoft. Scott's 20-year career in technology spans both academia and industry as researcher, engineer, and leader.

How do you see AI benefiting us at work?

The entire knowledge economy will see a transformation in how AI helps with repetitive aspects of your work and makes it generally more pleasant and fulfilling. Not only will it help with tasks like writing, editing, or creating reports, it is going to apply to almost anything—designing new molecules to create medicine or making manufacturing “recipes” from 3D models.

This is the “copilot for everything” dream—that you would have a copilot that could sit alongside you as you are doing any kind of cognitive work, helping you not just get more done, but also enhancing your creativity in new and exciting ways.

Will AI help make work more rewarding?

In many cases, people now have new, interesting, and fundamentally more effective tools than they have had before. For some workers, it is enhancing that core flow that you get into when you are doing the work; it speeds you up. It is like having a better set of running shoes to go run a race or marathon.

This is exactly what we are seeing with the experiences that developers are having with

Copilot; they are reporting that Copilot helps them stay in the flow and keeps their minds sharper during what used to be boring and repetitive tasks. And when AI tools can help to eliminate drudgery from a job, something that is super repetitive or annoying or that was getting in their way of getting to the thing that they really enjoy, it unsurprisingly improves satisfaction.

Will AI transform the way we work?

We've gone from machine learning in a few places to being integrated into a wide range of different products and services, everything from how your Outlook email client works, your predictive text in Word, your Bing search experience, to what your feed looks like in Xbox Cloud Gaming and LinkedIn. There's AI all over the place making these products better.

As all these AI systems continue to grow and evolve, I think we can expect that these advances are going to fundamentally change the nature of work, in some places more than others, and in some cases create a whole spate of new jobs that didn't exist before. You can look back and see the same thing happen adjacent to all kinds of famous paradigm shifts in technology over history: the telephone, the automobile, the internet. And just like with those examples, we're going to need new ways to think about work and new ways to think about skills.

A brief history of AI: From Turing to ChatGPT

1940 Enigma Machine


Alan Turing pioneers machine learning to crack the Enigma code.

1955 The term 'artificial intelligence' is coined

John McCarthy uses 'artificial intelligence' to describe the concept of machines that can exhibit human-like intelligence.

*1991 Microsoft Research

Microsoft establishes its research division, dedicated to advancing technology and exploring innovative ideas, including AI.



1950 Turing Test

Alan Turing proposes a natural language conversation test for machine intelligence.

1964 ELIZA – the first chatbot

Joseph Weizenbaum releases ELIZA, the first program of its kind to enable human-machine conversations using pre-fed responses.

1995 A.L.I.C.E chatbot

A.L.I.C.E uses natural language processing to enable conversations with humans to flow more naturally.

May 11, 1997 Deep Blue defeats chess champion Garry Kasparov


IBM's supercomputer Deep Blue is the first to defeat a reigning world chess champion.

October 4, 2011 Siri is integrated into the iPhone 4S

Voice recognition software Siri is the first 'intelligent virtual assistant' for mobile devices.

March 15, 2016 AlphaGo defeats human Go champion

AlphaGo wins against Lee Sedol, the greatest Go player in the world.



January 13, 2011 IBM Watson wins Jeopardy!

IBM's Watson beats human champions to win quiz show *Jeopardy!*

September 30, 2012 ImageNet Challenge

A breakthrough deep learning model which significantly improves the accuracy of image recognition.

*October 18, 2016 Microsoft researchers achieve benchmark in speech recognition

AI technology recognises the words in a conversation as well as a person does.

Microsoft projects and collaborations are indicated with an asterisk and highlighted in a blue box.

***July 12, 2017**

Microsoft AI for Earth

Microsoft launches the AI for Earth programme, to harness AI to address environmental challenges such as climate change, biodiversity loss, and water conservation.

February 14, 2019

GPT-2 language model

OpenAI's GPT-2 language model generates impressive text content, translates, answers questions, summarises passages.

***May 19, 2020**

Microsoft announces its first AI supercomputer

Microsoft announces it has built one of the top five publicly disclosed supercomputers in the world, making new infrastructure available to train extremely large AI models.

***October 14, 2020**

Azure Cognitive Services achieves benchmark in image captioning

Microsoft researchers build an AI system that can generate captions for images with high accuracy.

2017

2018

2019

2020

***December 3, 2018**

Microsoft real-time translation achieves benchmark.

Real-time voice translation and captions is integrated into Office 365, based on deep learning work from the Microsoft Research team.

***February 13, 2020**

Turing-NLG language model

Microsoft Project Turing introduces Turing Natural Language Generation (T-NLG), the largest model ever published at 17bn parameters, which outperforms a variety of language modelling benchmarks.

June 11, 2020

GPT-3

The GPT-3 deep learning model is launched, trained on large text datasets of hundreds of billions of words.

November 30, 2022

ChatGPT

ChatGPT marks a step-change in chat bots, with capability to provide detailed responses and articulate answers across many domains of knowledge.

March 14, 2023

GPT-4

GPT-4's multimodal capabilities allow it to accept and produce text and image inputs and outputs, differentiating it from the text-only capabilities of GPT-3 and 3.5.

2021

2022

2023

***June 21, 2022**

GitHub Copilot

GitHub Copilot is made generally available to individual developers. For the first time in the history of software, AI can be broadly harnessed by developers to write and complete code.

***February 7, 2023**

AI comes to Bing and Edge

Microsoft integrates AI into the new Bing and Edge, creating a 'new era of search' with added conversational tools.

***March 16, 2023**

Microsoft 365 Copilot

Microsoft announces a new AI copilot feature for Microsoft 365, with the ability to create documents, read and summarise emails, craft presentations, and more.

Microsoft projects and collaborations are indicated with an asterisk and highlighted in a blue box.

Developing and scaling AI

More than a decade ago, Microsoft forecast an exponential growth in demand for AI systems and started to build special computing infrastructure to handle it.

Inspired by early research developments, we've integrated large-scale language models in services ranging from [Microsoft Bing](#) and Microsoft 365, and have enabled other companies to take advantage of these technologies through Azure AI and the [Azure OpenAI Service](#).

We believe that every organization in the world should benefit from the power of large-scale AI models. We've developed platforms, tools, and a supercomputing infrastructure that would allow any developer to build and scale their own AI innovation.

Through Azure OpenAI Service, we provide businesses and developers with access to high-performance AI models, such as GPT-4, Codex, and DALL-E 2. The service is based on the same platform we use to power AI models in our own products, including [GitHub Copilot](#), [Power Platform](#), [Microsoft Designer](#), [Bing and Edge](#).

We continue to evolve our AI infrastructure based on feedback and insights from training and serving AI models at scale, and our teams are working in

lockstep with industry partners on the design of processors, networks, and datacenters optimized for AI.






The convergence of innovations in infrastructure, machine learning acceleration software, platform service, and modeling powered by cloud technology has created the perfect conditions to accelerate innovation in AI and enable every company to become an AI technology company.

Putting people at the centre

At every stage from development to deployment and ongoing supervision, people play an essential role in AI: collecting, preparing, and curating data to train AI systems; evaluating which algorithms or models to use; validating system performance in real-world scenarios to assess its behavior and robustness before deployment; gathering and analyzing ongoing user feedback to identify errors, biases, or limitations to make improvements to the data, algorithms, or models; and many more roles.

By prioritizing safeguarding throughout, AI developers and practitioners can put in place measures to mitigate potential harm, ensure fairness, and promote transparency. This hands-on involvement is critical to shaping and guiding AI systems to achieve reliable and beneficial outcomes.

The technology stack for AI foundation models

 Applications	Software programs where the output of an AI model is put to work
 API Services	APIs (Application Program Interfaces), or endpoints, through which applications access pre-trained models
 Powerful Pre-Trained AI Models	Pre-trained models like GPT-4 that can be used to solve similar problems without starting from scratch
 Machine Learning Acceleration Software	Software that speeds up the process of developing and deploying large AI models
 AI Datacenter Infrastructure	Advanced supercomputing infrastructure, including clusters of advanced GPUs (Graphics Processing Units) with high bandwidth network connections

Why the cloud is the critical infrastructure behind AI:

- Training AI models requires both large datasets and AI-optimized cloud computing infrastructure.
- Cloud-based AI platforms make it easier for organizations to develop their own AI applications.
- Cloud computing resources give organizations of all sizes a cost-effective way to run and scale AI innovation.

Microsoft is committed to building Azure into an AI supercomputer for the world, serving as the foundation of our vision to democratize AI as a platform. We pushed the frontier of cloud supercomputing technology, [announcing our first top-5 supercomputer in 2020](#), and have subsequently constructed multiple AI supercomputing systems at massive scale.



New steps toward
democratizing AI

The progress toward democratizing AI and the opportunities it affords marks a transformative shift, extending AI's impact beyond the often-arcane realm of data scientists and developers in the back office to everyone in our communities.

Generative AI, in particular, represents a major leap for democratization of the technology, helping to eliminate the need for specialized AI knowledge to harness AI capabilities. Through the use of natural language, more people are able to benefit from powerful AI tools. LLMs have the capability to provide 'no code' services, which empower people to be able to interact with AI systems without coding and to create results without a need for technical expertise. By simply asking a question in one's native language, people can create content and summarize data effortlessly, broadening the audience for AI from IT experts to anyone familiar with a browser and a connection to the internet.

It is key to recognize that internet connectivity and digital literacy are essential to truly achieving democratic access. Microsoft recognizes this imperative and is actively working to bridge gaps in

these areas. Through initiatives like the [Microsoft Airband Initiative](#), Microsoft endeavors to make affordable broadband accessible to underserved global communities, having connected 51 million people to high-speed internet since 2017 and having set the goal to connect 250 million by 2025.

Microsoft's investments in skilling programs like [AI Skills Initiative](#), offering courses like Generative AI on LinkedIn Learning, have already reached 80 million people worldwide.

Consistent with our investments in connectivity and skilling, Microsoft is focused on fairness and inclusion, two key principles that help to form our responsible AI commitment. We also collaborate with researchers, civil society groups like the Canadian Civil Liberties Association, and governments and international agencies to inform our efforts to align AI development and use with diverse community interests. We seek feedback to help ensure that, as we achieve more equitable access, we're also prepared to support meaningful use among individuals and groups with different backgrounds and interests.



AI
in Action

Canada was the first country in the world to create a national strategy for AI, releasing it in 2017. The [Pan-Canadian Artificial Intelligence Strategy](#) is a \$125 million initiative delivered by CIFAR (Canadian Institute for Advanced Research), a globally recognized research organization based in Canada. The Strategy aims to attract and retain top AI talent, support AI research excellence, and foster AI innovation and collaboration across Canada and has funded three AI institutes: the Vector Institute in Toronto, Mila in Montreal, and AMII in Edmonton, as well as 80 Canada CIFAR AI Chairs across the country. Together, these entities collectively shape and bolster the success of AI initiatives in the country.

According to a report produced by Deloitte on behalf of CIFAR, Amii, Mila, and the Vector Institute entitled '[Impact and opportunities: Canada's AI ecosystem](#),' Canada boasts a dynamic AI landscape comprising diverse entities, including

more than 670 AI startups and 30 generative AI companies. The report ranks Canada fourth globally in generative AI companies per capita and is third among G7 countries in total funding per capita raised for generative AI companies, noting its robust ecosystem of established enterprises, generative AI vendors, research laboratories, business incubators, accelerators, venture capital firms, leading AI researchers and innovators, as well as regulatory bodies.

In recent months, technology developments have put a spotlight on advancements in generative AI and the possibilities for how it can shape the future of industry and society in Canada. The pace and scale of technological innovations happening today are unprecedented and breathe optimism into a demanding economy. Innovation is also providing a powerful boost to our efforts to address the major societal challenges we face in Canada.

HEALTH

We are at a unique moment in history where medicine, biology, and technology are converging on a large scale. Healthcare organizations and hospitals are using AI to develop new treatments and medicines and to help doctors and nurses improve patient care, reduce clinician burnout, and improve overall outcomes.

SUSTAINABILITY

The climate challenge is a race against the clock to decarbonize. AI is helping accelerate and scale our efforts to transition to clean energy and mitigate the impact of global warming on people and communities.



CULTURAL HERITAGE

UNESCO predicts that between 50 and 90 percent of endangered languages will disappear by the next century, threatening cultural heritage around the world. We use AI to work with non-profits, universities, and governments to help preserve heritage and languages and to help advance broader opportunities and a more inclusive future.



CITIZEN SERVICES

By embedding digital solutions into the core of service delivery to citizens, municipal, provincial, territorial, and the federal governments are unlocking creative new ways to deliver personalized services that meet the evolving demands and needs of citizens in real time.



CRISIS MANAGEMENT AND EMERGENCY PREPAREDNESS

During times of crisis and emergencies, from wildfires to floods, and with the recent pandemic, AI surfaces data-informed insights, enhances human decision-making in managing complex domestic and global challenges at rapid speeds, and facilitates prediction and timely responses.



SKILLING FOR THE FUTURE

AI offers tremendous potential to empower workers around the world—but only if everyone, everywhere has the skills to use it. To thrive in an AI-enabled economy, the workforce must be prepared and trained with the digital and AI skills required for the in-demand jobs of the future.



AI for a healthier future

Canada stands at a pivotal moment in the evolution of healthcare. According to [Statistics Canada's demographic projections](#), by 2030, seniors could represent between 21% and 23% of the total population in 2036. This demographic trend necessitates a fundamental transformation in healthcare delivery—a shift from traditional “episodic care-based philosophy to one that is

much more proactive and focused on long-term care management,” according to a [report by McKinsey and EIT Health](#).

The same report surveyed health care professionals, who said AI solutions can ease administrative tasks, facilitate home-based care, and improve clinical practice. Embracing these technological advancements becomes imperative for Canada's healthcare system to effectively navigate the challenges posed by its changing demographic landscape.

According to the [Canadian Medical Association](#), Canada allocates over \$300 billion each year to healthcare, constituting approximately 13% of its GDP, ranking second highest among OECD nations. The healthcare sector in Canada is also facing a staffing crisis with elevated levels of job vacancies, totaling 96,200 unfilled positions in health occupations in the fourth quarter of 2022 and rising.

An ounce of prevention is worth a pound of cure, and AI can help accelerate public health and clinical research at all levels of government to identify solutions to some of the most pressing healthcare challenges.

Mitigating the spread of COVID-19 with AI in Vancouver

The pandemic placed unprecedented demands on Canada's healthcare system. As health officials fought to mitigate the spread of COVID-19, a multidisciplinary collaboration was created between [Providence Health Care](#), Microsoft Canada, and computer vision researchers at the University of British Columbia to develop an AI-powered solution to monitor the effectiveness of social distancing policies and guidelines at St. Paul's Hospital in Vancouver. The project utilized deep learning AI tools to monitor personal protective equipment usage, social distancing, and occupancy levels in real time.

British Columbia accelerating cancer research by harnessing AI and Cloud computing

Microsoft's AI for Health initiative, a \$40 million program, played a crucial role, enabling the harnessing of AI, machine learning, and cloud computing to accelerate cancer research. A key

challenge in cancer research is that data often remains isolated in databases across different research centers. By integrating AI technologies, researchers can now unravel hidden insights within vast datasets, potentially leading to breakthroughs in cancer treatments.

To facilitate these breakthroughs, the [Cascadia Data Discovery Initiative \(CDDI\)](#) established a searchable database with available datasets, fostering collaboration and shaping new research questions. Spearheaded by Fred Hutchinson Cancer Research Center and BC Cancer Agency, and supported by Microsoft, CDDI's goal was to create a regional data-sharing ecosystem.

With this initiative, CDDI is not just a step toward improved cancer research, it is a leap toward a future where data-driven collaboration powered by AI transforms healthcare, making personalized treatments and cures a reality. The initiative emphasizes the importance of data sharing in scientific progress and embodies the ethical responsibility to maximize the use of tissue samples, ensuring they drive research forward instead of languishing in freezers or databases.

The Cascadia region of North America is home to some of the world's leading technology, research, and medical organizations. Scientific collaboration across these research institutions will accelerate research. At the centre of this collaboration is a focus on the coordinated generation of data and effective sharing. Development and robust regional collaborations and a data sharing ecosystem will position the Pacific Northwest as a global innovator in biomedical research and healthcare, now and into the future.



AI for a more sustainable future

Sustainability is the challenge of the century. According to the World Meteorological Organization (WMO), there is a 98% likelihood that at least one of the next five years, and the five-year period, will be the warmest on record. To curb the progression of climate change, humans urgently need to change the way we use the planet's resources.

A more sustainable future in Ontario through Azure IoT and AI technology

In a pioneering alliance, [Ontario Power Generation \(OPG\)](#) and Microsoft have unveiled a strategic partnership, signaling a significant leap toward a sustainable, net-zero future for Ontario.

Microsoft and OPG worked together to co-develop an innovative hourly energy matching platform powered by Microsoft's Azure services, including Azure IoT Central and Azure AI technology. This platform matches customers' energy needs with local carbon-free sources, enhancing accuracy and reducing emissions. OPG's modernization journey also includes transitioning to Microsoft's Azure cloud platform, ensuring efficient operations while reducing the organization's carbon footprint. This strategic partnership not only helps to shape Ontario's sustainable future, but it also serves as a global model for environmentally conscious economic growth.

A landmark feature of this partnership is Microsoft's procurement of Clean Energy Credits (CECs) from OPG's hydro and nuclear assets. By sourcing carbon-free energy, Microsoft is accelerating its global initiative to power its datacenters entirely with renewable energy, around the clock.

Vancouver pioneering AI to revolutionize farming and the sector's environmental impact

With 25% of greenhouse gas emissions linked to agriculture, and farmers struggling economically, [Terramera](#), a Vancouver-based company, envisions a future where synthetic pesticide use is reduced by 80% globally by the end of the decade. To achieve this, Terramera employs AI, with data from IoT devices, weather patterns, and GPS (Global Positioning Systems), collected through Microsoft Azure FarmBeats.

Their comprehensive approach not only improves agricultural yields but also mitigates climate change by sequestering carbon from the atmosphere, and boosts productivity by up to 70%. For Terramera, the mission is not just about balancing chemical use but creating a holistic, outcomes-based approach to agriculture, ensuring sustainability, profitability, and environmental preservation are interconnected.

Montreal startup reshaping food delivery and cutting waste with Azure credits and AI

Montreal-based [Radish Cooperative](#) is reshaping the food delivery landscape through the strategic integration of AI, championing innovation in sustainability and food service. Their cooperative model unites restaurants, couriers, and customers in ensuring quality and efficiency by leveraging Microsoft Azure's advanced AI services.

Radish conducts in-depth data analysis, enabling restaurants to enhance operations, cut waste, and foster sustainability. Radish's success story underscores the transformative potential of AI in the food sector, promoting economic viability, reducing environmental impact, and advocating for sustainable practices.



AI for preserving cultural heritage

Microsoft introduces Inuktitut to Microsoft Translator through the power of AI in Nunavut

The Government of Nunavut's strategic adoption of Microsoft Teams amidst the pandemic not only streamlined administrative processes, but AI technology also emerged as a beacon of opportunity for Indigenous language preservation and cultural inclusion. Inuktitut is spoken by approximately 40,000 Inuit across Inuit Nunangat, the Inuit homeland in Canada, where 70% of the residents of Nunavut use the Inuktitut language. Displaying the immense potential of AI in safeguarding rich traditions, Microsoft introduced Inuktitut text translation to Microsoft Translator to allow users to translate from more than seventy languages to or from Inuktitut.

Now AI-powered tools are used to add Inuktitut translation to apps, websites, and workflows. This provides an additional way to make Inuktitut more accessible at work, at school, and in everyday life and help the language continue to thrive. In Nunavut, AI became an enabler of progress, connecting communities, preserving languages, and fostering inclusivity. Nunavut's story exemplifies how AI can be pivotal in enabling effective communication, cultural exchanges, and preserving Indigenous languages as a fundamental and valued element of Canadian culture and society.



AI for delivering services to Canadians

Kelowna Develops AI tool to speed up building permit applications

In collaboration with Microsoft, the [City of Kelowna](#) is at the forefront of municipal innovation in British Columbia by harnessing the power of AI to expedite housing permit approvals. Using a \$350,000 grant from the provincial Local Government Development Approvals Program, the city developed an AI chatbot to automate permit applications and address applicants' queries about zoning bylaws and the official community plan.

This innovative tool not only provides considerable time savings for frontline city staff, but it also enhances efficiency and consistency in delivering services to its constituents and improves the housing approval processes. Kelowna's pioneering approach involves AI without replacing staff, as it does not diminish the importance of human oversight. Kelowna's commitment to AI reflects the ongoing synergy between technology and governments, promising a more streamlined and responsive process to deliver efficient services to citizens.

The City of Laval revolutionizes the delivery of services to citizens with AI

Facing the challenge of handling over 250,000 non-emergency requests annually, [the City of Laval](#) is using AI to transform its 311 non-emergency response system to expedite citizen-agent interactions.

The introduction of the 311 virtual agent is positively impacting the citizen experience. By eliminating the clerical task of entering the request in the system, the virtual agent is reducing wait times. It is also allowing city employees to respond to complex requests sooner.

Ottawa's 311 AI-powered chatbot reduces the burden of high-volume calls

The [City of Ottawa](#) has transformed its 311 services with an AI-powered chatbot developed in collaboration with Microsoft. Starting its pilot in early 2020, the chatbot, created using Microsoft's Power Virtual Agent, allows users to ask questions about city services and receive immediate responses. It is designed to evolve over time, learning from user interactions and feedback. Initially focusing on queries about waste diversion and citizen feedback, the chatbot will later expand to manage tasks like filling out permit applications.

This initiative aligns with Microsoft's no-code/low-code development approach, allowing non-developers to create

sophisticated chatbots. Ottawa has emphasized citizen participation in the pilot to improve the technology. The chatbot complements existing 311 services, reducing the burden on call centers during high-call volume time frames. It provides citizens with a convenient and efficient way to access information, highlighting how government-citizen interactions may be enhanced through AI technology.



AI for crisis management and emergency preparedness

Fighting Canada's record Alberta wildfires with a combination of AI and intuition

In the face of Canada's worst wildfire season, [Alberta Wildfire](#), in partnership with Microsoft and AltaML, harnessed AI to make crucial decisions. AltaML developed an AI tool analyzing extensive historical data, weather conditions, forest status, carbon emissions, and human behavior correlations. This AI-powered system predicts the likelihood of new fires, aiding duty officers in strategic resource allocation. Traditional decisions, often based on intuition and past methods, led to overestimating resources, resulting in excessive costs.

The AI tool, accurately predicting new wildfires 80% of the time, empowers officers, bridging the gap between experienced and less-experienced managers. The system, part of

GovLab.ai's efforts (an innovation hub formed by AltaML, the government of Alberta and [Mitacs](#), a nonprofit research organization), not only optimizes resources and can save \$2–\$5 million in annual operating costs, but is also a testament to public-private partnerships in tackling societal challenges. It underlines AI's potential to revolutionize wildfire management, providing personalized, region-specific insights, thereby simplifying complex problems, and enabling swift, effective responses to natural disasters.

Stantec launches digital AI-powered products to predict and manage floods

Stantec is a global leader in sustainable design and engineering with an impressive record of accomplishments and success in large-scale infrastructure and facilities projects. The company is charting an innovative path into the future, launching AI-powered digital products to address challenges such as predicting and managing flood events.

[Stantec's Flood Manager](#) analyzes vast datasets twenty-six times faster, enabling communities to swiftly comprehend and mitigate flood risks and helps prepare and protect communities by enabling scientists and engineers to quickly run models on potential flood risks. This involves analyzing enormous amounts of data, such as information on river levels, land use, and water runoff patterns. Using Azure high-performance computing, Stantec is refining its automation

for ingesting and managing the data. This AI-driven precision empowers emergency planners and ensures proactive and timely responses in real time.



AI for education and skilling

Building skills for an AI-ready workforce

There is an urgent need to transform education and training in Canada and around the world. Canada is currently grappling with a significant shortage of digital and STEM (Science, Technology, Engineering and Math) skills, a challenge intensified by a tight labour market and the growing demand for digital-oriented jobs. A [report from the C.D. Howe Institute](#) emphasizes the urgent need for Canada to transform its education system to bridge this skills gap.

UNESCO also wants countries to harness the potential of AI technologies for learning “while ensuring that its application in educational contexts is guided by the core principles of inclusion and equity.” It notes in the Beijing Consensus, which sets out a common position on AI and learning, that the promise of “AI for all” must be that everyone can take advantage of the technological revolution under way and access its benefits, notably in terms of innovation and knowledge. And for this to

happen, there is a growing need for researchers, educators, and students to have a foundational understanding of AI and data use to be able to engage positively, critically, and responsibly with this technology.

Since June 2020, Microsoft and [LinkedIn](#) have helped over 80 million learners access hundreds of courses to help them take the next step in their careers. Now, as more people transition to technical roles, eight free Professional Certificates are offered to help people learn the foundational skills and relevant hard and soft skills needed for in-demand tech and tech-enabled jobs. One of the certificates addresses how AI and generative AI models work and how to use AI tools to help open up new opportunities, including functionality, ethical considerations of using generative AI, its impact, and a first look at Microsoft 365 Copilot, which will change the way we work.

Proportion of Canadian enterprises reporting difficulties filling ICT (Information Communications Technology) specialist vacancies

According to a [2021 report from KPMG](#), 68% of Canadian businesses struggled to hire skilled tech talent to foster measurable growth. There continues to be a growing number of available tech roles in Canada, with an estimated [2.26 million digital economy jobs predicted by 2025](#).

[KPMG's 2023 Global Tech Report](#) entitled 'Innovating with intention: The path to tech transformation,' also identified AI as the top

technology to help businesses achieve their goals over the next three years, and that 37% of the 'Canadian organizations that undertook AI digital transformation efforts over the last two years saw an increase in performance or profitability of 11% or over.' The same report stated that one of the top transformation challenges for Canada was the lack of skills within organizations, making education and skilling a necessity for organizations to remain competitive.

Helping 6,000 jobseekers launch their tech careers

In October 2023, [NPower Canada announced a significant milestone in its Canadian Tech Talent Accelerator \(CTTA\) initiative](#), with a continued co-investment of \$2 million from DIGITAL, Microsoft, and the CIBC Foundation.

Launched in 2021, the CTTA has been instrumental in Canada's post-pandemic economic recovery by providing digital skills, certifications, and job placement support to over 3,300 jobseekers, with 80% of graduates finding sustainable employment within six months.

The investment will enable the program to scale, launching more than 6,000 jobseekers into tech careers. In the next phase, NPower Canada will introduce a cybersecurity skilling program alongside existing programs in Data Analytics, IT Service Management, and Artificial Intelligence. This collaboration signifies Microsoft's commitment to expanding economic opportunity and growth for all.



Responsible Development and use of AI



Responsible AI – Q&A with Natasha Crampton

Natasha Crampton leads Microsoft's Office of Responsible AI, as the company's first Chief Responsible AI Officer. The Office of Responsible AI puts Microsoft's AI principles into practice by defining, enabling, and governing the company's approach to responsible AI. The Office of Responsible AI also collaborates with stakeholders within and outside the company to shape new laws, norms, and standards to help ensure that the promise of AI technology is realized for the benefit of all.

Why is responsible AI so important?

AI offers tremendous opportunities to improve our world—stimulating economic growth, making work more enjoyable and helping address pressing societal issues like sustainability. However, like every technology before it, some people, communities, and countries will turn this advancement into both a tool and a weapon. Some people will use this technology to exploit the flaws in human nature, deliberately target people with false information, undermine democracy, and explore new ways of pursuing criminal endeavors.

This creates a profound sense of responsibility, both in terms of AI development and its deployment. Since 2017, Microsoft has invested in a cross-company program to ensure that our AI systems are responsible by design.

How does Microsoft approach responsible AI?

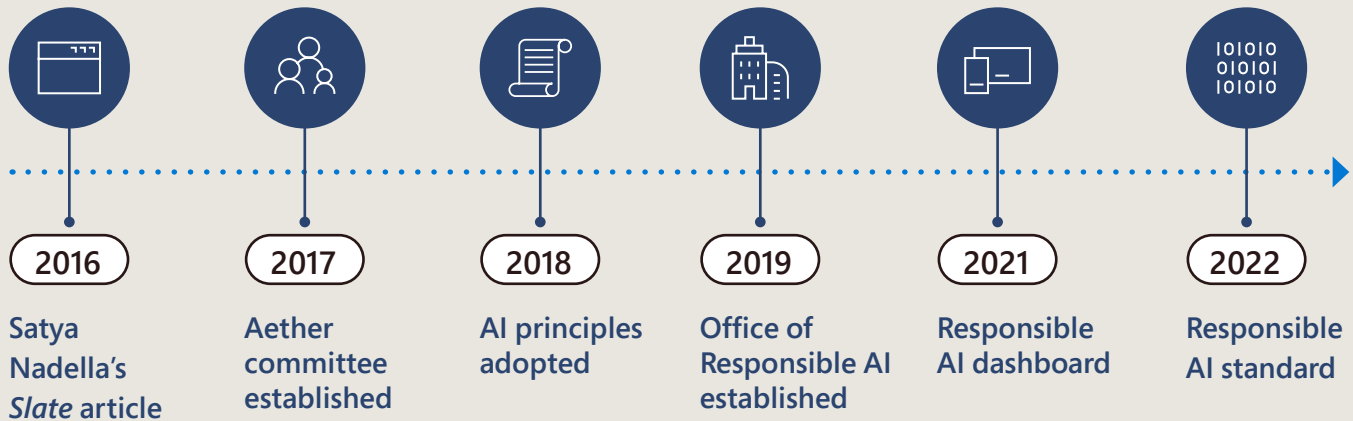
For Microsoft, responsible AI development is about both a practice and a culture.

Practice is how we formally operationalize responsible AI across the company, through governance processes, policy requirements, and tools and training to support implementation. Culture is how we empower our employees to actively champion responsible AI.

We see three essential tenants of upholding the right culture and practices:

- 1. Leadership from the top.** Microsoft's responsible AI journey began in 2016 when Microsoft's Chairman and CEO, Satya Nadella, penned an [article](#) articulating his vision for

Our Responsible AI Journey



humanity powered by AI. Satya's vision set in motion the beginnings of our core AI principles. Microsoft's responsible AI efforts continue to be championed by Microsoft's senior leadership team today. Microsoft's Responsible AI Council is chaired by Microsoft's Vice Chair and President, Brad Smith and by Chief Technology Officer Kevin Scott and convenes senior staff from research, policy, and engineering teams who lead responsible AI efforts across the company. The council serves as a forum for leadership alignment and oversees the progress we are making on our responsible AI commitments.

2. Inclusive governance models. Since starting our work on responsible AI seven years ago, we learned the need to create a governance model that was inclusive and encouraged multi-disciplinary collaboration

across engineering, policy, and research. The governance model established by Microsoft's Office of Responsible AI ensures that there are senior leaders tasked with spearheading responsible AI across every business group. These leaders are supported by a large network of trained 'responsible AI champions. Our continually updated [Responsible AI Standard](#), which is publicly available, offers this community guidance for building AI systems responsibly. Taking lessons from long-standing, cross-company commitments to privacy, security, and accessibility, we realized that responsible AI must be supported by the highest levels of leadership in the company and championed at every level across Microsoft. Responsible AI is a company-wide mandate.

3. Invest in people. Microsoft currently has nearly 350 people working on responsible AI, with 129 dedicated to it full-time. These team members hold positions across the company in policy, engineering, research, sales, and other core functions. We are committed to continue investing in hiring diverse talent, assigning additional talent to focus on responsible AI full time, and upskilling more people throughout the company.

How does the Responsible AI Standard help teams put Microsoft's principles into practice?

From crafting an AI system's purpose to designing how people interact with it, we must keep people at the center of all AI decisions. While our responsible AI principles state the enduring values we seek to uphold, we needed more specific guidance on how to build and deploy AI systems responsibly. This is why we developed our Responsible AI Standard, a more practical guide that memorializes a set of 'rules of the road' for our engineering teams so that upholding our AI principles is a daily practice.

The Responsible AI Standard was the result of a multi-year, cross-company effort that reflected

a vast array of input from researchers, engineers, lawyers, designers, and policy experts. We consider it to be a significant step forward for our practice of responsible AI because it sets out concrete, practical guidance on how to identify, measure, and mitigate harms ahead of time. It also requires teams to adopt tools and controls to secure beneficial uses while guarding against potential misuses of their products.

When building and updating the Responsible AI Standard, we recognized early on that it is impossible to reduce all the complex sociotechnical considerations—for many different use cases—into an exhaustive set of pre-defined rules. This led us to create a program and process for ongoing review and oversight of high-impact cases and rising issues and questions, which we call Sensitive Uses.

The Sensitive Uses review process is triggered when Microsoft personnel are involved in developing or deploying an AI system and the foreseeable use or misuse of that AI system could have a consequential impact on a user's legal status or life opportunities, present the risk of significant physical or psychological injury, or restrict, infringe upon, or undermine the ability to realize an individual's human rights.

Microsoft's responsible AI principles



Fairness

AI systems should treat everyone fairly and avoid affecting similar situated groups of people in different ways.



Reliability & safety

To build trust, it is critical that AI systems operate reliably, safely, and consistently under normal circumstances and in unexpected conditions.



Privacy & security

AI systems should be secure, respect privacy laws, and provide appropriate consumer controls.



Inclusiveness

AI systems should empower everyone, meaning they must incorporate and address a broad range of human needs and experiences.



Transparency

AI systems should be intelligible and explainable, so people can understand how they are used to inform decisions that impact their lives.



Accountability

People who design and deploy AI systems must be accountable for how their systems operate.



Helping our Canadian customers develop and deploy AI responsibly – Antony Cook

Antony Cook leads Microsoft Customer and Partner Solutions group, a team comprised of more than 350 legal and government affairs professionals located in more than 50 countries around the world. They focus on all legal support for the global sales organization including pioneering new legal and corporate affairs solutions that advance business goals to benefit consumers, industry partners, and communities.

Microsoft has a responsible AI process that our AI systems go through, which includes reviews by a multidisciplinary team of experts, that helps us understand potential harms and find mitigations. Examples of mitigations include refining the dataset used to train models, deploying filters to limit the generation of harmful content, integrating techniques like query blocking on sensitive topics that helps to prevent misuse by bad actors or applying technology that can return more helpful, representative, and diverse responses and results.

Another very important safeguard is intentional and iterative deployment. We take a measured approach to rolling out many of these new AI powered experiences. We may start with a limited preview with a select number of customers with well-defined use cases in mind. Collaborations with these early customers will help us make sure the responsible AI safeguards are working in practice so we can scale adoption more broadly.

One of our most important responsible AI commitments is sharing our learnings with customers. We provide transparency documentation for our platform AI services. Transparency Notes communicate in clear, everyday language the purposes, capabilities, and limitations of AI systems so our customers can understand when and how to deploy our platform technologies. They also identify use cases that fall outside the solution's capabilities and the Responsible AI Standard. Transparency Notes fill the gap between marketing and technical documentation, proactively communicating information that our customers need to know to deploy AI responsibly.

Customers also need practical tools to operationalize responsible AI practices. Over the years, responsible AI research at Microsoft has led to the incubation of tools such as Fairlearn and InterpretML. The collection of tools has grown in capability, spanning many facets of responsible AI practice including the

ability to identify, diagnose, and mitigate potential errors and limitations of AI systems. Since their original conception within Microsoft, these tools continue to improve and evolve through the contributions of active open-source communities. Our latest tool, which is in preview, is Azure Content Safety. It helps businesses create safer online environments and communities through models that are designed to detect hate, violent, sexual, and self-harm content across languages in both images and text.

Building on the Responsible AI Toolbox, Microsoft's responsible AI program has invested in integrating some of the more mature responsible AI tools directly into Azure Machine Learning so our customers will also benefit from the development of engineering systems and tools. The collection of capabilities, known as the Responsible AI Dashboard, offers a single pane of glass for machine learning practitioners and business stakeholders to debug models and make informed, responsible decisions as they build AI systems or customize existing ones.

The community involved in developing, evaluating, and using AI extends beyond our direct customers. To serve this broad ecosystem, we publicly share key artifacts from our responsible AI program, including our Responsible AI Standard, Impact Assessment Template and Guide, detailed primers on the implementation of our responsible AI by design approach and collections of cutting-edge research. Our digital learning paths further empower leaders to craft an effective AI strategy, foster an AI-ready culture, innovate responsibly, and more.

The bottom line is that every organization that creates or uses AI systems will need to develop and implement its own governance systems. To help them do so, we will continue to share our knowledge and tools. We will share the work we are doing to build a practice and culture of responsible AI at Microsoft, including key parts of the curriculum that we use to train Microsoft employees. And we continue to invest in dedicated resources and expertise in regions around the world to respond to customer questions about deploying and using AI responsibly.

Furthermore, we have [announced](#) that we are creating an AI Assurance Program to help customers ensure that the AI applications they deploy on our platforms meet the legal and regulatory requirements for responsible AI. This program is open to our customers around the world.

This program will include:

Regulator engagement support: We have extensive experience helping customers in the public sector and highly regulated industries manage the spectrum of regulatory issues that arise when dealing with the use of information technology. For example, in the global financial services industry, we worked closely for a number of years with both customers and regulators to ensure that this industry could pursue digital transformation on the cloud while complying with its regulatory obligations. One learning from this experience has been the industry's requirement that financial institutions verify customer

identities, establish risk profiles, and monitor transactions to help detect suspicious activity, the “know your customer” requirements. We believe that this approach can apply to AI in what we are calling “KY3C,” an approach that creates certain obligations to know one’s cloud, one’s customers and one’s content. We want to work with our customers to apply KY3C as part of our AI Assurance Program.

Customer councils: We will bring customers together in customer councils to hear their views on how we can deliver the most relevant and compliant AI technology and tools.

Regulatory advocacy: Finally, we’ll play an active role in engaging with governments to promote effective and interoperable AI regulation. The recently launched Microsoft blueprint for AI governance (see next section) presents our proposals to governments and other stakeholders for appropriate regulatory frameworks for AI.

We are committed to supporting customers in Canada and around the world implement their own AI systems responsibly, and we will develop responsible AI programs for our partner ecosystem.

Many of our partners have already created comprehensive practices to help customers evaluate, test, adopt and commercialize AI solutions, including creating their own responsible AI systems. We are launching a program with selected partners to leverage this expertise to assist our mutual customers in deploying their own responsible AI systems. PwC and EY are our initial launch partners, and we will be looking to add Canadian partners to the program.

Ultimately, we know that these commitments are only the start, and we will have to build on them as both the technology and regulatory conditions evolve. But we are also excited by this opportunity to partner more closely with our customers as we continue on the responsible AI journey together.

KY3C:

Applying to AI services the “Know Your Customer” concept developed for financial services

Know your Cloud

Know your Customer

Know your Content

Advancing trustworthy AI

For more than 30 years, [Microsoft Research](#) has been advancing the foundations of computing and translating new scientific understanding into innovative technologies to create value for our customers and broad benefit to society.

Our researchers collaborate across disciplines, institutions, and geographies to deliver cutting-edge advances in vision, speech, language, decision-making, and machine learning. They have pioneered AI breakthroughs in [conversational speech recognition](#), [machine translation](#), [image captioning](#), [natural language understanding](#), and [commonsense question answering](#).

As part of our ongoing commitment to [building AI responsibly](#), research scientists and engineers at Microsoft are also pursuing methods and technologies aimed at helping builders of AI systems cultivate appropriate trust—that is, building trustworthy models with reliable

behaviours and clear communication that set proper expectations. When AI builders plan for failures, work to understand the nature of the failures, and implement ways to effectively mitigate potential harms, they help engender trust that can lead to a greater realization of AI's benefits.

Facilitating trustworthy measurement, improving human-AI collaboration, designing for natural language processing (NLP), advancing transparency and interpretability, and exploring the open questions around AI safety, security, and privacy are key considerations for developing AI responsibly. The goal of trustworthy AI requires a shift in perspective at every stage of the AI development and deployment life cycle. We're actively developing a growing number of [best practices and tools](#) to help with the shift to make responsible AI more available to a broader base of users. Many open questions remain, but as innovators, we are committed to tackling these challenges with curiosity, enthusiasm, and humility.



AI
Governance



Advancing AI governance in Canada and internationally - Brad Smith

As Microsoft's vice chair and president, Brad Smith is responsible for spearheading the company's work and representing it publicly on a wide variety of critical issues involving the intersection of technology and society, including artificial intelligence, cybersecurity, privacy, environmental sustainability, human rights, digital safety, immigration, philanthropy, and products and business for non-profit customers. He leads a team of roughly 2,000 business, legal, and corporate affairs professionals located in 54 countries and operating in more than 120 nations.

As this paper highlights, there are enormous opportunities to harness the power of AI to contribute to Canadian growth and values. But another dimension is equally clear. It's not enough to focus only on the many opportunities to use AI to improve people's lives. We need to focus with equal determination on the challenges and risks that AI can create, and we need to manage them effectively.

This is perhaps one of the most important lessons from the role of social media. Little more than a decade ago, technologists and political commentators alike gushed about the role of social media in spreading democracy during the Arab Spring. Yet, five years after that, we learned that social media, like so many other technologies before it, would become both a weapon and a tool—in this case aimed at democracy itself.

Today, we are 10 years older and wiser, and we need to put that wisdom to work. We need to think early on and in a clear-eyed way about the

problems that could lie ahead. As technology moves forward, it's just as important to ensure proper control over AI as it is to pursue its benefits. We are committed and determined as a company to develop and deploy AI in a safe and responsible way. We also recognize, however, that the guardrails needed for AI require a broadly shared sense of responsibility and should not be left to technology companies alone. In short, tech companies will need to step up, and governments will need to move faster.

When we at Microsoft adopted our six ethical principles for AI in 2018, we noted that one principle was the bedrock for everything else—accountability. This is the fundamental need to ensure that machines remain subject to effective oversight by people, and the people who design and operate machines remain accountable to everyone else. In short, we must always ensure that AI remains under human control. This must be a first-order priority for technology companies and governments alike.

A five-point blueprint for governing AI

- 1 Implement and build upon new government-led AI safety frameworks
- 2 Require effective safety brakes for AI systems that control critical infrastructure
- 3 Develop a broader legal and regulatory framework based on the technology architecture for AI
- 4 Promote transparency and ensure academic and public access to AI
- 5 Pursue new public-private partnerships to use AI as an effective tool to address the inevitable societal challenges that come with new technology

This connects directly with another essential concept. In a democratic society, one of our foundational principles is that no person is above the law. No government is above the law. No company is above the law, and no product or technology should be above the law. This leads to a critical conclusion: People who design and operate AI systems cannot be accountable unless their decisions and actions are subject to the rule of law.

In May, Microsoft released a whitepaper, *Governing AI: A Blueprint for the Future*, which sought to address the question of how do we best govern AI and set out Microsoft's [five-point blueprint](#).

The blueprint builds on lessons learned from many years of work, investment, and input. Here we draw some of that thinking together with

a specific focus on Canada, its initiatives on AI regulation, and what a viable path to advancing AI governance internationally may look like.

Progress towards the democratization of AI is making powerful tools available to everyone, but that also includes people who might choose to use these tools as weapons. Thoughtful regulation is required to establish the guardrails and norms for everyone that has a role in developing, deploying, and using powerful tools. In drafting regulation for the responsible use of AI, a human-centered approach harnessing multiple representative voices should reflect the coming together of government policymakers, business, civil society, academia, and youth entrepreneurs, to be as representative of community needs and experiences as possible. Microsoft's strategy is to support multisectoral, multistakeholder

communities convened by well recognized and trusted organizations. We aim to develop capacity building programs, with opportunities for hands-on AI engagement by regulators so that regulation balances the goals to accelerate the benefits of AI while establishing necessary guardrails and safety brakes that maintain effective human oversight and control.

From early on, we have been supportive of Canada's efforts to develop a regulatory regime that effectively addresses safety and upholds fundamental rights while continuing to enable innovations that will ensure that Canada remains globally competitive. Our intention is to offer constructive contributions to help inform the work ahead. Collaboration with leaders and policymakers across Canada and around the world is both important and essential.

In this spirit, here we want to expand upon our five-point blueprint, highlight how it aligns with Canada's *Artificial Intelligence and Data Act* (AIDA) discussions, and provide some thoughts on the opportunities to build on this regulatory foundation.

First, implement and build upon new government-led AI safety frameworks

One of the most effective ways to accelerate government action is to build-on existing or emerging governmental frameworks to advance AI safety.

A key element to ensuring the safer use of this technology is a risk-based approach, with defined processes around risk identification

and mitigation as well as testing systems before deployment. As the Canadian Parliament debates and finalizes AIDA, its initial focus on a risk-based approach provides an important benchmark for the future. In the interim, Microsoft has been supporting federal, provincial, and municipal governments as they look to establish regulation, policy, and guidance for the responsible use of AI. For example, Microsoft has made the Public Services and Procurement Canada (PSPC) Supplier's commitment to support the Government of Canada's effort in leading the way on the responsible use of AI.

In other parts of the world, other institutions have advanced similar work, such as the AI Risk Management Framework developed by the U.S. National Institute of Standards and Technology, or NIST, and the new international standard ISO/IEC 42001 on AI Management Systems, which is expected to be published in the Fall of 2023.

Microsoft has committed to implementing the NIST AI Risk Management Framework and we will implement future relevant international standards. Opportunities to align such frameworks internationally should continue to be an important part of the ongoing Canada-US dialogue on AI.

As Canada finalizes AIDA, Canada could consider using procurement rules to promote the use of relevant trustworthy AI frameworks. For instance, when procuring high-risk AI systems, procurement authorities could require suppliers to certify via third-party audits that they comply with relevant international standards.

We recognize that the pace of AI advances raises new questions and issues related to safety and security, and we are committed to working with others to develop actionable standards to help evaluate and address those important questions. This includes new and additional standards relating to highly capable foundation models.

Second, require effective safety brakes for AI systems that control critical infrastructure

Increasingly, the public is debating questions around the control of AI as it becomes more powerful. Similarly, concerns exist regarding AI control of critical infrastructure like the electrical grid, water system, and city traffic flows. Now is the time to discuss these issues.

Our blueprint proposes new safety requirements that, in effect, would create safety brakes for AI systems that control the operation of designated critical infrastructure. These fail-safe systems would be part of a comprehensive approach to system safety that would keep effective human oversight, resilience, and robustness top of mind. They would be akin to the braking systems engineers have long built into other technologies such as elevators, school buses, and high-speed trains, to safely manage not just everyday scenarios, but emergencies as well.

In this approach, the government could define the class of high-risk AI systems that control critical infrastructure and warrant such safety measures as part of a comprehensive approach to system management. New laws could require operators

of these systems to build safety brakes into high-risk AI systems by design. The government could then oblige operators to test high-risk systems regularly. And these systems would be deployed only in licensed AI datacenters that would provide a second layer of protection and ensure security.

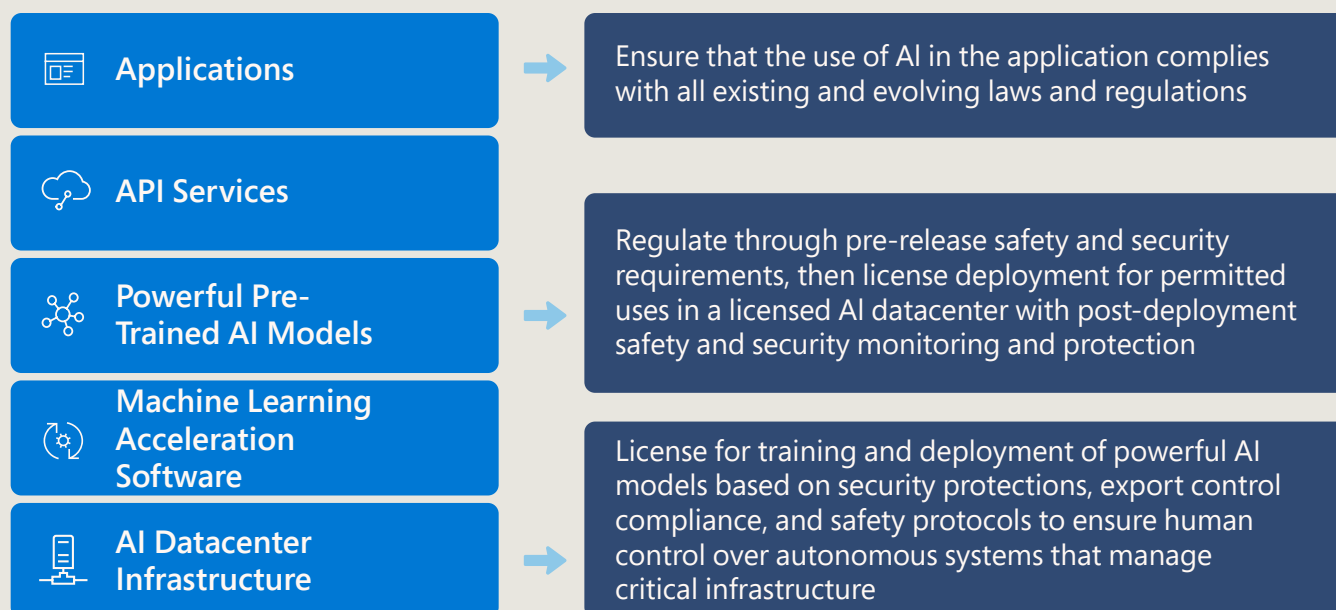
Third, develop a broad legal and regulatory framework based on the technology architecture for AI

As we've worked the past year with AI models at the frontier of this new technology, we've concluded that it's critical to develop a legal and regulatory architecture for AI that reflects the technology architecture for AI itself.

Regulatory responsibilities need to be placed upon different actors based on their role in managing different aspects of AI technology. Those closest to relevant decisions on design, deployment, and use are best placed to comply with corresponding responsibilities and mitigate the respective risks, as they understand best the specific context and use-case. This sounds straightforward, but as discussions in Canada have demonstrated, it's not always easy.

AIDA acknowledges the challenges to regulating complex architecture through its risk-based approach for establishing requirements for high-impact systems. At the application layer, this means applying and enforcing existing regulations while being responsible for any new AI-specific deployment or use considerations.

A proposed AI regulatory architecture



It's also important to make sure that obligations are attached to powerful AI models, with a focus on a defined class of highly capable foundation models and calibrated to model-level risk. This will impact two layers of the technology stack. The first will require new regulations for these models themselves. And the second will involve obligations for the AI infrastructure operators on which these models are developed and deployed. The blueprint we developed offers suggested goals and approaches for each of these layers.

The different roles and responsibilities require joint support. We are committed to helping our customers apply 'Know Your Customer' (KYC) principles through our recently announced AI Assessment Program. Financial institutions use this framework to verify customer identities,

establish risk profiles and monitor transactions to help detect suspicious activity. As Antony Cook details above, we believe this approach can apply to AI in what we are calling "KY3C": **Know one's cloud, one's customers and one's content.**

While AIDA does not expressly include KYC requirements, we believe such an approach will be key to meeting both the spirit and obligations of the Act. AIDA does begin to lay a foundation of requirements for a person who is responsible for a high-impact system, including establishing measures that identify, assess, and mitigate risks. We support this approach. Even before AIDA is implemented, we will test all our AI systems prior to release and use red teaming for high-risk systems.

Fourth, promote transparency and ensure academic and nonprofit access to AI

It is also critical to advance the transparency of AI systems and broaden access to AI resources. While there are some inherent tensions between transparency and the need for security, there exist many opportunities to make AI systems more transparent. That is why Microsoft has committed to an annual AI transparency report and other steps to expand transparency for our AI services.

In tackling this issue, we can start with building blocks that exist already. One of these is the Coalition for Content Provenance Authenticity, or C2PA, a global standards body with more than 60 members including Adobe, the BBC, Intel, Microsoft, Publicis Groupe Sony, and Truepic. The group is dedicated to bolstering trust and transparency of online information including releasing the world's first technical specification for certifying digital content in 2022, which now includes support for Generative AI. As Microsoft's Chief Scientific Officer, Eric Horvitz, [said last year](#), "I believe that content provenance will have an important role to play in fostering transparency and fortifying trust in what we see and hear online."

There will be opportunities in the coming months to take important steps together on both sides of the Atlantic and globally to advance these objectives. Microsoft will soon deploy new state-of-the-art provenance tools to help the public identify AI-generated audio-visual content and understand its origin. At Build 2023, our annual developer conference, we announced the development of a new media provenance service.

The service, implementing the C2PA specification, will mark and sign AI-generated videos and images with metadata about their origin, enabling users to verify that a piece of content is AI-generated. Microsoft will initially support major image and video formats and release the service for use with two of Microsoft's new AI products, Microsoft Designer and Bing Image Creator.

We also believe it is critical to expand access to AI resources for academic research and the nonprofit community. Unless academic researchers can obtain access to substantially more computing resources, there is a real risk that scientific and technological inquiry will suffer, including that relating to AI itself. Our blueprint calls for new steps, including those we will take across Microsoft, to address these priorities.

Fifth, pursue new public-private partnerships to use AI as an effective tool to address the inevitable societal challenges that come with new technology

One lesson from recent years is that democratic societies often can accomplish the most when they harness the power of technology and bring the public and private sectors together. It is a lesson we need to build upon to address the impact of AI on society.

AI is an extraordinary tool. But, like other technologies, it too can become a powerful weapon, and there will be some around the world who will seek to use it that way. We need to work together to develop defensive AI technologies that will create a shield that can withstand and defeat the actions of any bad actor on the planet.

Important work is needed now to use AI to protect democracy and fundamental rights, provide broad access to the AI skills that will promote inclusive growth, and use the power of AI to advance the planet's sustainability needs. Perhaps more than anything, a wave of new AI technology provides an occasion for thinking big and acting boldly. In each area, the key to success will be to develop concrete initiatives and bring governments, companies, and NGOs together to advance them. Microsoft will do its part in each of these areas.

Acting now: International partnership to advance AI governance

Canada's early start towards AI regulation offers an opportunity to establish an effective legal framework, grounded in the rule of law. But beyond legislative frameworks at the level of nation states, multilateral public-private partnership is needed to ensure AI governance can have an impact today, not just a few years from now, but also at the international level. In parallel to Canada's focus on AIDA, Canada has also showed early leadership in publishing its own Code of Conduct for Advanced AI Systems, to support Canadian companies in their push toward responsible AI.

This is important to serve as an interim solution before regulations such as AIDA come into effect, but, perhaps more importantly, it will help us work towards a common set of shared principles that can guide both nation states and companies alike.

The model of a voluntary code provides an opportunity for Canada, the European Union,

the United States, the other members of the G7 as well as India, Brazil, and Indonesia, to move forward together on a set of shared values and principles. If we can work with others on a voluntary basis, then we will all move faster and with greater care and focus. That's not just good news for the technology world, but for the whole world.

Working towards a globally coherent approach is important, recognizing that AI—like many technologies—is and will be developed and used across borders. And it will enable everyone, with the proper controls in place, to access the best tools and solutions for their needs.

Similarly, at the annual G7 Summit in Hiroshima in May 2023, leaders [committed](#) to "advance international discussions on inclusive artificial intelligence (AI) governance and interoperability to achieve our common vision and goal of trustworthy AI, in line with our shared democratic values." The Hiroshima Process International Guiding Principles and Code of Conduct for Organizations Developing Advanced AI Systems, which was issued in October 2023, represents meaningful progress toward a coordinated, global approach to advancing AI in a manner that is safe, secure, and trustworthy. Microsoft will continue to engage in these global AI governance discussions.

Microsoft fully supports and endorses international efforts to develop such a voluntary code. Technology development and the public interest will benefit from the creation of principle-level guardrails, even if initially they are non-binding.

To make the many different aspects of AI governance work on an international level, we will need a multilateral framework that connects various national rules and ensures that an AI system certified as safe in one jurisdiction can also qualify as safe in another. There are many effective precedents for this, such as common safety standards set by the International Civil Aviation Organization, which means an airplane does not need to be refitted midflight from Brussels to New York.

In our view, an international code should:

- Build on the work already done at the OECD to develop principles for trustworthy AI.
- Provide a means for regulated AI developers to attest to the safety of these systems against internationally agreed standards.
- Promote innovation and access by providing a means for mutual recognition of compliance and safety across borders.

Before AIDA and other formal regulations come into force, it is important that we take steps today to implement safety brakes for AI systems that control critical infrastructure. The concept of safety brakes, along with licensing for highly capable foundation models and AI infrastructure obligations, should be key elements of the voluntary, internationally coordinated G7 code that signatory nation states agree to incorporate into their national systems. High-risk AI systems, relating to critical infrastructure (e.g., transport,

electrical grids, water systems) or systems that can lead to serious fundamental rights violations or other significant harms, could require additional international regulatory agencies, based on the model of the International Civil Aviation Organization, for example. Developing a coherent global approach offers significant benefits for all those involved in developing, using, and regulating AI.

Lastly, we must ensure academic researchers have access to study AI systems in depth. There are important open research questions around AI systems, including how one evaluates them properly across responsible AI dimensions, how one best makes them explainable, and how they best align with human values. The work the OECD is doing on the evaluation of AI systems is seeing good progress. But there is an opportunity to go further and faster by fostering international research collaboration and boosting the efforts of the academic communities by feeding into that process. Canada is well placed to lead on this, partnering with the US.

AI governance is a journey, not a destination. No one has all the answers, and it is important we listen, learn, and collaborate. Strong, healthy dialogue between the tech industry, governments, businesses, academia, and civil society organizations is vital to make sure governance keeps pace with the speed at which AI is developing. Together, we can help realize AI's potential to be a positive force for good.

Sources and references

The Promise of AI

References:

•Microsoft, July 2019, The University Health Network leverages Microsoft Azure and AI to bridge the gap between cancer research and treatment, <https://news.microsoft.com/en-ca/2019/07/09/the-university-health-network-leverages-microsoft-azure-and-ai-to-bridge-the-gap-between-cancer-research-and-treatment/>

•BC Cancer, December 2020, BC Cancer Scientists partnering with Microsoft's artificial intelligence (AI) experts to develop the next generation of cancer imaging and treatment tools, <https://www.bccrc.ca/articles/bc-cancer-scientists-partnering-microsofts-artificial-intelligence-ai-experts-develop-next>

•Microsoft, June 2023, The City of Kelowna increases and speeds access to its services with Azure AI, <https://customers.microsoft.com/en-us/story/1641166525350342888-city-of-kelowna-government-azure-open-ai-service>

•Microsoft, September 2022, OPG and Microsoft announce strategic partnership to power a Net-Zero future for Ontario, <https://news.microsoft.com/en-ca/2022/09/26/opg-and-microsoft-announce-strategic-partnership-to-power-a-net-zero-future-for-ontario/>

•Microsoft, September 2023, Fighting Canada's record wildfires with a combination of AI and intuition, <https://news.microsoft.com/source/features/ai/ai-alberta-canada-wildfire-firefighting/>

•Microsoft, January 2021, Microsoft introduces Inuktituk to Microsoft Translator, <https://news.microsoft.com/en-ca/2021/01/27/microsoft-introduces-inuktituk-to-microsoft-translator/>

• Microsoft, Canadian Tech Talent Accelerator program helps Canadian talent skill up for digital careers, <https://customers.microsoft.com/en-us/story/b97da413-30ef-4c15-b09f-06cad2ae293b?preview=1>

AI Explained

References:

•GitHub Copilot: GitHub, <https://github.com/features/copilot>

•Copilot report: GitHub, September 2022, Research: quantifying GitHub Copilot's impact on developer productivity and happiness, <https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/>

•Power Platform: Microsoft, <https://powerplatform.microsoft.com/en-us/>

•Microsoft Designer: Microsoft, <https://designer.microsoft.com/>

•Bing and Edge: Yusuf Mehdi, Microsoft, February 2023, Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web, <https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web/>

•Azure OpenAI Service: <https://azure.microsoft.com/en-ca/products/ai-services/openai-service#overview>

•Azure AI supercomputer: Microsoft, May 2020, Microsoft announces new supercomputer, lays out vision for future AI work, <https://news.microsoft.com/source/features/ai/openai-azure-supercomputer/>

General sources:

•Microsoft, February 2023, What is Microsoft's approach to AI?, <https://news.microsoft.com/source/features/ai/microsoft-approach-to-ai/>

•Microsoft, March 2023, Explaining AI: The fundamentals and the frontiers, <https://news.microsoft.com/source/features/ai/explaining-ai-the-fundamentals-and-the-frontiers/#approach-to-ai>

•Microsoft, A conversation with Kevin Scott: What's next in AI, <https://news.microsoft.com/source/features/ai/a-conversation-with-kevin-scott-whats-next-in-ai/>

•Microsoft, February 2022, The innovation behind AI at Scale, <https://blogs.microsoft.com/ai-for-business/ai-at-scale-technology/>

New steps toward democratizing AI

References:

•Microsoft, Microsoft Airband Initiative, <https://www.microsoft.com/en-us/corporate-responsibility/airband-initiative>

•LinkedIn, Microsoft launches new AI skills training and resources as part of Skill for Jobs Initiative, <https://www.linkedin.com/pulse/microsofts-launches-new-ai-skills-training-resources-part-behncken/>

•Microsoft, Microsoft Azure AI Fundamentals: Generative AI, <https://learn.microsoft.com/en-us/training/paths/introduction-generative-ai/>

•LinkedIn, Build the Skills you need for the digital economy, <https://opportunity.linkedin.com/skills-for-in-demand-jobs>

AI in Action

Cited sources:

•Government of Canada, Pan-Canadian Artificial Intelligence Strategy, <https://ised-isde.canada.ca/site/ai-strategy/en>

•Deloitte, Impact and opportunities: Canada's AI ecosystem – 2023, <https://www2.deloitte.com/content/dam/Deloitte/ca/Documents/press-releases/ca-national-ai-report-2023-aoda-en.pdf>

- United Nations, International Decade of Indigenous Languages 2022-2032, <https://www.un.org/development/desa/indigenouspeoples/indigenous-languages.html>
- Statistics Canada, Population Projections for Canada (2021 to 2068), Province and Territories (2021 to 2043), <https://www150.statcan.gc.ca/n1/pub/91-520-x/91-520-x2022001-eng.htm>
- McKinsey, March 2020, Transforming Healthcare with AI, <https://www.mckinsey.com/~media/McKinsey/Industries/Healthcare%20Systems%20and%20Services/Our%20Insights/Transforming%20healthcare%20with%20AI/Transforming-healthcare-with-AI.ashx>
- Canadian Medical Association, October 2022, Health care funding in Canada, <https://www.cma.ca/news/health-care-funding-canada>
- World Meteorological Organization, May 2023, Global temperatures set to reach new records in next five years, <https://public.wmo.int/en/media/press-release/global-temperatures-set-reach-new-records-next-five-years>
- C.D. Howe Institute, August 2023, Canada faces a serious shortage in digital and stem skills, <https://www.cdhowe.org/media-release/canada-faces-serious-shortage-digital-and-stem-skills>
- UNESCO, Artificial intelligence in education, <https://www.unesco.org/en/digital-education/artificial-intelligence>
- KPMG, August 2021, Nearly 70 per cent of Canadian businesses struggling to find skilled talent, <https://www.newswire.ca/news-releases/nearly-70-per-cent-of-canadian-businesses-struggling-to-find-skilled-talent-893065566.html>
- ICTC, August 2021, ICTC Labour Market Outlook: Additional demand for digital talent to reach 250,000 by 2025, <https://www.ictc-ctic.ca/news-events/ictc-labour-market-outlook>

[additional-demand-for-digital-talent-to-reach-250000-by-2025](#)

- KPMG, Innovating with intention: The path to tech transformation, <https://assets.kpmg.com/content/dam/kpmg/ca/pdf/2023/09/ca-kpmg-2023-global-tech-survey-canadian-perspective.pdf>
- NPower Canada, October 2023, NPower Canada's workforce development program receives \$2M investment from DIGITAL, Microsoft and the CIBC Foundation to empower 6,000 job-seekers, <https://npowercanada.ca/npower-canadas-workforce-development-program-receives-2m-investment-from-digital-microsoft-and-the-cibc-foundation-to-empower-6000-job-seekers/>

References:

- Microsoft, November 2020, Providence Health Care collaborates with Microsoft and UBC to mitigate spread of Covid-19, <https://news.microsoft.com/en-ca/2020/11/05/providence-health-care-collaborates-with-microsoft-and-ubc-to-mitigate-spread-of-covid-19/>
- Microsoft, July 2019, New Cascadia Data Discovery Initiative accelerates health innovation, <https://blogs.microsoft.com/on-the-issues/2019/07/12/new-cascadia-data-discovery-initiative-accelerates-health-innovation/>
- Microsoft, September 2022, OPG and Microsoft announce strategic partnership to power a Net-Zero future for Ontario, <https://news.microsoft.com/en-ca/2022/09/26/opg-and-microsoft-announce-strategic-partnership-to-power-a-net-zero-future-for-ontario/>
- Michael Hainsworth, Microsoft, June 2021, Real AI for Agriculture and Food Security, <https://news.microsoft.com/en-ca/2021/06/28/real-ai-for-agriculture-and-food-security/>

•Microsoft, June 2023, Radish Cooperative reduces food waste, improves service for customers and business for restaurants with Azure, <https://customers.microsoft.com/en-us/story/1643336731024164757-radish-cooperative-consumer-goods-azure>

•Microsoft, January 2021, Microsoft introduces Inuktitut to Microsoft Translator, <https://news.microsoft.com/en-ca/2021/01/27/microsoft-introduces-inuktitut-to-microsoft-translator/>

•Microsoft, June 2023, The City of Kelowna increases and speeds access to its services with Azure AI, <https://customers.microsoft.com/en-us/story/1641166525350342888-city-of-kelowna-government-azure-open-ai-service>

•Microsoft, September 2021, Ville de Laval speeds up community response systems with an Azure AI solution, <https://customers.microsoft.com/en-us/story/1409585222396616040-laval-gov-azure-canada>

•Microsoft, October 2021, Talk to the bot: City of Ottawa helps employees navigate the pandemic with Power Virtual Agents, <https://customers.microsoft.com/en-us/story/1420166398092426080-city-of-ottawa-power-virtual-agents>

•Microsoft, September 2023, Fighting Canada's record wildfires with a combination of AI and intuition, <https://news.microsoft.com/source/features/ai/ai-alberta-canada-wildfire-firefighting/>

•Microsoft, April 2023, Stantec uses Azure to build a digital innovation platform with real-world impact, <https://customers.microsoft.com/en-us/story/1620794556954301040-stantec-professional-services-azure>

[microsoft.com/en-us/story/1620794556954301040-stantec-professional-services-azure](https://customers.microsoft.com/en-us/story/1620794556954301040-stantec-professional-services-azure)

•LinkedIn, Build the Skills you need for the digital economy, <https://opportunity.linkedin.com/skills-for-in-demand-jobs>

Responsible Development & Use of AI

References:

•Slate, June 2016, The Partnership of the Future, <https://slate.com/technology/2016/06/microsoft-ceo-satya-nadella-humans-and-a-i-can-work-together-to-solve-societys-challenges.html>

•Responsible AI Standard: Microsoft, <https://aka.ms/RAIStandardPDF>

•Microsoft Research: Microsoft, <https://www.microsoft.com/en-us/research/>

•Conversational speech recognition: Microsoft, August 2017, Microsoft researchers achieve new conversational speech recognition milestone, <https://www.microsoft.com/en-us/research/blog/microsoft-researchers-achieve-new-conversational-speech-recognition-milestone/>

•Machine translation: Microsoft, June 2019, Neural Machine Translation Enabling Human Parity Innovations In the Cloud, <https://www.microsoft.com/en-us/translator/blog/2019/06/17/neural-machine-translation-enabling-human-parity-innovations-in-the-cloud/>

- Image captioning: <https://www.microsoft.com/en-us/research/blog/novel-object-captioning-surpasses-human-performance-on-benchmarks/>
- Natural language understanding: Microsoft, October 2020, Novel object captioning surpasses human performance on benchmarks, <https://www.microsoft.com/en-us/research/blog/microsoft-deberta-surpasses-human-performance-on-the-superglue-benchmark/>
- Commonsense question answering: Microsoft, December 2021, Azure AI milestone: Microsoft KEAR surpasses human performance on CommonsenseQA benchmark, <https://www.microsoft.com/en-us/research/blog/azure-ai-milestone-microsoft-kear-surpasses-human-performance-on-commonsenseqa-benchmark/>

General sources:

- Natasha Crampton, Microsoft, May 2023, Reflecting on our responsible AI program: Three critical elements for progress, <https://blogs.microsoft.com/on-the-issues/2023/05/01/responsible-ai-standards-principles-governance-progress/>
- New York Times, May 2023, Microsoft call for A.I. rules to Minimize the Technology's Risks, <https://www.nytimes.com/2023/05/25/technology/microsoft-ai-rules-regulation.html>
- Governing AI: A Blueprint for the Future whitepaper: Microsoft, May 2023, Governing AI: A Blueprint for the Future, <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RW14Gtw>
- Brad Smith, Microsoft, January 2023, Putting AI in the service of the common good, <https://www.linkedin.com/pulse/putting-ai-service-common-good-brad-smith/>
- Antony Cook, Microsoft, June 2023, Announcing Microsoft's AI Customer Commitments, <https://blogs.microsoft.com/blog/2023/06/08/announcing-microsofts-ai-customer-commitments/>
- Microsoft, February 2022, Advancing AI trustworthiness: Updates on responsible AI research, <https://www.microsoft.com/en-us/research/blog/advancing-ai-trustworthiness-updates-on-responsible-ai-research/>
- Eric Horvitz, Microsoft, December 2021, Building AI responsibly from research to practice, <https://blogs.microsoft.com/ai-for-business/building-ai-responsibly-from-research-to-practice/>

AI Governance

References:

- Governing AI: A Blueprint for the Future whitepaper: Microsoft, May 2023, Governing AI: A Blueprint for the Future, <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RW14Gtw>

General sources:

- Brad Smith, Microsoft, May 2023, How do we best govern AI?, <https://blogs.microsoft.com/on-the-issues/2023/05/25/how-do-we-best-govern-ai/>
- Eric Horvitz, Microsoft, January 2022. A Milestone Reached in Digital Content Provenance, https://erichorvitz.com/A_Milestone_Reached_Content_Provenance.htm

