THE Information Accountability Foundation
COLLABORATIVE POLICY INNOVATION

Trusted Digital Transformation

Considerations for Canadian Public Policy

January 2019

Introduction

Canada's future is digital, and the legal infrastructure needs to support that future. Navdeep Bains, Minister of Innovation, Science and Economic Development, said it well:

> [W]e live in a world where data are the raw materials that drive the innovation economy. We not only live in a world with more data. We also live in a world where more data come from a wider variety of sources. . . . In fact, data are being collected and analyzed at a speed that is rapidly approaching real time. That means the time between what we know and when we act on what we know is getting shorter. . . .
>
> [O]pportunities [exist] to innovate, serve their customers better and create entirely new jobs and industries that never existed before. In short, big data analytics has the potential to affect the lives of Canadians more quickly and directly than ever before. Technology now allows information to be captured, copied, shared and transferred quickly and endlessly. . . . That means just about every business, regardless of what it sells, is now a data and software company. . . .
>
> Powerful software can now extract insights from small, seemingly disconnected pieces of data. Digital technologies that enable the continuous collection of large datasets have the potential to make companies innovative and valuable. But large-scale data collection also means that governments, businesses and citizens must continually review privacy and security policies and practices. Our government's goal is to encourage the free flow of data to spark innovation. But government also has a responsibility to protect the privacy of citizens, promote fairness, foster equality of opportunity for all Canadians and make itself more open and accountable.[1]

Canada has been served well by the Personal Information Protection and Electronic Documents Act ("PIPDEA")[2] that went into effect in 2000. PIPEDA's purpose is "to establish, in an era in which technology increasingly facilitates the circulation and exchange of information in a manner that recognizes the right of privacy of individuals with respect to their personal information and the need of organizations to collect, use or disclose personal information for purposes that a reasonable person would consider appropriate in the circumstances."[3] While having a foundation of ten principles, PIPEDA is heavily dependent on two pillars: consent and accountability. Consent is the tool that affords individuals the opportunity to exercise control over their personal information, and accountability requires organizations to develop and implement policies and practices to uphold the fair information principles set forth in PIPEDA.[4]

---

[1] Speech by the Honourable Navdeep Bains, PC, MP, Minister of Innovation, Science and Economic Development, Toronto, Ontario, July 28, 2017

[2] S.C. 2000, c. 5

[3] PIPEDA, Purpose, Part 1(3)

[4] Office of the Privacy Commissioner Report on Consent [hereinafter Report on Consent] at 1, 16

PIPEDA was based on a consensus standard ratified in 1996 that became PIPEDA's fair information principles. However much has happened in the digital world since 1996:

- Acceleration of the Internet
- Y2K common processing platforms that created the base for cloud computing
- 2007 introduction of smart phones
- Big data
- Internet of Things, and
- Artificial Intelligence

Never in the world's long history has such technology acceleration taken it, including Canada, into another industrial age, this time the fourth.

This paper is intended as a discussion starter in Canada, nothing more. It is authored by the Information Accountability Foundation (IAF), a global think tank. The IAF has conducted similar work in Europe, Asia and Latin America. The paper will set out the challenge, establish the concept of enhanced accountability, propose principles, including more effective individual engagement, and make high level recommendations for the policy discussion that are supported by the paper. Those recommendations are:

1. Data pertaining to individuals should be used to create real value for identified stakeholders in a balanced, fair fashion that serves individuals, society and private organizations. This balancing should take place in all sectors, and the risk of not using data should be as important a consideration as the risk of negative consequences.
2. Individuals have clear rights related to data and its uses, and those rights should be explicit and actionable as well as theory.
3. Accountability requires organizations to be reasonable and responsible in what they do with data pertaining to individuals and to be answerable for how they demonstrate that they are acting as effective data stewards.
4. Organizations should have checks and balances in place, so their data stewardship is conducted effectively. When organizations cause negative consequences that are consequential, they should take actions to mitigate those consequences.
5. Enforcement agencies should have the powers and resources so that they may act in a manner trusted by the public and seen as predictable by those subject to enforcement.

Canada has been and still is the paragon for a principles-based approach to privacy. The question is how should principles be applied in this digital transformation? This paper's intent is to set a discussion path that might be as fruitful as the one in the 1990's that created PIPEDA.

The Impact of the Digital Economy

Digital transformation is changing the way Canadians work, play, share, shop and even the way they may choose to experience their world. As the digital economy becomes imperative for productivity and growth, Canada's digital innovators will underpin its future prosperity. Today's Canadian digital industries include a combination of information and communication

technology (ICT), digital and interactive media, content industries, and manufacturers and service companies that use creativity, talent and digital skills to capture, transmit and display data and information electronically in ever evolving ways. Increasing domestic uptake of digital innovation may be the single most important element to improving productivity. A one percent increase in digital technology adoption could generate $2.5 billion for Canada.[5]

Data is the fundamental "natural resource" of the digital economy and is at the heart of the innovation economy. Digital technologies transform data into insights that allow companies to be more innovative with products and services and compete more effectively in the global market. Marketplace frameworks that leverage and promote Canada's data at home and abroad will boost innovation and wealth creation in the global economy. Companies need to pay close attention to the strategic importance of data and the returns they can reap on their investments in data tools and analytics.[6]

The Value of Data and Advanced Data Processing Activities and the Role of Technology

"The world's most valuable resource is no longer oil, but data."[7] In today's digital world, data, like oil, goes through many refinements and iterations as it becomes useful in a multiple of ways. Smartphones and the internet have made data abundant, ubiquitous and far more valuable. Artificial intelligence (AI) techniques such as machine learning (ML) extract more value from data.[8] Sensors have amplified the amount of data available for use. However, the "data is oil" analogy is an imperfect one because data itself is not a commodity. The same data can be hosted or processed by multiple entities.

However, the "data is oil" analogy does demonstrate that data intensive activities involving, for example, AI need data to run. AI takes in raw data and converts it into something useful for decision-making. The organization uses data it has accumulated over time as input to train an algorithm and uses the algorithm to generate predictions to inform actions. The ongoing value of data usually comes from the actions the organization takes in its day-to-day business – the new data it accrues each day. New data allows operation of the AI model after it is trained, and ongoing operational data enables improvements of the AI model through learning.[9]

However, under privacy frameworks, such as PIPEDA, there is a difference between personal data and nonpersonal or anonymized data. Yet, the boundary between nonpersonal data and personal data may become increasingly blurred when AI is capable of learning underlying relationships between datasets. Furthermore, while these innovative applications of technology may not use personal data, the data they do use can have an impact on individuals

[5] Digital Industries, The sector today and opportunities for tomorrow, Interim Report at 2 [hereinafter Interim Report]
[6] Id. at 4
[7] "Regulating the internet giants," The Economist, May 6, 2017
[8] Id.
[9] A. Agrawal, J. Gans & A. Goldfard, "Is Your Company's Data Actually Valuable in the AI Era?", Harvard Business Review, January 17, 2018; Microsoft transforming work with data and AI

(e.g. personalized advertisements and targeted sales activities). These ramifications suggest that an evolved approach to data protection is needed to balance the benefits and the risks to individuals of more data intensive impacting activities.

Advanced Data Processing Activities, such as AI and ML

The current wave of progress and enthusiasm for AI began in around 2010, driven by three factors that built upon each other: the availability of big data (from sources including e-commerce, businesses, social media, science, and government) which provided raw material for dramatically improved machine learning approaches and algorithms which in turn relied on the capabilities of more powerful computers.[10]

There is no single definition of AI, and the concept of what defines AI has changed over time, but at the core, AI is generally any software which approximates some significant fraction of some aspect of human intelligence.[11] In some cases, a problem is considered as requiring AI before it has been solved, but once a solution is well known, it is considered routine data processing.[12] AI is divided broadly into two stages: Narrow AI which uses the principles of pattern recognition to carry out one specific task (e.g. language translation, self-driving vehicles), and General AI which exhibits apparently intelligent behavior at least as advanced as a person across a full range of cognitive tasks.[13] Narrow AI is already providing breakthroughs (e.g. in medicine where it is used to diagnose patients based on genomic data and in industry where it is employed in the financial world for uses ranging from fraud detection to improving customer service by predicting what services customers will need.)[14] The current consensus is that full General AI will not be achieved for at least decades.[15]

ML is a statistical process that starts with a body of data and tries to derive a rule or procedure that explains the data or can predict future data. An advantage of ML is that it can be used even in cases where it is infeasible or difficult to write down explicit rules to solve a problem. In a sense, ML is not an algorithm for solving a specific problem but rather a more general approach to finding solutions for many different problems, given data about the problems.[16] ML is about teaching computers to learn in the same way humans do, by interpreting data from the world around humans, classifying the data and learning from the computer's successes and

[10] Preparing for the Future of Artificial Intelligence, Executive Office of the President of the United States, National Science & Technology Council, Committee on Technology (2016) [hereinafter U.S. Future of AI Report] at 6; European Group on Ethics in Science and New Technologies Statement Artificial Intelligence, Robotics and 'Autonomous' Systems (2018) [hereinafter EU Statement on AI] at 7; AI in the UK: ready, willing and able? House of Lords Select Committee on Artificial Intelligence (2018) [hereinafter AI in the UK] at 13-14; CNIL Report at 16.
[11] U.S. Future of AI Report at 7; Krupansky, Jack (2017, June 13). Untangling the Definitions of Artificial Intelligence, Machine intelligence, and Machine Learning. *Medium* [hereinafter Krupansky].
[12] U.S. Future of AI Report at 7.
[13] Id.; AI in the UK at 15; CNIL Report at 16.
[14] Marr, Bernard. What is Artificial Intelligence and How Will it Change our World?
[15] U.S. Future of AI Report at 7.
[16] Id. at 7-9.

failures.[17] With ML, everything about the decision procedure is known, but there may be too much information to interpret it clearly.[18] ML is a subset of AI.[19]

Regulatory Response to Advanced Data Processing Activity, such as AI and ML

Regulatory responses to privacy and data protection risks raised by advanced data processing activities, such as AI and ML, range from nonexistent to related.

AI Related Guidance

In the EU, the EU General Data Protection Regulation (GDPR) regulates the processing of personal data[20] and prohibits solely[21] automated decision-making,[22] including profiling,[23] unless the decision is necessary for the performance of a contract, is authorized by law, or is based on explicit consent.[24] In addition, parts of the GDPR address the "full range of rights and interests of an individual" which tracks to the differences between Privacy and Data Protection set forth in the EU Charter of Fundamental Rights.[25] Also the EU Statement on AI calls for a common, internationally recognized ethical and legal framework and proposes a set of fundamental ethical principles based on the values laid down in the EU treaties and the EU Charter of Fundamental Rights.[26] In AI in the UK and the CNIL Report, blanket AI-specific regulation is considered inappropriate or unnecessary, partly because the GDPR appears to address many of the concerns regarding the handling of personal data.[27] Instead, AI in the UK recommends consistent and widely-recognized ethical guidance, in the form of an AI code of conduct,[28] and the CNIL Report recommends two founding principles for the development of algorithms and AI: fairness and continued attention and vigilance.[29]

---

[17] Marr, Bernard. What is Machine Learning – A Complete Beginner's Guide [hereinafter Marr on ML]; CNIL Report at 16.
[18] U.S. Future of AI Report at 9.
[19] Marr on ML.
[20] GDPR Article 1
[21] Solely automated decision-making is the ability to make decisions by technological means without human involvement. Article 29 Data Protection Working Party Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679. WP251 rev. 01 [hereinafter Article 29 Guidelines] at 8.
[22] Automated individual decision-making has a different scope and may partially overlap with or result from profiling, Article 29 Guidelines at 8, and means decisions which produce legal effects concerning an individual or similarly significantly affects an individual, GDPR Article 22(1).
[23] Profiling means any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to an individual, in particular to analyse or predict aspects concerning that individual's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements. Id. Article 4(4). Profiling is composed of three elements: (1) it has to be an automated form of processing; (2) it has to be carried out on personal data; and (3) the objective of the profiling must be to evaluate personal aspects about an individual. Article 29 Guidelines at 6-7.
[24] GDPR Article 22(2).
[25] Charter of Fundamental Rights of the European Union, Ttile II, Articles 7 & 8
[26] EU Statement on AI at 5.
[27] AI in the UK at 116 ¶ 386; CNIL Report at 45.
[28] AI in the UK at 125 ¶¶ 419-420.
[29] CNIL Report at 48-50.

In the U.S., the U.S. Future of AI Report recommends that the approach to regulation of AI-enabled products should be informed by assessment of the risk that the addition of AI may reduce alongside the assessment of the risk that it may increase. If the risk falls within the bounds of an existing regulatory regime, the policy discussion should start by considering whether the existing regulations already adequately address the risk or whether they need to be adapted to the addition of AI. Also, where regulatory responses to the addition of AI threaten to increase the cost of compliance or slow the development or adoption of beneficial innovation, policymakers should consider how those responses could be adjusted to lower costs and barriers to innovation without adversely impacting safety or market fairness. Because the use of AI to make consequential decisions about people, often replacing decisions made by human-driven bureaucratic processes, leads to concerns about how to ensure justice, fairness and accountability, the U.S. Future of AI Report also recommends ethical training for AI practitioners and students, augmented with technical tools and methods for putting good intentions into practice by doing the technical work needed to prevent unacceptable outcomes.[30]

In Singapore, the Personal Data Protection Commission has issued a Discussion Paper on Artificial Intelligence (AI) and Personal Data (Singapore Paper).[31] The objective of the Singapore Paper is to propose an accountability-based framework for discussing ethical, governance and consumer protection issues related to the commercial deployment of AI in a systematic and structured manner. Key preliminary views are that governance frameworks around AI should be technology-neutral and "light-touch", AI developers and user companies should be provided with regulatory clarity when developing AI technologies and translating them into AI solutions, and policies and regulations that promote explain ability, transparency and fairness, as well as human-centricity, as clear baseline requirements can build consumer trust in AI deployments.[32] The Singapore Paper proposes a four-stage governance framework: A. Identifying the objectives of an AI governance framework, B. Selecting appropriate organizational governance measures, C. Considering consumer relationship management processes, and D. Building a decision-making and risk assessment framework.[33]

In Canada two significant AI specific initiatives are underway. The Canadian Institute for Advanced Research (CIFAR) is leading the Government of Canada's $125 million Pan-Canadian AI Strategy,[34] working in partnership with three provincial AI institutes, the Alberta Machine Intelligence Institute, Mila in Montreal, and the Vector Institute in Toronto. Announced in the 2017 federal budget, the strategy has four major goals: (1) increase the number of outstanding AI researchers and skilled graduates in Canada, (2) establish interconnected nodes of scientific excellence in Canada's three major centers for AI, (3) develop global thought leadership on the economic, ethical, policy and legal implications of advances in AI, and (4) support a national

[30] U.S. Future of AI Report at 1-3

[31] Singapore Personal Data Protection Commission Discussion Paper on Artificial Intelligence (AI) and Personal Data – Fostering Responsible Development and Adoption of AI 5 June 2018

[32] Singapore Paper at 2-3

[33] Id. at 7-14

[34] Pan-Canadian Artificial Intelligence Strategy, https://www.cifar.ca/ai/pan-canadian-artificial-intelligence-strategy

research community on AI. Also, the Treasury Board of Canada Secretariat is developing a Directive on Automated Decision-Making[35] to enable the Government of Canada's use of technology and automated systems to make, or assist in making, administrative decisions to improve service delivery and an Algorithmic Impact Assessment[36] to help assess and mitigate the risks associated with deploying an automated decision system.

General Regulatory Reaction

The privacy and data protection risks raised by advanced data processing activities may not be adequately addressed by existing regulatory regimes.[37] Indeed, history has shown that generally privacy and data protection law has lagged technological advances.[38]

Mainframe computers were invented during the 1940s, but it was not until 1976 that the United Nations recognized the basic right to privacy in Article 17 of the International Covenant on Civil and Political Rights (UN International Covenant) which proclaimed that no one should "be subjected to arbitrary or unlawful interference with his privacy, family, home or correspondence, nor to unlawful attacks on his honour and reputation." Relational databases were invented during the 1970s, but it was not until 1980 that the Organization for Economic Cooperation and Development (OECD) created a set of eight fair information principles and codified them in the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (OECD Privacy Guidelines). The first web browser was invented in 1990 and released in 1991, and the use of virtual computers became popular in the 1990s leading to the development of the modern cloud computing infrastructure, but the EU Data Protection Directive (EU Directive), the EU's version of the OECD Privacy Guidelines, was not adopted until 1995, and PIPEDA was not adopted until 2000. Advanced analytics (big data) began in 2006, the first smart phones were released in 2007, when Watson won Jeopardy in 2011, AI broke through, and the Internet of Things started in approximately 2013. The GDPR, which replaces the EU Directive and is the most notable change to data protection legislation in well over two decades, began being drafted in 2012, was not adopted until 2016, and did not go into effect until 2018.[39] Despite its recency and despite its goal of dealing with advanced data driven technologies, when the GDPR was finalized, AI and ML were not in wide use. As evidenced by the GDPR, changes to privacy and data protection laws require lead time, funding and resources due to the complex nature of the subject and the rapid change of the technology. As this chronology shows, privacy and data protection legislation are ill-equipped to keep up with, let alone anticipate, technological changes such as advanced data processing activities.

[35] Treasury Board Directive on Automated Decision-Making, https://docs.google.com/document/d/1LdciG-UYeokx3U7ZzRng3u4T3IHrBXXk9JddjjueQok/edit#heading=h.umd3sgrbb3d9

[36] Government of Canada Algorithmic Impact Assessment (v0.2), https://canada-ca.github.io/digital-playbook-guide-numerique/views-vues/automated-decision-automatise/en/algorithmic-impact-assessment.html

[37] According to the House of Lords Select Committee on Artificial Intelligence, many of the concerns regarding the handling of personal data appear to be addressed by the GDPR. AI in the UK at 116 ¶ 386.

[38] This lag is true in Canada as well as other parts of the world.

[39] According to the House of Lords Select Committee on Artificial Intelligence, the GDPR addresses the privacy and data protection risks by prohibiting solely automated decision-making. AI in the UK at 116 ¶ 386.

Revolutionary and innovative developments in advanced data processing activities, like AI and ML, which leverage the tremendous availability of data from the plethora of sensors increasingly being used, present challenges to the regulatory strengths and effectiveness of existing personal data and consent-based data protection laws. These laws have remained largely unchanged since the creation of the OECD Privacy Guidelines in 1980. For AI and ML, the performance of human intelligence processes by machines (especially computer systems), which may include harvesting of data from multiple sources, a consent-based approach to privacy is not fully effective because often there is no human involvement. The growth in sensors and the Internet of Things has made the intricacy of complex data systems challenging for any individual to understand and, consequently, in a meaningful way, to "consent". As a result, while the consent-based approach will remain an important feature for future governance approaches,[40] given the significant impact nonpersonal data can have on individuals, a broader view of governance, in which all types of data can be considered, as appropriate, is needed. Rather than trying to change the definition of personal data to keep up with the technological developments, for example, it is more workable to develop governance mechanisms that help organizations determine that data is not being used in an inappropriate fashion. This result suggests a need to evolve the accountability expectations on the organizations that use data. Accountability increasingly includes being responsible for the consequences of all data practices. In addition, foundational mechanisms designed to ensure meaningful individual participation, including consent, need to also evolve.

Challenges to Canadian Law

Technological innovations, such as AI and ML, have created challenges for the cornerstones of Canada's federal private sector privacy law – accountability and consent – the tools that afford individuals the opportunity to stake their autonomy and exercise control over their information and that enable organizations to collect, use and disclose personal information. Organizations that wish to collect, use or disclose that data must, by law, seek and obtain consent but cannot always pinpoint or predict every reason for which personal information may be used or disclosed in today's rapidly changing, data-driven marketplace. Consent may be a poor fit in certain circumstances, for example, where consumers do not have a relationship with the organization using their data and where uses of personal information are not known at the time of collection or are too complex to explain to individuals in an effective way. The increasingly complex digital environment poses challenges for the protection of privacy and the consent model.[41]

Consent is one foundational element of PIPEDA, but when PIPEDA was adopted, interactions with businesses were more predictable, transparent and bidirectional. In this digital environment, it is no longer entirely clear to consumers who is processing their data and for what purposes.[42] Situations in which consent may be simply impracticable (or at the very least

[40] OPC Report on Consent at 1
[41] Id.
[42] Id.

challenging) include, for example, where there is no relationship between an individual and the organization collecting and using personal information, such as search engines, and in the case of big data initiatives or Internet of Things devices that individuals have no choice but to use. In the current digital ecosystem, it is no longer fair to ask consumers to shoulder all of the responsibility of having to deconstruct complex data flows in order to make an informed choice about whether or not to provide consent.[43] Everyone – individuals, organizations, regulators and legislators – needs to play their part for privacy to be protected effectively.[44]
While accountability remains a second core foundational element of PIPEDA, the expectations on organizations as to accountability were not provided until well after PIPEDA came into effect in 2000. The Office of the Privacy Commissioner of Canada and provincial commissioners in Alberta and British Colombia adopted accountability guidance in 2012.[45] Recently, the OPC emphasized that companies should operationalize and respect their obligations under PIPEDA. Specifically, the OPC reiterated that accountability is a fundamental PIPEDA principle that requires organizations to develop and implement policies and practices to uphold the fair information principles set forth in PIPEDA.[46] The OPC also expressed its support for two elements of privacy by design: its temporal requirements (as early as possible and continuously assessed) and the addressing of both technological and organizational factors. The OPC also emphasized that accountability requires organizations to be able to demonstrate the steps they have deliberately taken to design and implement the requirements of PIPEDA.[47]

As the OPC has recognized, the two pillars of PIPEDA – consent and accountability – need to evolve to address the challenges presented by data intensive activities and technologies. In fact, evolving accountability can help address the situations when consent may be challenged or impracticable. Accountability can address how individuals will trust advanced data processing, so they will trust organizations with their data. Being responsible for the consequences of data practices increases trust in the organization's use of data.

Enhanced Data Stewardship

Uses of data that do not easily enable meaningful consent, uses that may not be within the individual's expectation, uses that cannot be explained effectively through transparency alone, can raise issues about trustworthiness of advanced data processing activities. How does the individual trust that the organization is not using the data in a way that adversely impacts his or her rights or interests? As the OPC's Report on Consent observed, current privacy and data protection laws do not effectively or fully address these issues.[48] Yet, much innovation, much

[43] Id. at 7
[44] Id.
[45] The Office of the Privacy Commissioner of Canada (OPC) and the Offices of the Information and Privacy Commissioners (OIPCs) of Alberta and British Columbia, "Getting Accountability Right with a Privacy Management Program," April 17, 2012
[46] Id. at 16
[47] Id. at 17
[48] OPC Report on Consent at 7. To some extent, the legitimate interest basis for processing under the EU Directive and the GDPR attempts to addresses this issue but only for personal data.

transformation of data into information and information into insight, depend on the organization's use of data that the individual may not anticipate but might benefit from (e.g. collision avoidance systems on cars).[49]

In order to encourage innovation, digital information strategies are being adopted which recognize that the internet and digital technologies are transforming the world, that the needs of business, government and the general public impact the competitiveness of their country's economy, and that the protection of personal data and fair data processing are needed for the development of Internet-based economies.[50] If individuals do not trust how organizations are using their data and how organizations are transforming data into information and information into knowledge, and the law is not keeping up with the technology, organizations need guidance on how to act ethically and apply equitable principles. This guidance is needed particularly with respect to advanced data processing activities, such as AI and ML and sensor driven data processing, and to the application of knowledge that enables data driven innovation to reach its full potential.[51]

Acting ethically, fairly and responsibly means organizations need to understand and evaluate advanced data processing activities and their positive and negative impacts on all parties. This approach means organizations will need to be effective data stewards and not just data custodians. A data custodian is someone who manages describable data duties (a compliance function). A data steward considers the interests of all parties and uses data in ways that create maximum benefits for all parties while minimizing risks to individuals and other parties (a broader accountability function). A data steward asks whether the outcomes of advanced data processing activities are legal, fair and just.[52] Legal, fair and just is a proxy for ethical and associated and describable values. To determine whether advanced data processing activities, such as AI and ML, that may be impactful on people in a significant manner and/or that directly impact people, are ethical or fair, organizations should define values that are reduced to core or guiding principles and then are translated into organizational policies and processes.

This approach is similar to corporate social responsibility which encompasses the economic, legal and ethical expectations that society has of organizations at a given point in time.[53] Like corporate social responsibility, organizations should have a corporate data responsibility which encompasses the economic, legal and ethical responsibilities they have at a given point in time with respect to the data they collect, use, or disclose. These responsibilities form the basis for data stewardship.

---

[49] IAF, "Artificial Intelligence, Ethics and Enhanced Data Stewardship", September 20, 2017 [hereinafter IAF Paper], at 5-7

[50] E.g. Interim Report

[51] OPC Report on Consent at __

[52] IAF Paper at 6-7

[53] Schwartz, Mark S. & Carroll, Archie B, "Corporate Social Responsibility: A Three-Domain Approach", Business Ethics Quarterly, 2003, Vol 13, Issue 4, at 503, 509

Like corporate social responsibility, ultimately, data stewardship is predominantly driven by organizational policies, culture and conduct and not just technological controls. Thus, the core question is: what does an appropriate trustworthy accountability framework look like for a data steward?

Enhanced Data Stewardship Accountability Elements

In 2009, the accountability principle in the OECD Privacy Principles formed the basis for the Essential Elements of Accountability (Essential Elements).[54] In 2010, the EU Article 29 Data Protection Working Party issued opinion 3/2010 on the principle of accountability.[55] The Office of the Privacy Commissioner of Canada and provincial commissioners in Alberta and British Colombia adopted accountability guidance in 2012.[56] Hong Kong and Colombia issued accountability guidance in 2015.[57] Now, accountability is the foundation of the GDPR.[58] The guidance and the adoption of the GDPR has elevated accountability from check-box compliance to a risk-based approach but has not kept up with the advanced data processing activities, such as AI and ML, and the way sensor driven data can be accessed and used, that may be impactful on people in a significant manner and that directly impacts people. To be able to transform data into information and information into knowledge and insight and knowledge into competitive advantage, for individuals to be able to trust data processing activities that might not be within their expectations, enhanced data stewardship accountability elements are needed. Against this backdrop, the question is: how can data intensive activities and technologies that have an impact on individuals be conducted in a fair, responsible and ethical manner while achieving the desired benefits?

The Enhanced Data Stewardship Accountability Elements for Data Processing Activities, such as AI and ML, that Directly Impacts People (Enhanced Elements)[59] call for organizations to:

1. Define data stewardship values that are reduced to guiding principles and then translated into organizational policies and processes for ethical data processing.
2. Use an "ethics by design" process to translate their data stewardship values into their data analytics and data use design processes so that society, groups of individuals, or individuals themselves, and not just the organization, gain value from the advanced data processing activities, such as AI and ML, and require Ethical Data Impact

[54] Essential Elements, http://www.informationaccountability.org
[55] WP 173
[56] The Office of the Privacy Commissioner of Canada (OPC) and the Offices of the Information and Privacy Commissioners (OIPCs) of Alberta and British Columbia, "Getting Accountability Right with a Privacy Management Program," April 17, 2012
[57] PCPD, "Privacy Management Programme: A Best Management Guide," February 18, 2014; Columbia Superintendence of Industry and Commerce, "Guidelines for the Implementation of the Accountability Principle," May 2015
[58] GDPR Article 5(2).
[59] The Enhanced Elements are published separately at http://informationaccountability.org/wp-content/uploads/Enhanced-Data-Stewardship-EDIA-FINAL-10.22.18.pdf.

Assessments (EDIAs)[60] when advanced data analytics may be impactful on people in a significant manner and/or when data enabled decisions are being made without the intervention of people.[61]

3. Use an internal review process that assesses whether EDIAs have been conducted with integrity and competency, if the issues raised as part of the EDIA have been resolved and if the advanced data processing activities are conducted as planned.
4. Be transparent about processes and where possible enhance societal, groups of individual or individual interests; communicate the data stewardship values that govern the advanced data processing activities, such as AI or ML systems developed, and that underpin decisions widely; address and document all societal and individual concerns as part of the EDIA process and design individual accountability systems that provide appropriate opportunities for feedback, relevant explanations and appeal options for impacted individuals.[62]
5. Stand ready to demonstrate the soundness of internal processes (i.e. demonstrate that internal processes are in compliance with ethical standards that are over and above compliance with the law) to the regulatory agencies that have authority over advanced data processing activities, including AI or ML processes, as well as certifying bodies to which they are subject, when data processing is or may be impactful on people in a significant manner.

The Enhanced Elements are the foundational support to an overall data stewardship and values driven framework that addresses uses of data that do not easily enable meaningful consent, uses that may not be within the individual's expectation, uses that cannot be explained effectively through transparency alone.

[60] EDIA's are a label proxy. Other names such as Responsible Date Impact Assessments could also encompass the same process objectives.

[61] Like the socially beneficial activities exception to consent proposed by the OPC, OPC Report on Consent at 15, in the EDIA, benefits must outweigh privacy impacting consequences as well as other risks.

[62] Like the socially beneficial activities exception to consent proposed by the OPC, OPC Report on Consent at 15, the organization issues a public notice describing its practices.

This connected framework is also supported by a Process Oversight Model to enhance the trustworthiness and effectiveness of Data Stewardship accountability including the effectiveness of EDIA's.[63]

The Role of an EDIA

An EDIA is a process that looks at the full range of rights and interests of all parties in a data processing activity to achieve an outcome when advanced data analytics may be impactful on people in a significant manner and/or when data enabled decisions are being made without the intervention of people. An EDIA assists an organization in looking at the rights and interests impacted by the data collection, use and disclosure in advanced data-driven activities.

To determine whether an EDIA may be necessary, the organization should consider, before the activity begins and when there are any changes which affect the scope of the activity, whether the data processing activity involves advanced analytics or the potential for a significant impact on individuals. For example, an EDIA might be necessary for activities such as: evaluation or scoring (including profiling and predicting), automated individual decision-making, systemic observation or monitoring, data processed on a large scale, matching or combing data sets, innovative use or applying new technological or organizational solutions (such as AI and ML). If the data processing activity may have an impact on an individual or on a group of individuals that may not be anticipated or easily known, then an organization should consider whether an EDIA should be done either at the concept stage or at the service/product/analytical development stage or at both stages. If the organization determines that an EDIA is not

[63] An example or Model EDIA and the Process Oversight Model are published separately at http://informationaccountability.org/wp-content/uploads/Enhanced-Data-Stewardship-EDIA-FINAL-10.22.18.pdf.

appropriate for the data processing activity, then only a privacy impact assessment (PIA) may need to be completed.[64]

The very nature of an ethical and values-based assessment requires a careful consideration of the data activity benefits as well as the risks to individuals and society, considering the interests of all the parties who may be part of the activity. While open, structured questions can help, a way to organize the ultimate decision as to whether to proceed can be evaluated by using a well-established risk modeling process where the outcome of the analysis (significance, likelihood and effectiveness of controls) is qualified and quantified

Successful implementation of an EDIA assumes and depends on the full implementation of the Enhanced Elements and, in particular, on highly qualified and competent, accountable roles and responsibilities with appropriate separation of duties. For example, EDIAs could be conducted by the privacy group, engineers in the business, a combination of several parts of the organization while the review and decision-making process should be separately done by people not accountable for the actual data processing activity.

The EDIA is broader in scope than the typical PIA. For example, all data are considered in an EDIA and not just personal data. Therefore, all aspects of the EDIA include data in the aggregate, non-identifiable form that may be outside the scope of PIPEDA and many other privacy and data protection laws. However, to the extent the EDIA can be used to consider and appropriately mitigate the impact of a personal data practice, the EDIA process may supplement (or be woven into) the organization's PIA process. In this regard, the EDIA process may enhance and augment an organization's privacy management program and compliance with its legal obligations under PIPEDA or similar regulatory frameworks.

An EDIA does not replace a PIA; it is designed to be used in conjunction with PIAs; it is not a complete PIA. Organizations my incorporate the EDIA in whole or in part into their own unique processes and programs so as to supplement or evolve with their PIA processes.

Process Oversight

Assessments conducted solely by the parts of an organization implementing intensive data activities may raise issues of trustworthiness. Where the oversight of the assessment and accountability process is done by the organization itself, the oversight process should look at how the organization has translated organizational ethical values into principles and policies and into an "ethics by design" program. The oversight process should also consider how well-established internal review processes, such as EDIAs and effective individual accountability systems, have been implemented. It is presumed that the oversight process is independent from the assessment process. The oversight process could be a function performed by, for example, an internal audit group or an internal control function.

[64] The OPC has issued guidance on PIAs. See https://www.priv.gc.ca/en/privacy-topics/privacy-impact-assessments/gd_exp_201103/

The function of the oversight process may be likened to an assessment of "controls and controls effectiveness" by the internal audit group. Examples of oversight control areas that could be evaluated are:

I. Accountability for the oversight process,
II. Translation of organization values into principles and policies,
III. Translation of organizational values into an "ethics by design" program,
IV. Utilization of the EDIA,
V. Internal review process,
VI. Individual accountability system,
VII. Transparency of process.

The oversight process does not oversee the conduct of individual EDIAs but rather the conduct of the overall EDIA process as well as key elements of Data Stewardship Accountability. When conducting an EDIA, the participants are expected to evaluate the activity to the best of their ability. When overseeing the EDIA process, the overseers (e.g. the internal audit group) are evaluating the integrity with which EDIAs were conducted. If EDIAs repeatedly (not occasionally) are inaccurate in balancing risks and benefits, then perhaps the process is not operating correctly (not that individual EDIAs were conducted incorrectly).
Evidence of oversight is important. Whether this oversight occurs internally, for example by the internal audit group, or externally, for example by a consulting firm, it is necessary that documentation exist that demonstrates how the oversight was conducted and that, in fact, it was conducted.

The oversight process should measure whether the EDIA process is being conducted with honesty and recognizes the full range of interests of all parties to evidence that the interests of the organization were not placed in front of the interests of other parties.[65] The organization should stand ready to demonstrate its assessment governance process and individual assessments to regulators with appropriate authority.[66]

Recommendations:

The following recommendations are intended as a starting point for discussions in Canada as it evolves the application of privacy principles and/or evolves other public policy approaches. They are divided into two parts. The first is a set of governance principles that in many ways are or could be used in a similar manner to PIPEDA's existing Schedule 1. They support the goals of accountable data stewardship and evolved ways to ensure individuals have meaningful and effective ways of engaging and participating in arenas where data impacts them.

[65] IAF Oversight Report at 21
[66] Id. at 23-24.

The second portion is a set of five high level policymaking objectives that track to the principles themselves. They also include an objective for regulatory guidance and enforcement. This addresses a need for trusted and predictable regulatory oversite.

Fair Processing Principles to Facilitate the Fourth Industrial Revolution in Canada

New information and communications technologies applied to the physical, digital and biological worlds are becoming a major driver of Canadian economic and opportunity growth. Toward that objective the Canadian government has adopted an artificial intelligence strategy to take full advantage of this opportunity to drive economic growth. To facilitate this growth Canada will need a common criterion for assuring business and academia may think and learn with data and implement insights in a socially responsible manner. Canada has been a privacy trend setter, but this goes beyond privacy to fairness that will require an ethical data stewardship design. Canada will need rules of the road for business using data beyond common understanding, and updated rights to protect individuals and society.

To advance the ethical data stewardship deign model, the Information Accountability Foundation (IAF) proposes a set of principles. Their operational objective is to:

- *Responsibly and safely facilitate the fourth industrial revolution in Canada where digital, physical and biological come together;*
- *Preserve the rights and interests of Canadians;*
- *Be interoperable with other new and emerging information governance regimes, and*
- *Enable all the benefits of the 21st century information age.*

While interoperable with other regimes, this framework is Canadian in its vision. It is heavily influenced by core concepts such as reasonable, appropriate, proportional, legitimate and importantly accountability that are well developed in Canada.

The principles framework is divided into two parts. The first part describes the rights necessary for individuals to function with confidence in our data driven world. The second part focuses on the obligations that organizations must honor to process and use data in a legitimate, fair and responsible manner. This framework draws heavily and expands on PIPEDA's accountability principle. This original OECD principle has been best developed in Canada, and this framework continues Canadian leadership. While the framework outlines principles, in some cases it includes means and outcomes to better illustrate a principle's intent.

Individual Rights

1. **Transparency** Individuals have the right to be free from secret processing of data that pertains to or will have an impact on them. Organizations should provide understandable statements about their data collection, creation, use and disclosure

practices and about their policies and governance. While these statements may be directed at enforcement agencies, they should also be publicly available. To augment this level of transparency, Organizations should also provide summaries and other means that make their data collection, creation, use and disclosure practices understandable to individuals.

2. **Access and Redress** Individuals have the right to request access to and information on the data they provided, to understand what data is observed by the organization that pertains to them, and to be told what types of data are inferred by analytical algorithms. They also have the right to request changes to data to ensure accuracy, provide feedback, and receive relevant explanations on data use. Individuals have the right to object if they believe that the data about them is inaccurate or being used out of context, is not being undertaken in an accountable manner, or if they believe that uses of data are not legitimate. The right to object to processing does not pertain where data processing and use are permitted by law. Because intellectual property rights may prevent individuals from having full access or disclosure of inferences made by the organization, and where inferences such as scores potentially have negative consequences for individuals, organizations should provide relevant explanations about their processing, appropriate opportunities for feedback, and the ability for individuals to dispute such processing.

3. **Engagement and Appropriate Control** Individuals have the right to control data uses that are highly consequential to them. This should be facilitated through an appropriate level and contextual application of consent where possible. Where consent isn't possible or less impactful, they have the right to know that accountability processes assure the data uses are fair and responsible. Individuals also have the right to know that data is disclosed to third-parties beyond the context of the relationship or the legitimate purpose of the data use and to request such disclosure not take place, with the exception of data shared to assure security or for public purposes required by law. Where highly consequential uses, such as health, financial standing, employment, housing and education, are governed by specific laws, those laws take priority.

4. **Beneficial Purposes** Individuals have the right to expect that organizations will process data that pertains to them in a manner that creates benefits for the individual or for a broader community of people. In cases where the organizations receive most of the benefit, a demonstrable vetting process should determine there is minimal risk or impact to an individual. All data processing including for beneficial purposes should be part of understandable summaries required under the Transparency Principle. Where there are benefits and the potential for negative consequences to individuals, individuals should expect an explanation of the results and the ability to dispute the findings, as provided in the Access and Redress Principle.

Accountable Data Stewardship

5. **Assessed and Mitigated Impacts** All collection, creating, use and disclosure of data should be compliant with all applicable laws, industry codes, and internal policies and

practices, and should be subject to privacy, security and fair processing by design. Employees should receive appropriate training for their specified roles, and accountable employees should be identified to oversee privacy, security and fair processing obligations. Specifically, fair processing assessments should identify individuals and groups of individuals who are impacted, both negatively and positively, by the processing, and should guard against identifiable negative consequences. Where there are negative consequences, organizations should mitigate those consequences to the degree possible. If unacceptable consequences still persist for some individuals or groups, the organization should document why the benefits to other individuals, groups and companies are not outweighed by the unacceptable consequences.

6. **Secure** Data should be kept secure at a level that is appropriate for the data.
7. **In Context** Data should be collected, created, used and disclosed within the context of the relationship between the individuals to whom the data pertains and the organization, based on the reasonable expectations of individuals as a group. Public safety, security and fraud prevention are considered within context.
8. **Legitimate Uses** Data should be processed only for legitimate uses that have been disclosed, would be expected or are consistent with those uses. When the data is no longer necessary for the legitimate use, it should not be retained in an identifiable manner. Legitimate uses include the following:
    A. Where individuals have provided informed consent.
    B. Ongoing business processes such as fraud prevention, accounting and product improvement that would be expected of an enterprise.
    C. Freely thinking and learning with data by organizations that demonstrate effective accountability, including mitigating risks to individuals, consistent with the societal objective of encouraging data driven innovation, and that honor the Onward Responsibility Principle.
    D. Uses that create definable benefits for individuals, groups, organizations and society that are not counterbalanced by negative consequences to others, and that are based on assessments established by external criteria.
    E. Designated public purposes, including public safety and in response to an appropriate legal request.
    F. Organizations that stand ready to demonstrate why they believe other uses that are based on assessments established by external criteria are legitimate.
    G. Where permitted by law.
9. **Accurate** Data should be accurate and appropriate for all legitimate uses and that level of accuracy should be maintained throughout the life of the data.
10. **Onward Responsibility** Organizations that originate data should be responsible for assuring the obligations initially associated with the data are maintained within the accountability chain. As the data chain expands the previous data originator bears responsibility for the accountability chain. When data leaves the accountability chain, for example when requested by the government the party providing the data to the government is only accountable for assuring the government has a legal right to request that data and the disclosure is as limited as possible.

11. **Oversight** Organizations should monitor all uses of data to ascertain that the uses are legitimate, the data is processed fairly, the data is accurately used within the context of the relationship with those to whom the data pertains, and processes that support individual rights and accountable data stewardship are effective and tested. The oversight process, whether conducted by an internal body or an external agent, should be separate from and independent of those persons associated with the processing.
12. **Remediation** Organizations should stand ready to demonstrate the effectiveness of policies, practices and internal oversight to those that have external authority for oversight. Organizations should consider rectifying negative consequences where they reach a level of significant impact to individuals.

Recommended Policymaking Objectives

1. Data pertaining to individuals should be used to create real value for identified stakeholders in a balanced, fair fashion that serves people, society and private organizations. This should take place in all sectors and the risk of not using data should be as important to considerations as the risk of negative consequences.
2. Individuals have clear rights related to data and its uses and those rights should be explicit and actionable in reality as well as theory.
3. Accountability requires organizations to be reasonable and responsible in what they do with data pertaining to individuals, and answerable for how they demonstrate that they are acting as effective data stewards.
4. Organizations should have checks and balances in place to assure that their data stewardship is conducted effectively. When organizations cause negative consequences that are consequential, they should take actions to mitigate those consequences.
5. Enforcement agencies should have the powers and resources so that they may act in a manner trusted by the public and seen as predictable by those subject to enforcement.

Conclusion

Together, these revised principles and policymaker recommendations provide a framework for future legal structures that accomplish the goal of enabling Canada's digital transformation. The recommendations enable the rich opportunities and potential of leveraging data to provide the multitude of benefits advanced data processing creates while protecting the privacy of citizens, promoting fairness, and fostering equality of opportunity for all Canadians. The recommendations also enable the benefits of leveraging the potential of data in a way that makes Canada and Canadian business more open and accountable. Recognizing the strengths of individual rights made possible through instruments such as PIPEDA, the recommendations evolve these strengths to provide more effective privacy protection and equity commensurate with more complex and individual impacting data processing. This paper is intended to establish a basis for a Canadian public policy discussion.