The rising importance of data ethics

The key to earning and retaining customer trust

Introduction

Data is integral to the way businesses, governments, and other organizations function in today's world. Technological advances allow for incredible amounts of data to be created, shared, analyzed, and used to generate insights, foster innovation, improve products and services, and more. But as organizations collect, share, and use ever-growing volumes of data about individuals, concerns have grown about how that data is being used and whether it is right to do so. This is prompting governments, regulators, and companies to grapple with the issue of data ethics: how to protect people's right to privacy and control over their personal data while still enabling organizations to use that data to deliver fresh insights, make smarter decisions, and provide better experiences.

Data ethics is fast becoming a fundamental part of the conversation about data collection and usage—and a matter of interest for boards of directors, C-suite executives, and business leaders. A commitment to strong AI governance, supported by a data ethics framework, will be essential, not only to ensure organizations comply with applicable laws and regulations, but also to enable those organizations to manage their risks, differentiate themselves in the market, and, crucially, earn and retain the trust of customers and others. Without that trust, individuals may choose to withhold their data, cutting off the information organizations increasingly depend on to compete and succeed.

While data ethics is the subject of growing discussion today, it will quickly become an essential part of doing business. That's why it's vital for organizations to begin to take action now. In this report, we explore how to define data ethics and why it needs to be treated as a strategic priority. We review the current regulatory landscape as well as pending developments. We offer suggestions on how to implement or improve a responsible AI program and provide examples of Canadian organizations' own data ethics journeys. And we peer ahead to a number of issues that are likely to have a significant effect on the data ethics conversation in the near term.

Data ethics defined

Organizations increasingly recognize that to stay competitive, innovate, and drive growth, they need to collect and use data. In particular, they need to harness data in a way that helps them understand their customers and make the right customer-focused decisions. And that is where data ethics comes in.

Data ethics is fundamentally a matter of morality. It's about doing the right thing with respect to collecting, sharing, protecting, and using data, especially when that data involves personally identifiable information (PII) about customers, employees, or anyone else. It's about earning the trust of customers and others regarding the organization's use of data—and retaining that trust over the long term. Data ethics means taking time to consider the human impact of collecting, sharing, and using data, whether that data is sourced internally or externally from partners, open sources, or third-party vendors. It means providing individuals with a say in what data they share and how it can be used—and respecting and meeting their wishes. Data ethics involves grounding data-related decisions in the organization's brand values and a clear understanding of the potential financial and reputational impacts of a data misstep. And it means understanding and complying with applicable regulatory requirements for data collection, storage, and usage.

Why data ethics needs to be a strategic priority

Data ethics isn't a niche concern of interest only to chief data officers and data analytics teams. As organizations increasingly depend on harnessing internal and external data to drive growth, improve the customer experience, make better business decisions, and more, they can become increasingly vulnerable to new financial and reputational risks. Failing to understand and manage these data-related risks can have major consequences, including eroded customer trust, a tarnished brand, regulatory penalties, and significant financial damage. Ensuring ethical data use is key to mitigating these risks and reducing the potential for costly missteps—which is why responsible AI needs to be a strategic priority for the business and a key agenda item for C-suite executive and board members alike.

Examples of dubiously ethical data use—and its consequences—abound. Twitter is paying a US $150 million fine in a settlement over the US Federal Trade Commission's allegations that the company misused users' personal information to serve targeted advertising.[1] Google faces a lawsuit alleging the company is collecting sensitive information about students, without proper disclosure or parental consent, in the course of supplying free Chromebooks to schools.[2] The Supreme Court of Canada opened the door to a $400 million class-action lawsuit against Uber after ruling that workers can settle disputes with local jurisdictions.[3]

Implementing an enterprise-wide commitment to AI governance can help organizations avoid running into similar trouble. A well-founded data ethics program helps organizations make their customers feel safe in providing personal data. It makes sure that data-driven insights are auditable, explainable, fair, and representative. It assists the organization in complying with regulatory requirements for data capture, maintenance, and use. An effective data ethics program can also serve as a powerful competitive differentiator as citizens and lawmakers alike grow more concerned about how organizations collect and use individuals' data. Organizations that are clear and transparent about the data they collect and how it will be used and protected will be better positioned to gain individuals' trust and thus benefit from their data.

Figure 1: Key considerations

CUSTOMER IMPACT

Consent/privacy
How does the organization gain individuals' consent? What are individuals' expectations about how their data will be used?

Autonomy
To what extent are individuals given a say in how their data is collected, stored, and used? What are the ethical implications of this?

TRUSTWORTHINESS

AI solutions
What are the ethical implications of using off-the-shelf solutions for analytics activities? What is the impact of these solutions in terms of transparency, fairness, and explainability?

Data-sharing
What are individuals' expectations regarding data-sharing? What are the ethical implications of sharing individuals' data with other parties?

Fairness
How fair are the AI models used by the organization? How can biases be mitigated?

Vendor services
What are individuals' expectations regarding outsourcing data-related services to third parties?

Vendor data
What are the ethical implications of using purchased third-party data? What is the impact of third-party data on consent, transparency, and bias in AI techniques?

Explainability
Are individuals' expectations regarding the transparency and explainability of AI use cases being met?

FOUNDATIONAL ENABLERS

Data acquisition
What are the ethical implications regarding individuals' data privacy? What constitutes "ethical use"?

Data security
Is acquired/shared data managed, stored, and secured in compliance with applicable regulations and individuals' expectations?

Data quality
How does data quality and completeness impact business decisions made using AI techniques?

Data availability and representativeness
What are the ethical implications of using partial/incomplete shared data in AI use cases?

The regulatory landscape

In response to growing concerns about organizations' collection and use of personal data and the ethical implications of using AI and related technologies, regulators around the world have introduced new rules and legislation, and continue to work on new guidance.

Organizations should strive to understand and adhere to both the spirit and the letter of these regulations as they implement or strengthen data ethics programs, collect and use individuals' data, and deploy AI, machine learning, and similar technologies. This includes, for example, providing customers with clear explanations about what data is collected and how it will be used and establishing appropriate oversight and controls to ensure that automated decision systems (ADS) work to mitigate bias and similar risks.

No single piece of legislation, regulation, or guidance will fully address the many issues that arise from the use of personal data and AI technologies, of course. Organizations will need to navigate a matrix of regulations that target separate but overlapping areas (e.g., privacy, competition, and consumer protection) and operate in different national, provincial/state, and even municipal jurisdictions. While it can be challenging to stay on top of the ever-evolving regulatory landscape, it's essential if organizations are to avoid penalties for non-compliance. There are also other benefits: organizations that are aware of regulators' areas of inquiry or concern can proactively strengthen controls, policies, and processes ahead of legislative action.

Figure 2

Jurisdiction	Enacted	Emerging
Canada (federal)	The Personal Information Protection and Electronic Documents Act (PIPEDA) governs how private sector organizations collect, use, and disclose personal information.	• Bill C-26, an Act Respecting Cyber Security (ARCS), targets federally regulated businesses providing critical infrastructure services in the finance, telecommunications, energy, and transportation sectors. If enacted, the law would require those businesses to prepare for, address, and report cyberattacks as well as authorize the government to order businesses to respond to anything deemed a threat to national cybersecurity interests without publicly disclosing the order. Violations of the Critical Cyber Systems Protection Act, included in Bill C-26, can result in penalties of up to $15 million. • The government introduced Bill C-27, an expanded, updated version of the Digital Charter Implementation Act (DCIA). The three-part law includes the Consumer Privacy Protection Act (CPPA), a new private sector privacy law modernizing personal information collection; a new tribunal under the Personal Information and Data Protection Tribunal Act (PIDPTA); and the Artificial Intelligence and Data Act (AIDA), regulating international and interprovincial trade and commerce in AI systems. Severe violations of the CPPA could result in penalties of up to $25 million or 5% of global gross revenues, whichever is higher.

Jurisdiction	Enacted	Emerging
		• The government has pledged to establish a transparent, accountable regulatory framework for online safety and committed to review competition laws to ensure they reflect the realities and challenges of the digital economy and address emerging and potentially harmful business behaviours.[4] • The Office of the Privacy Commissioner of Canada released key recommendations for regulating AI, drawing attention to the risks entailed in automated decision-making.[5]
Canada (provincial)	Bill 64 in Quebec introduces new obligations for organizations doing business in the province, including the need to establish and implement a privacy framework and a requirement to provide individuals with much more transparency about how their personal data would be used. The legislation also introduces new rules regarding consent and new individual rights pertaining to automated decision-making that involves personal information. Violations can incur penalties of up to $25 million or 4% of global annual revenue, whichever is greater.	• The BC and Yukon information and privacy commissioners issued a joint report in 2021 to raise ethical and legal concerns regarding governments' use of AI and highlight how failing to address issues of fairness and privacy could damage public trust in governments. The commissioners offered a number of recommendations, including using synthetic or de-identified data in ADS whenever possible; notifying individuals when ADS is used to make decisions and explaining why and how ADS is used; and providing individuals with a means to object to the use of ADS.[6]
European Union	The General Data Protection Regulation (GDPR) affects organizations that collect and process personal data of individuals located in the EU. The ultimate aim of GDPR is to strengthen and unify data protection for EU citizens. GDPR violations can incur penalties of up to 4% of worldwide annual revenue.	• The EU Artificial Intelligence Act, anticipated to enter into force in 2023 or 2024, emphasizes the ethical application of AI. Under the Act, some AI applications are prohibited owing to unacceptable risk, including social scoring applications and the manipulation of human behaviours, opinions, and decisions. Other AI applications deemed to pose high or moderate risk are permitted, subject to compliance with certain requirements or specific transparency obligations. • Europe's Digital Markets Act will, among other things, ensure digital companies can only combine personal data to facilitate targeted advertising with an individual's express consent.[7]
United Kingdom		• The UK government introduced the Data Protection and Digital Information Bill in mid-2022 to encourage innovation while promoting responsible AI use. The government also published its proposed rules for regulating AI in the United Kingdom as part of its National AI Strategy, which aims to help organizations harness AI's potential while addressing the challenges AI use presents.[8]
United States	The California Consumer Privacy Act (CCPA) of 2018 created new consumer privacy rights and business obligations with respect to the collection and sale of personal information. The California Privacy Rights Act (CPRA), passed in 2020, amends and expands the CCPA; it adds rights to limit the use and disclosure of sensitive personal information, to opt out of both the sale and sharing of personal information, and to correct inaccuracies.	• A growing number of states are introducing measures to study the impact of AI usage and identify the role policymakers should play. In 2021, at least 17 states enacted general AI bills or resolutions.[9] • As of January 2, 2023, New York City will ban employers in the city from using automated employment decision tools to screen job candidates, unless the technology has been subject to a bias audit within the year prior to its use.[10]

Figure 3: Data ethics framework

What actions do we prioritize?

Data life cycle process: Assess principles/risks in the context of each stage of the data life cycle. Understand the breadth of data ethics considerations.

Who is accountable?

Accountability model: Identify appropriate teams from across the organization to be responsible for assessing the relevant and appropriate actions to address data ethics.

How do we operationalize actions?

Measurement and monitoring: Document ethical considerations to understand adoption issues and aggregate as possible/appropriate.

Processes and controls: Ensure an understanding of the existing mechanisms (processes) that can be used to embed the data ethics actions.

Tools and techniques: Identify and leverage tools and techniques to systematize a set of relevant data ethics considerations.

How do we accelerate adoption?

Communications, training, and change management: Ensure customers, employees, shareholders, and other stakeholders are kept informed of the organization's perspectives and actions as they relate to data ethics. Upskilling and training are key components to embed from the start.

Making data ethics real

Every organization that depends on individuals' data to innovate, grow, and prosper must take data ethics seriously to earn and keep customers' trust, stay onside with regulators, avoid potential financial and reputational harm, and stand out in the market. But how should an organization go about implementing or upgrading its data ethics program?

In this section, we offer our perspectives on how to move forward. However, it's important that organizations recognize they are not starting from scratch when it comes to AI governance. It's about augmenting the work they've already done in the areas of data governance and management, data privacy, data security, and more. This provides a solid base for a strengthened commitment to responsible AI. To build on that foundation, organizations can:

Establish responsibility and accountability

Make data ethics a strategic priority for the entire enterprise. To take a truly holistic approach, it's important to bring together multiple perspectives and disciplines: data, legal, privacy, business lines, marketing, the corner office, and more. Gathering all the key stakeholders into a working group, committee, or other body will help ensure the organization addresses data ethics from all angles.

In most cases, data ethics falls within the purview of the chief data and analytics officer (CDAO), reflecting both data's increasingly critical role to the organization and the evolution of the CDAO role itself, which has historically focused on regulatory reporting and compliance. Today, it also involves supporting data-driven decision-making, innovation, and monetization across the enterprise—activities that have the potential for ethical missteps if not appropriately governed.

But other leaders also have vital and ongoing parts to play in making data ethics come to life within the organization. Chief privacy officers and chief legal officers are key to looking at AI governance from a defensive perspective, ensuring the organization stays on side of existing and emerging legislation and other rules. Business line leaders—who typically use data-driven insights to drive revenue and create value for the organization—can approach data ethics from an offensive perspective, leveraging responsible AI to deepen customers' trust and strengthen relationships.

Standard-setters and other groups weigh in

- The Data Management Association (DAMA International) has updated its Data Management Book of Knowledge (DMBOK) 2.0 to include data ethics and address how to incorporate ethical principles into data procurement, storage, management, use, and disposal.[11]
- In 2019, the Institute of Electrical and Electronics Engineers (IEEE) launched a standards project to address the ethical issues of designing autonomous and intelligent systems. In 2020, it introduced two complementary frameworks for integrating ethical analysis into engineering practice.[12] In 2021, it unveiled a methodology that would support ethical systems engineering.[13]
- The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) issued a technical report to analyze factors that impact the trustworthiness of systems providing or using AI.[14] ISO 26000 guidance on social responsibility is also being used as a lens through which to examine the ethical risks of AI use and identify potential countermeasures.[15]
- In the United States, the National Institute of Standards and Technology (NIST) released a first draft of a privacy framework that sets out an ethical foundation for data usage in 2020.[16] The organization also proposed guidelines for identifying and managing bias in the use of AI in 2021.[17]
- Information Systems Audit and Control Association (ISACA) has put forward its well-known governance-of-enterprise framework, COBIT, as a useful enabler for organizations aligning to the GDPR.[18]
- Canada's CIO Strategy Council launched its AI Ethics Assurance Program in 2020. The program is designed to help organizations obtain independent assurance on whether their controls meet established criteria with respect to the ethical design and use of automated decisioning systems (ADS).
- Work is also being done to develop AI management system standards to enable organizations to demonstrate they've implemented, and continually strive to improve, processes specific to AI development and usage. Management system standards have been deployed in various industries to facilitate innovation and technology development through well-structured governance and careful risk management.
- A proposed global AI certification program is in development that reflects Organisation for Economic Co-operation and Development (OECD) AI principles and various international AI standards and guidance. The proposed program would assess the overall maturity of an AI system's data, model, and deployment and recognize it with a tiered certification level.

Embed it into every stage of the data life cycle

An organization's data ethics framework is executed throughout the data life cycle: acquisition/generation, management, aggregation and analysis, use, sharing, and archiving and disposal. Each stage presents opportunities to strengthen an organization's commitment to AI governance. The organization should begin by assessing how data ethics is currently embedded into each life cycle stage. Depending on the results of this assessment, the organization should then identify the practical, concrete actions that can be taken to implement or improve data ethics considerations. These actions may involve creating new policies or standards for explainability, or crafting explanations for automated decision-making or consent communications for customers. The organization may decide to update processes, develop a new control, adjust analytical models, or even rethink the kinds of customer data collected from the outset.

Establish a group to challenge your data use

Mature organizations typically set up a group responsible for independently challenging the use of data and evaluating the soundness of analytical models used to make business decisions. Traditionally, these groups focused on material or monetary impacts, but they are increasingly considering issues such as fairness, bias, transparency, explainability, and privacy.

Tailor the approach to your organization’s needs

It’s important that organizations address responsible AI in a way that makes sense for their business. Accordingly, organizations should pursue a fit-for-purpose approach that focuses on those aspects of data ethics that are especially pertinent and relevant. A tailored, balanced approach to AI governance will enable them to balance effective data governance with data-driven innovation.

Case study: Scotiabank strengthens enterprise-wide data ethics capabilities

In 2020, Scotiabank approached Deloitte to collaborate on designing and implementing a holistic, enterprise-wide data ethics framework that would enable the bank to ensure the appropriate ethical considerations were being applied at every stage of the data life cycle. This effort would involve carefully reviewing and identifying key processes and procedures across the data life cycle—including those related to artificial intelligence and machine-learning model development—where data ethics could be embedded to enable robust, timely coverage and ease of use for practitioners.

“Including ethical considerations early in the design stage of a use case helps us take a proactive approach throughout the development process," says Anna Hannem, director of Data Ethics & Use. "The result is more ethical, trustworthy use cases, which benefit our customers and our businesses.”

The project involved consultations with a broad set of stakeholders representing more than 15 internal bank functions and business lines. The team assessed the current processes at the bank and outlined a set of key actions to be taken to operationalize the application of data ethics across each stage of the bank’s data life cycle. Next, the team developed enhancements to prioritized processes as well as new artifacts to embed ethical considerations in areas that warranted particularly stringent considerations when collecting, managing, or using data—including in artificial intelligence or machine-learning use cases. Finally, a data ethics governance operating model was developed and established to measure the bank’s data ethics processes, and clarify the role that each individual plays in data ethics.

As part of this process, Deloitte’s EthiKit was adapted to fit within the bank’s own technology environment and is now the basis of Scotiabank’s Ethics Assistant. The tool was piloted, tested by model developers, iterated based on user feedback, and ultimately deployed and operationalized.

“The tool provides great insights to data scientists about the potential bias and risks associated with their model development, highlighting opportunities to apply treatments to mitigate risk before deploying the models in production,” explains Tanaby Zibamanzar Mofrad, director of Data Science and Analytics, Global Wealth Analytics, at Scotiabank.

The project has enabled Scotiabank to formalize its application of data ethics across applicable enterprise policies, frameworks, and use cases. While practitioners still have a personal obligation to build in an ethical manner, this tool ensures they have support in place to apply ethical considerations more consistently. Ethical assessments and guidance are now automated and codified in the artificial intelligence and machine-learning development processes. This also accelerates use case delivery, as data ethics risk can be proactively identified and addressed early in the development and design life cycle. Finally, stakeholders and executives are aligned around operational responsibilities for prioritized data ethics processes and support the ongoing promotion, adoption, and standardization of data ethics across the enterprise.

"This new solution gives us the confidence that the enhancements we are making to develop personalized solutions for our customers are aligned with and strengthen our data ethics practices," says Grace Lee, senior vice president and chief data and analytics officer.

When using AI, consider the full range of risks

AI, machine learning, and related technologies currently represent a small proportion of organizations' overall data use, but the deployment of such tools to process and generate data insights is growing rapidly. As a result, organizations will need to update their existing risk management mechanisms and processes to reflect the new risks introduced by AI use. There's a wide range of risk areas to be considered, including fairness/ impartiality, robustness/ reliability, privacy, safety/security, responsibility/accountability, acceptable use, and third-party liability.

Leverage the work of standard-setters and similar groups

Organizations can also capitalize on the work being done by standard-setting bodies and other groups developing best practices to address issues of data ethics, responsible AI use, and similar targets. In recent years, for example, groups such as DAMA, ISO, NIST, and ISACA have sharpened their perspectives on data ethics and related matters and are working on developing frameworks and standards to help other organizations operationalize data ethics themselves.

The imperative for the public sector

Public sector organizations have access to some of the most sensitive and important data that individuals possess. As these organizations strive to use this and other data to serve citizens more nimbly and responsively, they must ensure that citizens believe they will do so in an ethical, trustworthy way.

They may wish to use a risk-based approach to implementing a data ethics framework, using the robust measures already in place in the data life cycle (e.g., data disposal rules and processes) to focus resources on addressing ethical issues elsewhere in the life cycle. Mature organizations are already searching for ways to streamline the application of data ethics policies and processes wherever feasible.

Understanding how citizens' data is being—or will be—used is critical, as is ensuring that these use cases meet ethical standards. This is particularly important when the intended data use changes for legitimate reasons, or when the intended data use was poorly explained at the outset. As well, organizations should strive to maintain a regular dialogue with other public sector entities to stay abreast of what others are doing to make data ethics a reality in a similar context.

Looking ahead to the future

Today, much of the discussion and effort regarding AI governance is understandably focused on issues surrounding data collection, data analysis, and data management. Many organizations are just beginning to grapple with questions of data ethics in an environment of growing regulatory scrutiny and public concern.

But the field of data ethics will continue to shift and evolve as citizens, organizations, governments, and regulators contend with technological advances, business innovations, increasing connectivity, and an ever-expanding volume of data. Understanding these changes is vital if boards, C-suite executives, CDAOs, and others are to ensure their organizations respond nimbly and proactively to retain customers' trust, secure or develop competitive advantages, and position themselves for future success.

Below are a number of developments that are likely to have an important impact on the responsible AI conversation in Canada in the near term.

Customer-driven consent

As organizations grow more connected to customers, vendors, and other third parties, the amount of personal data being collected, stored, used, and shared has greatly increased—as has the risk of that data being disclosed or misused, often without individuals' knowledge or consent. In response, regulators and companies are embracing the idea of customer-driven consent.

Customer-driven consent is an approach that promotes trust between organizations and individuals by providing the individual with control over what data they provide and how it is used. It requires companies to clearly and transparently explain why they wish to collect personal data, what that data will be used for, and with whom that data would be shared. And it allows the individual to decide for themselves what data they will provide and deliberately opt in to—and out of—particular data uses or exchanges.

For organizations, customer-driven consent means understanding what customer data is collected, how that data is used and shared, and perhaps most importantly, why. Steps must be taken to identify privacy issues that may arise and to ensure customer data isn't being used or shared for unintended purposes. Organizations have to clearly explain what data they want to collect and how and why they intend to use it—and, crucially, provide customers with the means to provide or withdraw consent to that data collection and usage. This matter of consent extends to externally sourced data as well: organizations will need to confirm that third-party data providers have also sought and obtained the requisite informed consent.

Digital IDs and credentials

Governments and the private sector are both looking at digital IDs—electronic versions of trusted government identification documents—to provide a safer, more secure, more private, and more convenient alternative to traditional ID documentation. Digital IDs are designed to be stored in digital wallet applications on smartphones, tablets, and computers, enabling people to prove who they are whether online or in person. While the technology remains shrouded in misinformation, properly designed digital IDs are actually privacy-enabling because only the data required for the intended exchange is presented. Proof-of-age data exchanges, for example, need only provide a green check mark, rather than exposing the individual's name, address, date of birth, and other personal details.

The Government of Ontario announced its digital ID ambitions in late 2020. Since then, the province has shared its technology road map and technical tools and held a series of consultations with citizens and businesses. While most digital IDs would be provided by each province independently, Canadians had their first real digital ID experiences with the Canadian COVID-19 proof of vaccination during the pandemic. In addition, Ontario has begun to trial digital ID tools in a health care setting: Niagara Health patients can access their diagnostic records (e.g., x-rays and MRIs) through a digital identity service for patients, while three Ontario hospitals have introduced a service that allows patients to prove their identity online.

Open banking

Open banking allows individuals to share their banking and financial data with third parties—commonly fintech companies—without providing the third party with their online banking usernames or passwords. On an open banking platform, the bank uses APIs to securely share the individual's data with the fintech or other third party on their behalf. While it has not yet been introduced in Canada, open banking is already used in countries such as the United Kingdom and Australia.

Financial services organizations will need to take the lead in driving the effective integration of AI governance into any open banking strategy. To ensure that customers' financial data is shared with third parties properly and in keeping with data ethics standards, they will need to establish a privacy and security structure for all enterprise data and identify and secure all data categories appropriately.

Marketing and hyper-personalization

In every industry, marketing has always sought to build a stronger connection between brands and customers. Today, businesses are expected to do more than simply meet customers' needs: they must anticipate and exceed them. As marketing grows ever more customer-centric and data- driven, traditional customer segmentation is fast giving way to AI- powered hyper-personalization.

Hyper-personalization is the most advanced way for brands to tailor their marketing to individual customers. It uses data, analytics, AI, and automation to create custom, targeted experiences and send highly contextualized communications to specific customers at the right place and time through the right channel. Hyper-personalized marketing creates opportunities for organizations to meaningfully engage customers, deepen existing relationships and build new ones, and improve the overall customer experience.

Brands and retailers aren't the only ones looking at using data and AI tools to hyper-personalize their services and messaging. Public sector organizations can use AI-powered personalization to improve outcomes by better understanding the needs of citizens and delivering the right services and information to them at the right time. Health care organizations are looking at how AI can be used to ensure each individual patient receives the specific treatment they require. Obviously, such use cases involve incredibly personal data about individuals, making it more important than ever for organizations to get responsible AI right.

Case study: A provincial government drives forward with responsible AI

One of Canada's provincial governments aims to lead the country in applying AI responsibly to both improve citizen outcomes and foster economic development. Deloitte worked with the province to develop one of the public sector's most comprehensive AI strategies, testing the strategy using real-world applications.

Dialogue and collaboration were essential to the project. We worked with stakeholders drawn from 23 ministries across the provincial government to align on overall priorities, design the responsible AI strategy, and create an operating model that incorporated our EthitKit toolkit. In developing the strategy, operating model, and framework, the team drew inspiration from other governments and global profit and non-profit organizations as well as applied lessons from a global scan of responsible AI risk areas.

To test and evaluate the strategy and framework, five proof-of-concept use cases were built using agile methodologies to demonstrate the ethical use of AI technologies within a government context, addressing issues such as wildfire resource planning and flood monitoring. The use cases identified $3 million in cost savings—and the lessons learned from them enabled the team to refine and improve the responsible AI strategy.

In addition, over the course of the project, more than 100 government stakeholders participated in a Deloitte AI Academy, which was designed to equip them to identify and support AI projects in their own areas of expertise and influence.

EthiKit by Deloitte

Accelerate your data ethics journey

Achieve your data ethics goals faster with EthiKit—the trusted toolkit and etiquette for responsibly managing and using data.

It's a powerful asset for any organization that wants to build a best-in-class data ethics program. Whether you're launching your first official data ethics program or upgrading what's already in place, EthiKit gives you the tools, guidance, and best practices you need to get it done.

The toolkit enables organizations to govern and manage AI risk effectively and efficiently throughout the AI development life cycle. It assesses data-related risks and provides practical guidance to AI teams, risk managers, and other stakeholders so you can manage and mitigate them. And because it comes pre-aligned with relevant Canadian and international regulations, guidance, and best practices, EthiKit keeps your organization onside in the ever-evolving regulatory landscape.

EthiKit helps you answer the following:

- What is the AI risk profile of my line of business? Of the enterprise?
- What is the velocity of my AI system development cycle and deployment success ratio? Where are the bottlenecks?
- Are risk levels being signed off prior to deployment?
- Which AI systems in development are showing multiple high-risk indicators?
- What are the risks that require the most attention and what can be done to mitigate them?

KEY FEATURES

Tailored data ethics focused on what's in your organization's control, based on your responses to a dynamic questionnaire that covers the AI risk spectrum.

Real-time, action-oriented guidance ensures you develop responsible, ethical AI systems. EthiKit outlines key considerations, identifies decisions that need to be made, and indicates who to consult and which tools or processes to use.

Developer playbooks for fairness and explainability provide AI developers with tactical, situation-specific, and model-agnostic guidance and give model developers the math, code sets, and other assistance needed to refine AI models.

Clear data visualization dashboards collect and display enterprise-wide insights in a way that's easy to understand and act on.

An embedded workflow engine facilitates real-time decisions and approvals to keep you moving forward.

Incorporates the latest regulatory guidance from Canada, the European Union, the United States, the OECD, the UK Financial Conduct Authority, and more.

Conclusion

As organizations collect and use ever-increasing amounts of data to innovate, drive growth, improve customer relationships, and more, the risk that this data could be misused also grows. And the consequences of unethical data use can be significant, both financially and reputationally.

If you haven't yet, now is the time to start your organization's data ethics journey. A commitment to ethical data use will quickly become essential to earning and retaining trust and to staying competitive, no matter the industry. In short, embracing data ethics is a crucial way of future-proofing your business.

Where to begin? Start by asking these questions:

Where are we in our data ethics journey?

- Are we looking to scale our use of AI and analytics?
- How do we ensure trust when using customers' data?
- Do we know which current or emerging regulations will impact our business?

What is our desired future?

- How will we ensure consistency across governance and decision- making?
- How will we ensure customer consent?
- How will we identify inherent bias?
- How will we ensure our decision-making processes are transparent and explainable?

How will we get from this present to that future?

- How do we identify risk when using customer data?
- What resources and tools do we need in place to support consistency?
- What skills do our AI teams need to eliminate bias, ensure transparency, and use data responsibly?

The answers to these questions will give you a sense of where your organization needs to go and how to get there.

Endnotes

1. Cat Zakrzewski, “Twitter to pay $150 million fine over deceptively collected data,” Washington Post, May 25, 2022.

2. Richard Nieva, “Google sued by New Mexico AG for allegedly collecting location data, contact lists from students,” CNET, Feb. 20, 2020.

3. Olivia Stefanovich, “Supreme Court sides with Uber drivers, opening door to $400M class-action lawsuit,” CBC News, Jun 26, 2020.

4. Canadian Heritage, “The Government’s commitment to address online safety,” Government of Canada, September 2022; Innovation, Science and Economic Development Canada, “Minister Champagne maintains the Competition Act’s merger notification threshold to support a dynamic, fair and resilient economy,” Government of Canada, February 7, 2022.

5. Professor Ignacio Cofone, ”Policy Proposals for PIPEDA Reform to Address Artificial Intelligence Report,” Office of the Privacy Commissioner of Canada, November 2020.

6. British Columbia Office of the Ombudsperson, British Columbia Office of the Information and Privacy Commissioner, and Office of the Yukon Ombudsman and Information and Privacy Commissioner, ”Getting ahead of the curve: Meeting the challenges to privacy and fairness arising from the use of artificial intelligence in the public sector,” Joint Special Report No. 2, June 2021.

7. Committee on the Internal Market and Consumer Protection (IMCO), “Deal on Digital Markets Act: EU rules to ensure fair competition and more choice for users,” European Parliament News, March 24, 2022.

8. Department for Digital, Culture, Media & Sport and Damian Collins MP, “UK sets out proposals for new AI rulebook to unleash innovation and boost public trust in the technology,” GOV.UK, July 18, 2022.

9. National Conference of State Legislatures, “Legislation Related to Artificial Intelligence,” NCSL, August 26, 2022.

10. Erin Mulvaney, “NYC Targets Artificial Intelligence Bias in Hiring Under New Law,” Bloomberg Law, December 10, 2021.

11. DAMA New England, “CDMP Study Group Session 3: Data Handling Ethics,” Data Management Association, March 4, 2020.

12. D. Peters, K. Vold, D. Robinson, and R. A. Calvo, “Responsible AI—Two Frameworks for Ethical Design Practice,” IEEE Transactions on Technology and Society, vol. 1, no. 1, March 2020.

13. IEEE, “IEEE Launches New Standard to Address Ethical Concerns During Systems Design,” IEEE Standards Association, September 15, 2021.

14. Elizabeth Gasiorowski-Denis, “Towards a trustworthy AI,” ISO News, July 7, 2020.

15. Weiwei Zhao, “Artificial Intelligence and ISO 26000 (Guidance on Social Responsibility),” IntechOpen, February 17, 2021.

16. Jory Heckman, “NIST privacy framework looks to underpin ethical use of artificial intelligence,” Federal News Network, February 19, 2020.

17. IEEE, “NIST: Recommendations for Identifying and Managing Bias in AI,” IEEE Innovation at Work, July 21, 2021

18. Joanna Karczewska, “COBIT 5 and the GDPR,” ISACA GDPR Working Group, May 24, 2017.

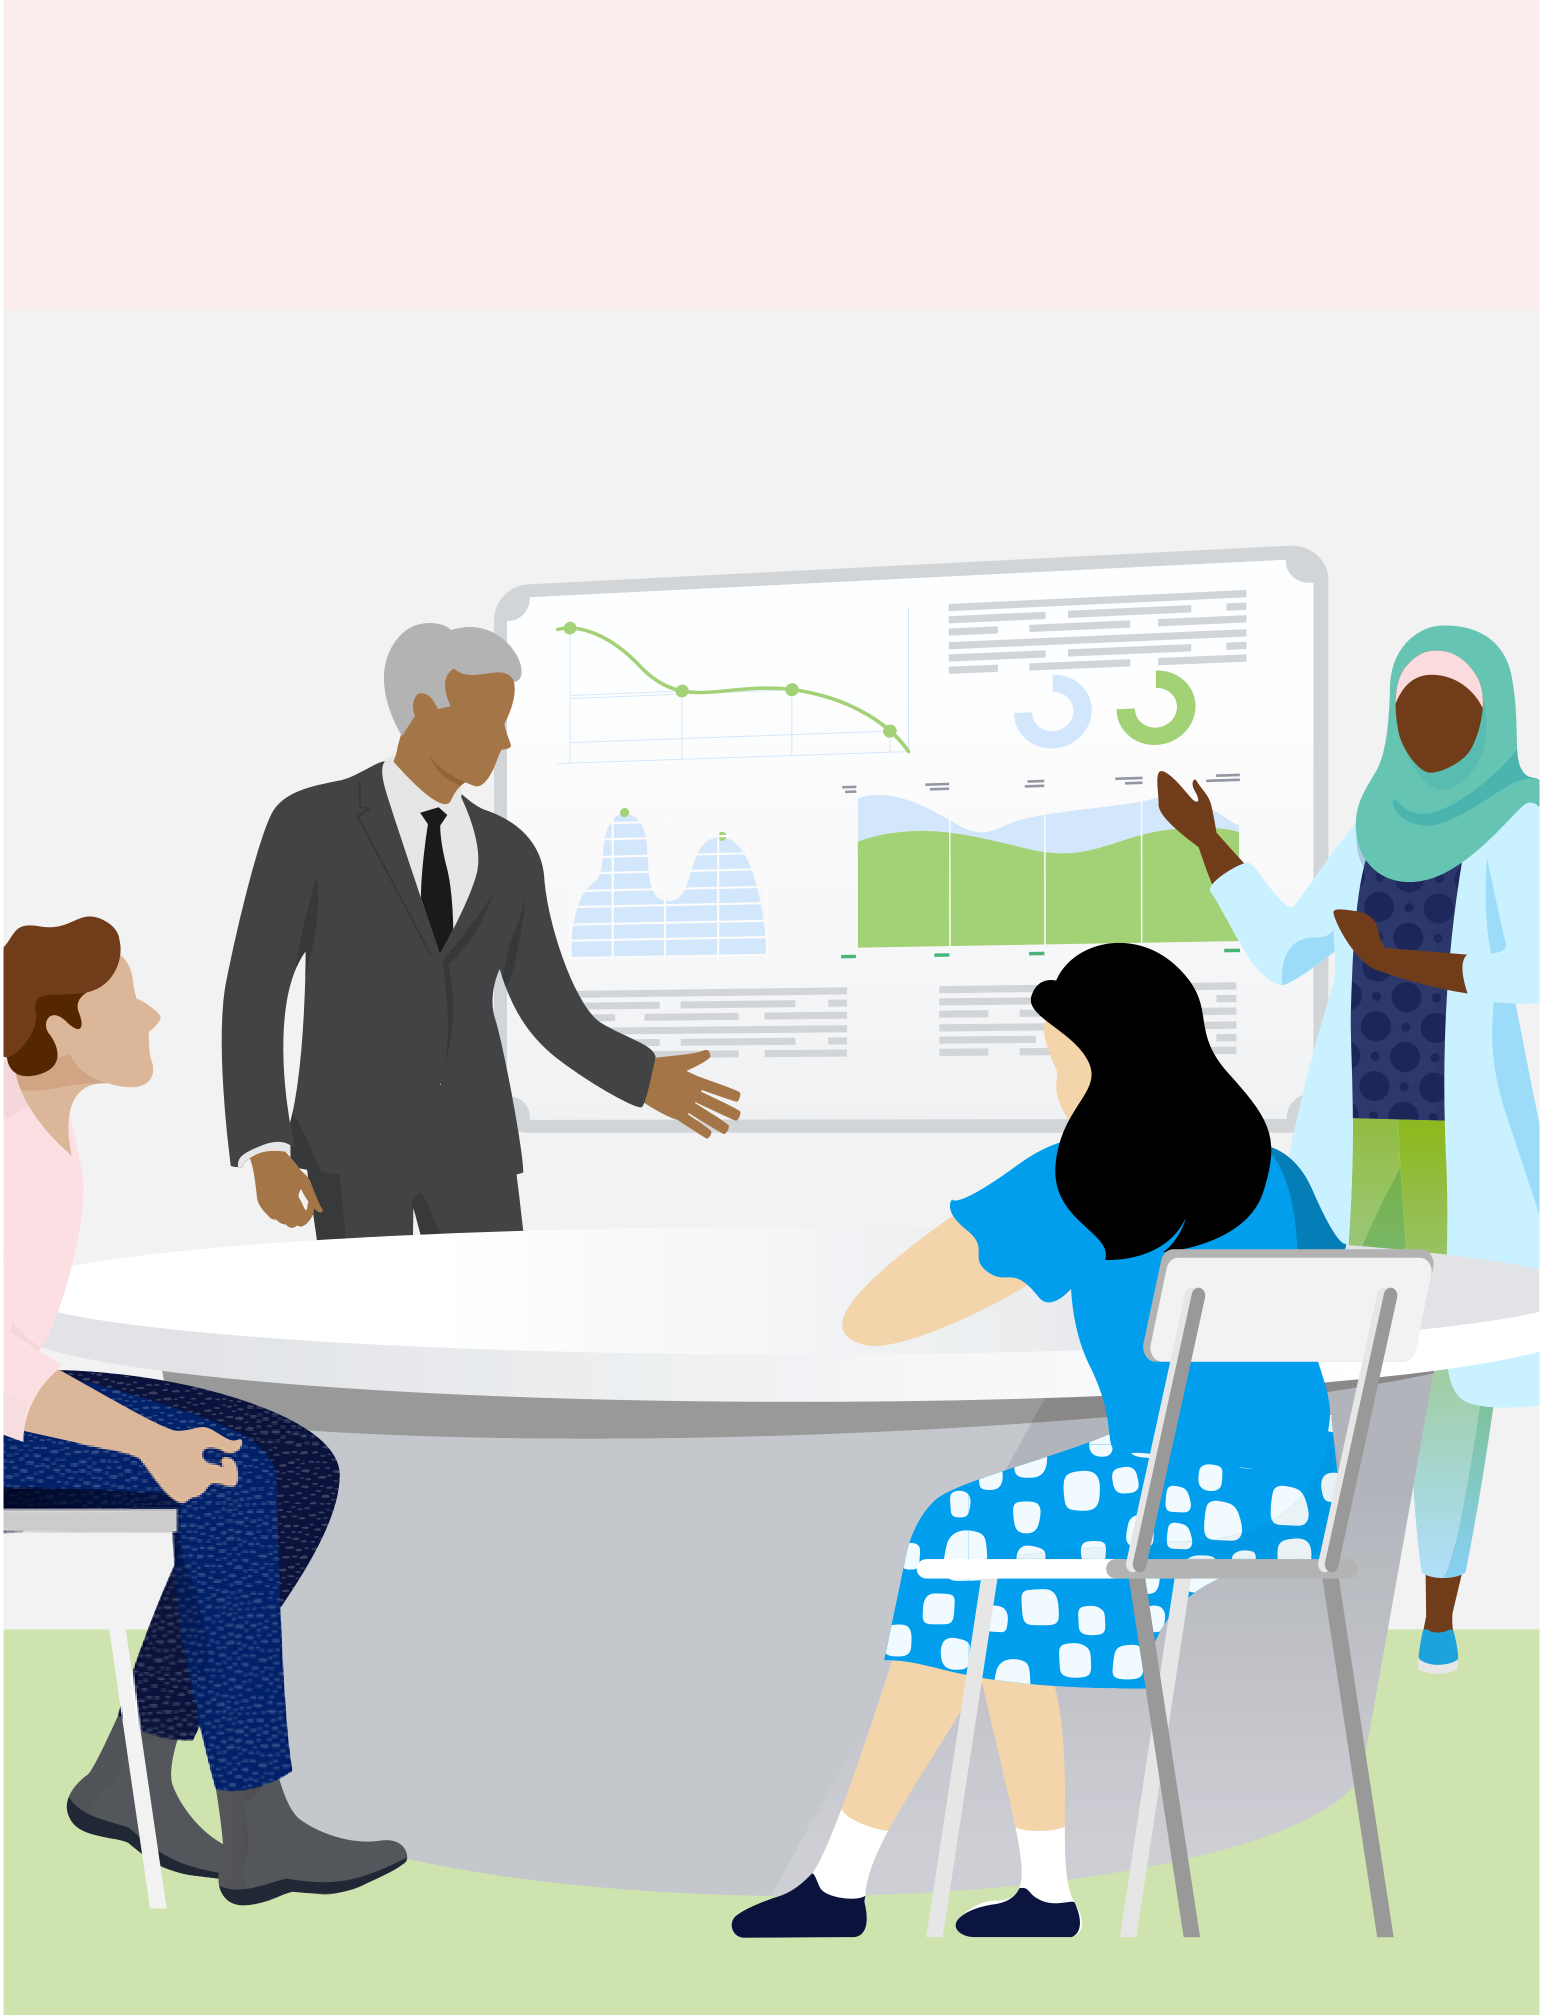

Contacts

To learn more about how your organization can become secure, vigilant, and resilient, please contact:

Preeti Shivpuri
Trustworthy AI Leader
pshivpuri@deloitte.ca

Mukul Ahuja
AI & Data Strategy Leader
mukulahuja@deloitte.ca

Adnaan Sikandar
Data Governance & Architecture Leader
adsikandar@deloitte.ca

Acknowledgements

Special thanks to those who participated in the interviews and provided support for this report.

Mona Balesh Abadi
Senior Manager, Data Ethics
Scotiabank

Anna Hannem
Director, Data Ethics
Scotiabank

Charlene Viscek
Director, Global Risk Management Communications
Scotiabank

Tanaby Zibamanzar Mofrad
Director, Data Science & Analytics, Global Wealth Analytics
Scotiabank

Grace Lee
Senior Vice President, Chief Data & Analytics Officer
Scotiabank

We wish to thank the following people from Deloitte Canada for their contributions to this report.

Noel Corriveau
Manager, Omnia AI

Chetan Phull
Associate Lawyer, Tax & Legal

Hwan Kim
Partner, Financial Services Strategy & Innovation

Giselle D'Paiva
Partner, Government & Public Sector Digital ID Leader

Sime Pavlovic
Senior Advisor, Consulting

Noemi Chanda
Partner, Data Privacy Leader

Deloitte.

www.deloitte.ca
This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

About Deloitte
Deloitte provides audit and assurance, consulting, financial advisory, risk advisory, tax, and related services to public and private clients spanning multiple industries. Deloitte serves four out of five Fortune Global 500® companies through a globally connected network of member firms in more than 150 countries and territories bringing world-class capabilities, insights, and service to address clients' most complex business challenges.

Deloitte LLP, an Ontario limited liability partnership, is the Canadian member firm of Deloitte Touche Tohmatsu Limited. Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please see www.deloitte.com/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms.

Our global Purpose is making an impact that matters. At Deloitte Canada, that translates into building a better future by accelerating and expanding access to knowledge. We believe we can achieve this Purpose by living our Shared Values to lead the way, serve with integrity, take care of each other, foster inclusion, and collaborate for measurable impact.

To learn more about Deloitte's approximately 330,000 professionals, over 11,000 of whom are part of the Canadian firm, please connect with us on LinkedIn, Twitter, Instagram, or Facebook.