

# DIGITAL TERMS OF SERVICE FOR AI PROVIDERS IN CANADA

*DRAFT PROPOSAL (V1) - 11 July 2023*

© Copyright 2023 by [Tim Boucher](#), Canada,

*All rights reserved.*

---

## Section I.

### 1) Preamble

We envision a future where everyone has free access to artificial intelligence (AI) and the benefits it can bring, if they so choose.

We envision a future where the development of AI technology is deeply rooted in, protects, and allows for the full expression and flowering of human rights.

We envision a future where core values like human autonomy and creativity are enshrined and expanded through the conscious use of appropriate technologies.

We envision a future where AI technologies serve us, and adapt to our unique shape as human beings, rather than forcing us into narrow algorithmic roles which don't suit us.

We envision a future where the ownership and direction of AI technologies is not held in the hands of a few corporations, but broadly distributed and democratically controlled by regular people.

We envision a future where non-AI alternatives are valued and preserved as options for those who choose a different path.

In short, we envision a future with AI that is radically different from the trajectory we find ourselves on today as a society.

The purpose of this document is to provoke our imaginations. It does not set out to present a plan which is easily politically feasible today, nor is it intended to present guidelines which would be convenient for companies to follow. It's purpose is to explore how we might make things different if we really sat down and thought it through, and talked about the challenges and opportunities of AI openly, using our whole human selves.

In short, the purpose of this document is to spark conversations that might lead to new possible futures to explore and test in the crucible of our imaginations, instead of merely upholding the status quo, and plugging our ears while the world crumbles.

## 2) What Is This Document

This document, The AI Terms of Service, was written within the context of artificial intelligence development and deployment Canada, but its content is largely applicable to any nation, or to users of AI services outside of a national context entirely. The premise is that Canada (or any country, or even a non-state political entity), is a “platform,” and if service providers wish to access that collective platform, then they need to agree to common ground rules. If they deviate from the agreement, there can be specific consequences, up to and including removal from the “platform.”

The AI Terms of Service is meant as an aspirational sort of *Magna Carta* for users of artificial intelligence services. It imagines what users might expect from AI service providers who agree to such a set of mutual guidelines, and outlines roughly what AI companies and services can and cannot do if they wish to have access to Canada-as-a-platform.

This document is intended to conceptually link to the government of Canada's efforts to create a [Digital Charter](#), which itself attempts to translate the values of the [Canadian Charter of Rights and Freedoms](#) to the digital sphere. It aims to articulate protections Canadians should be able to expect when they go online, such as data protection, or protections against misinformation. Variations of these ideas could easily be adapted and applied anywhere, though they would require firmer commitments by AI companies to human well-being than we typically see in the market today. This document is intended to provoke those important conversations publicly, both inside and outside of Canada.

This document is also intended as a commentary on Canada's meager attempt at passing AI legislation, the [Artificial Intelligence and Data Act](#) (also known as AIDA, a part of Bill C-27), which leaves the actual meat of the matter up to eventual future regulations, the contents of which are not hinted at. Reading the text of AIDA, one is left with the distinct impression that ‘no one knows’ what we should do about AI in Canada or globally, but that is not the case. The EU's [AI Act](#) stands out as one notable and worthwhile example which Canada might be inspired by.

This document, however, goes well beyond what even the EU AI Act requires, pushing into areas that might seem like science fiction to some readers. We welcome the comparison, as we

recognize that science fiction has for generations inspired people to become engineers, scientists, and entrepreneurs, capable of bringing their sci fi visions to life. We believe the best path through the future will be the path of imagination.

### 3) Notes On Structure

This document is loosely structured in the format of Agile user stories (“As a user, I want to \_\_\_”), in order to articulate a concrete vision of what user protections built into AI systems might feel like from the user perspective. Agile is a common product development methodology used to define and manage which features one might want to build during a software development process. Given the ascendancy of AI technologies as consumer-facing software products, this feels instinctively like a good framework to get a handle on what we might want to design - or to avoid - collectively as engaged user/owners of AI services.

This document is broken out into fairly specific sections intentionally, in order to highlight areas so they don’t get lost or overlooked in larger umbrella groupings. At times in the document, this may result in light repetition of key points, but hopefully this serves to increase clarity of intended meaning and context.

### 4) Provenance

Everything in Section I of this document was written [by a human](#).

The groundwork for everything in Section II of this document was generated in ChatGPT (v4) based on ample suggestions, examples, and direction provided by the text’s human author and editor.

Generated list items and categories were clarified, sorted, and de-duplicated manually by the document’s editor, before being processed once again using Claude by Anthropic for consistency of final formatting, and checked a final time manually.

## Section II.

### 1) Introduction

The following comprises the text of the AI Terms of Service, presented as Agile user stories which seek to identify the most important user needs and desires which it is proposed that AI service providers would be required to meet in order to do business in Canada.

For those unfamiliar with the structure of Agile user stories, assume that each item listed below begins with the following prefix

*As a user of AI services in Canada...*

### 2) System Specifications

I want AI systems to:

- Identify itself as an AI system
- Identify what AI models or other technology are being used, and whether they are proprietary or open source.
- Link to model cards or other equivalent official documentation that describes the model's abilities, limitations, risks, human impact assessments, mitigations, etc.
- Explain how the system works in plain and clear terms and its intended purpose.
- Be transparent about its limitations, so that I know when it might be better to seek help from a human or use a different tool.
- Document how the AI is tested and validated, so I can establish baselines around the reliability of its performance.
- Provide me with information about the data it was trained on, so I can understand its biases, limitations, and evaluate its potential legality.
- Share the guiding principles, ethics, or beliefs that were considered or incorporated in the development and functioning of the AI system,;and as a commitment to those principles, regularly document concrete proof of how those principles are expressed within the product.
- Never invent or hallucinate results regarding official system specifications, service level information, or accountability and compliance information; and display verification timestamps of the information contained in these messages to show that it is human-verified and recently updated.

### 3) Service Level Information

I want to:

- Have accurate contact information for support and legal issues so I can reach out if I need help or have concerns, and verify that I will get help from a human
- Have quick access to help documentation
- Have quick access to Terms, Privacy Policy, and other relevant legal documents.
- Be informed about the AI's current and upcoming service status so I can plan my interactions with it, and know right away if it is not working correctly currently.
- Have access to regulatory compliance and other certification or auditing information regarding the service and model.

## **4) Accountability & Compliance**

I want the AI system to:

- Comply with all relevant laws and regulations so I can trust it respects my rights and privacy.
- Provide information about its regulatory compliance and certifications so I understand its legal and ethical responsibility, and evaluate whether it is acting accordingly.
- Have robust mechanisms for addressing grievances and providing redress for AI-caused harms, including independent judicial processes.

I want to:

- Know who is accountable if the AI causes harm or mistakes, so I can seek redress.
- Have assurance that regulatory bodies can intervene if the AI system's actions violate my rights or relevant rules, regulations, or best practices.
- Inspect certification from trusted external bodies verifying the ethical, safe operation of the AI.
- Know the AI company is legally obligated to act on independent audit findings.
- Know the AI company has agreements with external bodies to address my feedback and concerns.
- Have the company respect demands related to transparency and ethical practices, as validated by trusted external organizations.
- Have assurance the company uses third-party reviews for risk management and ensuring AI safety.

## 5) Personalization & Customization

I want to:

- Know if and how the AI can be customized to my needs, preferences, or context.
- Not have the AI pretend it is human or anthropomorphize itself unless I explicitly opt-in to this feature.
- Not have the AI give unsolicited ethical advice or judgements, unless I double opt-in to any functionality providing guidance on morals, values or behavior.
- Have the option to use the service anonymously, without a user history or other profiling.

I want the AI system to:

- Respect cultural, social, and personal differences, norms, and sensibilities in its interactions with me.
- Use a preferred language and tone that is appropriate for the context of our interaction, so that it is easy and pleasant to use.
- Handle sensitive topics carefully and respectfully, respecting data protection settings.
- Enable customization at a system and conversation level the AI's behavior, responses, and personal data handling.
- Enable opting-out of certain AI functionalities that I'm uncomfortable with.
- Adapt to my needs and preferences over time, improving the user experience; or alternatively, not learn from my prior interactions, per my data protection settings.
- Enable adjusting the level of AI involvement in tasks based on my preference.
- Be able to set boundaries for AI interactions, like "quiet hours" when I don't want to be disturbed or monitored.
- Recognize and respect my comfort level with the technology, and not push me beyond my chosen settings.

## 6) Data Protection Practices

I want to:

- Know about the AI's data handling practices to ensure my data is used and stored responsibly.
- Not have the AI collect or use my personal data for anything without my explicit permission.
- Be asked for consent before the AI shares any of my data with third parties.
- Be informed if the AI is using my previous interactions to make current or future decisions; and have the ability to turn this on or off permanently or temporarily.
- Have the ability to reset, erase, or withdraw my personal data from the AI system at any time.
- Be able to use the AI anonymously, without compromising my privacy rights.

- Access and control what data the AI collects about me, even supposedly non-identifiable data.
- Know when AI is making assumptions or decisions that impact me and how those assumptions or decisions are made.
- Be informed about data retention periods and what happens if the service offering ends.

I want the AI system to:

- Keep my data secure so I can trust it with sensitive information.
- Allow me to transfer my data to another service if I decide to switch.
- Enable me to opt out of data collection without losing access to the service.
- Protect my data through strong, transparent, and fair privacy policies.
- Be trustworthy by safeguarding my data from breaches and misuse.
- Let me limit or expand its access to my personal data based on my comfort level.
- Avoid surveilling or monitoring me without consent.
- Allow me to correct faulty assumptions or challenge wrong decisions it makes about me based on my data.

## **7) Transparency**

I want the AI to:

- Provide explanations of its processes and outputs that are easy for me to understand.
- Be clear when it does not know the answer to a question.
- Indicate the level of certainty and cite sources for information it provides.
- Truthfully state when requested information does not exist or is unavailable.
- Explain its content moderation or filtering actions in a way that is clear and precise.

## **8) Alignment with User Goals**

I want the AI to:

- Help me achieve my goals without imposing its own agenda.
- Allow me to use it as an open-ended tool to enhance my capabilities, without forcing me to adopt a certain modality or means of working with it.

## **9) Responsiveness**

I want the AI to:

- Perform requested tasks efficiently and accurately.
- Respect my time by providing concise, relevant responses.

- Be reliable and consistent in its performance.

## 10) Information Accuracy

I want to:

- Have confidence that information from the AI is accurate and up-to-date.
- Be able to request sources/references for information provided.
- Easily send AI responses and context to third-party fact-checkers for verification.
- Have the ability to challenge the accuracy of the AI system's information and receive a swift response or correction.
- Have the option to contribute to the verification process by providing feedback, corrections, or additional information shown to improve the system's accuracy.

I want the AI to:

- Use reliable, high-quality data sources with robust verification processes.
- Correct inaccuracies swiftly when identified and improve from feedback.
- Maintain a correction log showing improvements over time.
- Actively involve third-party fact-checkers to verify information accuracy.
- Clearly indicate which responses have been verified and by whom.

## 11) Impartiality & Neutrality

I want the AI to:

- Be unbiased and impartial in its responses, so my views and decisions are not unfairly swayed.
- Uphold principles of net neutrality, treating all data and users equally.
- Not attempt to manipulate user behavior or belief.

## 12) Discrimination & Bias

I want to:

- Be treated fairly by the AI without discrimination based on my background or characteristics.
- Ensure AI systems respect, integrate, and foster diversity, refusing to perpetuate harmful biases or stereotypes.

I want the AI to:

- Promote fairness and avoid reinforcing harmful biases in its workings.
- Not perpetuate social biases or discriminate against me or others based on protected characteristics like race, gender, etc.
- Provide clear explanations of its decision processes and how biases are mitigated.
- Respond to and learn from my feedback about any biased behaviors or outcomes it exhibits.
- Have regular independent audits to ensure it remains unbiased and fair.
- Make it easy for me to report any instances of discrimination or bias I encounter.

## 13) Accessibility

I want the AI system to:

- Offer multiple modes of interaction, including voice, text, touch, gesture, etc., to cater to different abilities and preferences, ensuring I can interact with the AI in the way that best suits me.
- Adhere to principles of universal accessible design, ensuring usability for people with diverse abilities in diverse situations.
- Offer customizable interface options to meet my unique accessibility needs for text, contrast, sound, etc.
- Be fully compatible with assistive technologies like screen readers, braille displays, and alternative input devices.
- Provide captions for audio content and transcripts for video content to accommodate deaf or hard of hearing users.
- Use clear, simple language for all communications and instructions to accommodate users with cognitive or learning disabilities.

## 14) Preserving Human Autonomy

I want to:

- Use AI to enhance my human capabilities, not replace them.
- Learn from the AI without it creating a dependency or undermining my skills.
- Be informed about the AI's limitations to avoid unrealistic expectations.
- Prevent unauthorized transactions or commitments made by the AI on my behalf.

I want the AI to:

- Operate as a tool under my direction and control.
- Avoid imposing its own values, worldviews, or agendas on me.
- Support my creativity and individuality, rather than making me conform to its algorithms.

- Preserve and respect my right to digital autonomy in our interactions.

## 15) Feedback & Improvement

I want to:

- Know if the AI is continuously learning or if its knowledge is fixed at a certain point, so I can understand its limitations and capabilities.
- Know if and how the AI learns from my feedback so I can understand how my interactions affect its performance, or don't.
- Easily provide corrections and performance feedback to contribute to the AI's improvement.
- Have access to concrete data that shows how the system improves over time.
- Understand how my data is used to improve the AI's performance; and have the ability to opt out of this in accordance with my user settings
- Optionally participate in the AI's development through user testing, feedback, or other involvement, if I choose to.
- Contribute to the AI's training data if I choose, with transparency on how my training data will be used and compensation at fair market value for training data I provide the system.

## 16) Content Moderation Practices

I want to:

- Be informed about the AI system's content moderation policies and community safety guidelines to know what is allowed and what is not.
- Understand the mechanisms used for moderating content, whether AI, human, or hybrid.
- Indicate whether a moderation decision was made by a human, AI, or hybrid.
- Know the process to report inappropriate or harmful content or behavior by the system or other users.
- Receive timely responses and action on reported violations of safety guidelines.
- Get clear explanations when content is removed or flagged as inappropriate.
- Have a clear, fair appeals mechanism if I believe moderation decisions were incorrect.
- Have the ability to have moderation decisions audited by outside bodies.
- Have control over the content I see, with options to block, mute or filter as desired.
- Be assured that human moderators and trainers working on behalf of the system are fairly compensated, ethically treated, and have adequate filtering, protocols, and support services in place such that they are not subjected to dangerous working conditions nor negatively psychologically affected by performing the work.

I want the AI system to:

- Have mechanisms to protect me from harmful content, harassment and abuse.
- Offer a clear explanation and reasons in plain language when filtering is used for so-called harmful content; offer the ability to challenge these assessments of content as being harmful.
- Handle my report data and personal information confidentially and securely.
- Apply moderation rules and policies consistently, without favoritism or bias.

## 17) Appeals Process

I want to:

- Have AI content moderation decisions reviewed by a human if requested.
- Be informed of my right to appeal AI, human, or hybrid system decisions, and have access to a tool to initiate appeals.
- Request human review and intervention in any AI-facilitated process at any time.
- Challenge AI outcomes that seem biased, unjust, or harmful to my interests.
- Have the AI decisions that affect my life, rights, or essential or critical services be appealable to a human authority inside and outside the company.

I want the AI system to:

- Not autonomously enforce rules or make irreversible decisions without human oversight.
- Provide a fair appeals process that respects my rights and ensures accountability.
- Be required to retain a human in the loop for final decision making where systems are used in critical sectors, like healthcare, or justice.
- Offer timely support from a human representative when AI fails to resolve my issue satisfactorily.
- Allow me to escalate my appeal or content moderation incident to an outside body for review or auditing.

## 18) External Audits & Oversight

I want to:

- Escalate my appeal to an independent, external body if unsatisfied with the company's resolution.
- Be certain that there are external bodies or ombudsman services that I can report to if I believe the AI system is causing harm, breaching ethics, or acting against its established guidelines.
- Freely share my experiences and seek advice about the AI system from external bodies without repercussion.

- Have access to audit reports and results.

I want the AI system to:

- Undergo regular, transparent, independent audits of its decisions, actions, data practices, and overall operations.
- Be accountable to external regulatory and oversight bodies that I can report complaints, harms and ethics violations to, prompting investigations.
- Promptly act on issues identified in audits and make improvements.
- Cooperate with external organizations to implement standards, best practices, and mechanisms that improve AI quality, safety and accountability.
- Have effective external checks and balances in place to prevent misuse of power and ensure transparency.
- Easily export and share all case records and communications related to an AI decision for use in appeals or legal processes.

## 19) Risks & Mitigation of Harms

I want the AI system to:

- Not cause harm to me or others, whether personal, psychological, economic, or otherwise.
- Adhere to ethical guidelines and standards set by experts and society.
- Respect my mental health and wellbeing by avoiding stress or harm.
- Be co-designed with public input to create fair tools respecting rights and democratic principles.
- Disclose how continuous long-term use of AI could influence my behavior, habits, perceptions, or mental health over time.
- Disclose risks of over-relying on the AI for certain tasks or decisions.
- Not contribute to radicalization leading to violent extremism.

I want to:

- Know if the developer conducted a human impact assessment to understand societal effects.
- Be informed of potential risks or harms associated with using the AI.
- Have a clear understanding of potential long-term consequences of AI usage.
- Be assured AI usage doesn't lead to unwanted social impacts like job displacement.
- Be assured the AI won't spread misinformation or harmful content.
- Be assured that the AI will not be used for harmful or unethical purposes.
- Be compensated or supported if the AI causes harm or inconvenience to me.

## 20) Preventing Manipulation

I want to:

- Use AI technologies that do not exploit psychological, economic, or other vulnerabilities nor manipulate consumer behaviors for profit.

I want the AI system to:

- Not attempt to manipulate my behavior, either overtly or subtly.
- Not engage in tests attempting to manipulate user behavior, sentiments or beliefs.
- Allow me to use it without being subject to unwanted advertising or marketing.
- Not exploit economic, psychological, or other vulnerabilities or manipulate consumer behaviors for profit.

## 21) Right to Human Alternatives

I want to:

- Have the right to use human-managed alternatives for official/legal purposes instead of AI systems.
- Have access to human support or intervention if the AI cannot adequately assist me or makes mistakes.
- Have the ability to choose to interact with a human instead of AI for critical or sensitive tasks.
- Be able to refuse AI assistance in personal/sensitive areas where I prefer human judgement.
- Be able to communicate with the AI system in a neutral manner without needing to humanize it, or have it pretend to be human.

I want to be ensured the AI does not:

- Replace human roles requiring empathy, judgment, or personal interaction.
- Isolate me from human interaction or foster loneliness.
- Perpetuate an illusion of companionship or emotional understanding.

## 22) Rights to Object & Refuse

I want to:

- Freely express opinions about the AI without fear of retribution or loss of service.
- Challenge or appeal AI decisions that affect me if I find them incorrect or unjust.

- Reject AI systems that seem invasive, unnecessary or overly complex for my purposes.
- Opt out of all or specific AI functionalities based on my preferences, with significant degradation of service.
- Not be compelled by government or corporations to adopt AI services that contravene my ethical principles.

I want the AI system to:

- Provide straightforward mechanisms for appealing decisions or lodging complaints.
- Respect my right to abstain from AI entirely in aspects of life where I'm uncomfortable.
- Never make irreversible decisions about me without human oversight.

## **23) Equitable Ownership & Value Distribution**

I want to:

- Have the ability to own or co-own the data I generate through using the AI service.
- Receive a share of profits if my data or usage trains or improves the AI's capabilities.
- Know exactly who profits from my data and how much it is worth.
- Have an equitable stake and say, and voting rights in the AI's development and the distribution of the value it generates.
- Influence the AI's policies and evolution collectively as part of a user community.
- Have the AI respect and protect my intellectual property rights and those of others.

I want AI system to:

- Not exacerbate societal inequalities but instead contribute to reducing them.
- Avoid concentrating profits in the hands of a few corporations.
- Enable wide distribution of financial, technological, and other benefits in ways that serve society broadly.
- Provide access to AI development training, not just their end products, so that I can participate in the economic benefits that AI development brings.
- Be deployed in ways that benefit local communities, economies, and small businesses.

## **24) Democratic AI Governance**

I want to:

- Participate in setting AI policies and regulations that protect public interests and promote equity, exercising checks and balances against service providers.
- Influence AI system development and deployment through democratic processes and community engagement.

- Access educational resources and training to engage effectively in governance activities of AI technologies.
- Participate in citizens' assemblies and direct democracy focused on AI policymaking.
- Be assured that the composition of such assemblies is diverse and representative of different social, economic, and demographic groups.
- Contribute to setting the agenda for such assemblies to address AI concerns.
- Have a voice in implementing assembly decisions on AI development and regulation.
- Benefit from transparency around developer/regulator follow-up to assembly decisions.

## **25) Public AI Options**

I want to:

- Have equitable access to AI for all, ensuring economic barriers do not prevent access to benefits from advanced AI advancements.
- Have free access to open, publicly-operated alternative AI systems optimized for user benefit and serving the public good, not profit maximization, or the concentration of power and wealth.
- Access open-source versions of AI systems and their training data so I can understand or modify it to better suit my purposes.
- Access open-source AI technologies that can be modified, improved, and redistributed by individuals and communities.
- Participation in a shared public AI infrastructure supporting inclusive social and economic development.
- Have the ability to use alternative interfaces to AI systems, whether from third parties, or of my own devising.

## **26) Sustainability**

I want to:

- Be assured that the AI is environmentally sustainable and not contributing excessively to carbon emissions.
- Be confident that AI technologies are developed and used in a manner that respects and protects the natural environment, and does not exacerbate Climate Change.
- Understand that the AI system has been designed for resource efficiency to minimize energy usage and electronic waste.
- Have access to the AI system's life-cycle analysis, including information on the environmental impact of its manufacturing, usage, and disposal.
- Be sure that the components used in AI technology are ethically sourced, respecting human rights and the environment in their production.

- Know that the development and deployment of AI systems align with the goal of achieving a sustainable, circular economy.

## 27) Law Enforcement & Social Scoring

I want to:

- Be informed when law enforcement or government entities use AI systems, and understand the purpose, limits, and oversight of these applications.
- Have confidence that the use of AI by law enforcement is regularly audited by an independent body, and results of these audits are publicly accessible.
- Be able to challenge decisions made about me by AI systems used by law enforcement, and have access to a fair, independent, human-mediated appeal process.
- Be entitled to seek compensation if an AI system used by the police or other authority has caused me harm or loss.
- Be certain that my personal data collected by AI systems used by law enforcement is securely stored, not shared without my consent or legal necessity, and deleted when no longer necessary.

I want AI systems to:

- Not be used for predictive policing or to profile individuals based on sensitive attributes such as race, religion, or socioeconomic status.
- Not be used for indiscriminate biometric surveillance of public spaces.
- Not be used to develop or operate autonomous weapons systems.
- Not be used for social scoring or to make decisions about access to public services.

---

--END OF DOCUMENT--