

Joint Investigation Report PIPA2026-IR-01

Joint Investigation of OpenAI OpCo, LLC by the Office of the Privacy Commissioner of Canada, the Commission d'accès à l'information du Québec, the Office of the Information and Privacy Commissioner for British Columbia, and the Office of the Information and Privacy Commissioner of Alberta

May 6, 2026
OpenAI OpCo, LLC

OPC PIPEDA-045141; CAI P23-93177; OIPC BC 1032001-S; OIPC AB 030562

Table of contents

- Background*..... 3
 - Initiated complaints 5
 - Issues 6
 - Methodology..... 7
- Analysis*..... 8
 - Jurisdiction 8
 - The Objectives of the Acts 12
 - Technical background..... 14
 - OpenAI’s collection, use and disclosure of personal information 17
 - Issue 1: Did OpenAI collect, use and disclose personal information for an appropriate purpose? 20
 - Issue 2: Did OpenAI obtain valid consent and meet its obligation to inform individuals with respect to its collection, use and disclosure of personal information? 33
 - Issue 3: Was OpenAI sufficiently open about its models? 81
 - Issue 4: Did OpenAI take reasonable steps to ensure that the information it generates about individuals is as accurate, complete, and up-to-date as is necessary for the purposes for which it is to be used? 85
 - Issue 5: Did OpenAI provide individuals with the ability to obtain access to, correct and delete their personal information? 94
 - Issue 6: Did OpenAI establish appropriate retention and disposal procedures for the personal information that it collects, uses and discloses? 105
 - Issue 7: Did OpenAI meet its accountability requirements in respect of the personal information under its control? 108
- Recommendations*..... 111
- OpenAI’s response* 115
- Conclusion* 117
- Appendix A – Summary of current key mitigation measures implemented by OpenAI at the various stages of its models’ development and deployment* 121

Background

1. This report of findings examines **OpenAI OpCo, LLC’s (referred to in this report as either “OpenAI” or “the Respondent”)** compliance with Canada’s *Personal Information Protection and Electronic Documents Act* (“PIPEDA”), Quebec’s *Act Respecting the Protection of Personal Information in the Private Sector* (“Quebec’s Private Sector Act”), British Columbia’s *Personal Information Protection Act* (“PIPA-BC”) and Alberta’s *Personal Information Protection Act* (“PIPA-AB”) – referred to collectively as the “Acts”.

OpenAI

2. Parent company, OpenAI, Inc., is an Artificial Intelligence (“AI”)¹ research and deployment company registered in the state of Delaware in the United States. It was founded as a non-profit organization in 2015 with the stated goal of “ensuring that artificial general intelligence benefits all of humanity”.² It has no shareholders and is controlled by a Board of Directors.
3. The Respondent, the operating company that provides OpenAI’s products to end users, is a subsidiary of OpenAI, Inc. It was also registered in Delaware, three years later, in 2018. It is headquartered in San Francisco, California, United States.
4. In 2019, OpenAI, Inc. created another subsidiary, a for-profit enterprise with a “capped profit” structure, OpenAI LP, to raise the capital required to pursue the development of its technology.³ Shortly after, OpenAI LP entered into a strategic partnership with Microsoft.⁴
5. On October 28, 2025, OpenAI announced the completion of the recapitalization of its for-profit enterprise, which is now a public benefit corporation, called OpenAI Group PBC. The non-profit, OpenAI, Inc., is now called OpenAI Foundation and continues to control the for-profit enterprise.⁵

¹ The Government of Canada’s [Directive on Automated Decision-Making](#) defines AI as information technology that performs tasks that would ordinarily require biological brainpower to accomplish, such as making sense of spoken language, learning behaviours or solving problems.

² See [OpenAI’s website](#).

³ Before October 28, 2025, OpenAI LP was fully controlled and governed by OpenAI, Inc. via another entity OpenAI GP, LLC. The ‘capped profit’ structure placed a limit on economic returns for investors and employees. According to OpenAI, returns for its first round of investors were capped at 100x their investment. See [OpenAI’s announcement](#), March 11, 2019.

⁴ [OpenAI forms exclusive computing partnership with Microsoft to build new Azure AI supercomputing technologies](#), Microsoft, July 22, 2019; [OpenAI and Microsoft extend partnership](#), OpenAI, January 23, 2023; [Microsoft’s \\$13 billion bet on OpenAI carries huge potential along with plenty of uncertainty](#), CNBC, April 8, 2023.

⁵ See [Built to benefit everyone](#), OpenAI, October 28, 2025. Under United States corporate law, a benefit corporation (or in certain states, a public benefit corporation) is a type of for-profit corporate entity whose goals include making a positive impact on society.

6. According to OpenAI, as of March 31, 2025, the OpenAI group of companies was valued at US \$300 billion⁶. In October 2025, several media outlets reported that OpenAI had reached a US \$500 billion valuation following a secondary share sale.⁷ More recently, OpenAI announced US \$110 billion in new investment at a US \$730 billion pre-money valuation.⁸
7. OpenAI performs AI research and develops generative AI models. Generative AI is a subset of AI. Generative AI models can produce content such as text, audio, code, videos and images. This content is generated based on information that the user inputs, called a “prompt”, which is typically a question or short instructional text (e.g., “Write me a wedding speech delivered by a best man” or “Tell me about [a famous person]”). The prompt can also include images of what the user is looking for.
8. OpenAI provides the option of either free or paid access to its models, which are used by individuals and organizations (for Enterprise, Team and Education applications).⁹ Since its creation, the company has launched several products, including ChatGPT, described in more detail below (which produces text from text, images or voice prompts)¹⁰, DALL-E (which produces images from text or image prompts)¹¹ and Sora (which generates videos from text instructions).¹² Our investigation focused exclusively on specific versions of ChatGPT, as detailed further in paragraph 16.

ChatGPT

9. ChatGPT was released in November 2022. It is a conversation-style service that can respond to users’ prompts and create various types of content, including content as varied as articles, computer code and poems.
10. ChatGPT is powered by a foundational Large Language Model (“LLM”). LLMs are extremely large, complex machine learning systems capable of routinely generating elaborate, plausible-sounding – but not necessarily factually accurate – content in response to queries on virtually any topic.

⁶ See [New funding to build towards AGI](#), OpenAI, March 31, 2025.

⁷ See [OpenAI wraps \\$6.6 billion share sale at \\$500 billion valuation](#), CNBC, October 2, 2025.

⁸ See [Scaling AI for everyone](#), OpenAI, February 27, 2026 and [OpenAI’s New \\$110B Raise At A \\$840B Valuation Marks The Largest Venture Deal Ever](#), Crunchbase news, February 27, 2026.

⁹ The [ChatGPT Team](#) subscription is designed for organizations and businesses wishing to adopt ChatGPT for use among their teams. [ChatGPT Enterprise](#) is aimed at global companies and includes and offers enterprise-grade security, privacy and deployment tools. [ChatGPT Edu](#) allows universities to deploy AI to students, faculty, researchers and campus operations. For more information, see [OpenAI’s pricing](#).

¹⁰ See [ChatGPT can now see, hear, and speak](#), OpenAI, September 25, 2023.

¹¹ See for example [DALL·E 3](#), OpenAI. DALL·E 3 is built natively on ChatGPT. When prompted with an idea, ChatGPT will automatically generate tailored, detailed prompts for DALL·E 3.

¹² OpenAI has [announced](#) that it would discontinue Sora in 2026.

11. At the time of its release, ChatGPT was powered by an LLM called GPT-3.5. In March 2023, OpenAI introduced GPT-4, which was enhanced with additional image input capabilities in September 2023. Since then, OpenAI has regularly upgraded, and released new versions of, its models.¹³

OpenAI's economic model

12. OpenAI's economic model with respect to ChatGPT is primarily based on two major business lines:

- i. Direct access to ChatGPT, and other products (which are outside the scope of our investigation), via a free or premium subscription, including ChatGPT Team, Edu and Enterprise versions, which offer advanced capabilities and customization options. In April 2024, OpenAI announced that users had the option to use the free version of ChatGPT without an account.¹⁴
- ii. An Application Programming Interface ("API")¹⁵ platform, which allows API customers to build applications powered by OpenAI's models. The APIs enable these customers to integrate the capabilities of OpenAI's AI models into their own applications, products or services, which they can then make available to their own end users and customers. OpenAI charges for API usage based on consumption (i.e., "pay for what you use" approach).¹⁶

13. According to OpenAI, as of January 2024, there were several million monthly active ChatGPT users¹⁷ and several hundred thousand paid subscribers (i.e., Chat GPT Plus) in Canada. Quebec, British Columbia and Alberta all had a significant user base.

Initiated complaints

14. In April 2023, the Office of the Privacy Commissioner of Canada ("OPC") launched an investigation into OpenAI in response to a complaint alleging that the organization had collected, used and disclosed the complainant's personal information without consent.

15. In May 2023, given the significant privacy impact of Generative AI and its relevance to all Canadians, the OPC decided to close the initial complaint and, along with the Commission d'accès à l'information du Québec ("CAI"), the Information and Privacy Commissioner for British Columbia ("OIPC-BC"), and the Information and Privacy Commissioner of Alberta ("OIPC-AB"),

¹³ See [OpenAI's models' overview](#). e.g., GPT-3.5 Turbo, GPT-4 Turbo, etc. At the time of writing this report, OpenAI has released new models, including [GPT-4o](#), [OpenAI o1-mini](#), [GPT-5](#) and [GPT-5.5](#), which the Offices have not tested.

¹⁴ [Start using ChatGPT instantly](#), OpenAI, April 1, 2024.

¹⁵ An API is a set of rules or protocols that enables software applications to communicate with each other to exchange data, features and functionality.

¹⁶ For more information, see [OpenAI's pricing](#).

¹⁷ 'Monthly active users' means the number of unique users who have used ChatGPT within the past month.

collectively referred to as “the Offices”, initiated investigations pursuant to section 11(2) of PIPEDA, section 81 of Quebec’s Private Sector Act, section 36(1)(a) of PIPA-BC, and section 36(1)(a) of PIPA-AB respectively. The Offices decided to conduct the investigations jointly in order to leverage their respective expertise and resources, while avoiding duplication of efforts for the Offices and OpenAI.

16. The Offices specifically focused on ChatGPT and the underlying models that powered it at the time that the investigation was launched (i.e., GPT-3.5 and GPT-4, excluding recent releases). The Offices did not assess later models (we did, however, consider the adequacy of new measures implemented by OpenAI in response to our Preliminary Report of Investigation), or OpenAI’s other AI services (such as image or video generation). However, the findings in this report may still be relevant to them if their development and deployment process is similar to that used for GPT-3.5 and GPT-4 (e.g., models designed to mimic human conversations, using training techniques such reinforcement learning, which is discussed later in this report).
17. Our scope did not extend to the unlimited potential applications of the tool by OpenAI’s clients (e.g., API customers, developers of GPTs¹⁸, individual users).

Issues

18. Our investigation sought to determine whether OpenAI:
 - i. collected, used and/or disclosed personal information for purposes that a reasonable person would consider appropriate in the circumstances, and limited this collection to information which was necessary for these purposes;¹⁹
 - ii. obtained valid consent for the collection, use and disclosure of the personal information of individuals based in Canada via or in relation to ChatGPT;
 - iii. fulfilled its obligations to be open and transparent;
 - iv. took reasonable steps to ensure that the information it generates about individuals is as accurate, complete, and up-to-date as is necessary for the purposes for which it is to be used;
 - v. provided individuals with the ability to obtain access to, and correct, their personal information;

¹⁸ GPTs are custom versions of ChatGPT that anyone can build for their specific purposes. See OpenAI’s post [Introducing GPTs](#) for more information.

¹⁹ Throughout this report, the term “appropriate purpose” will be considered inclusive of “reasonable purpose” under PIPA-AB and PIPA-BC and “legitimate, real and important purpose” under Quebec’s Private Sector Act. See s. 5(3) of PIPEDA; ss. 2, 11, 14, and 17 of PIPA-BC, [SBC 2003 c 63](#); and ss. 2, 3, 11, 16 and 19 of PIPA-AB, [SA 2003 c P-6.5](#). In addition, we note that s. 4 of Quebec’s Private Sector Act, [CQLR c P-39.1](#) requires that “[a]ny person carrying on an enterprise who, for a serious and legitimate reason, collects personal information on another person must determine the purposes for collecting the information before doing so.”

- vi. fulfilled its obligation to establish appropriate retention and disposal procedures for the personal information that it collects, uses and discloses; and
- vii. was accountable for the personal information under its control.

Methodology

19. To conduct the investigation, the Offices considered information from a variety of sources, including:
- i. extensive written representations provided to the Offices by the Respondent through its legal counsel. This included descriptions of the privacy protective measures and tools implemented by OpenAI at the various phases of its models' development and deployment, as well as the results of OpenAI's internal evaluations of these measures and tools;
 - ii. information that the Offices gathered during interviews with several OpenAI employees, conducted at the Respondent's headquarters in San Francisco, United States and virtually, as well as at the offices of the OPC in Gatineau, Canada;
 - iii. internal testing on Chat-GPT (versions 3.5 and 4) from the perspective of a user, which the investigative team conducted in the OPC's technical laboratory between November 2023 and May 2024. This included testing the account creation process, examining the content and frequency of OpenAI's notification about accuracy, using OpenAI's export data tool, gaining a better understanding of OpenAI's threshold for determining who is a public individual, and more generally interacting with ChatGPT to replicate the user experience; and
 - iv. information that the Offices gathered and analyzed from publicly accessible sources concerning issues relevant to the investigation (e.g., media articles, studies published by OpenAI or independent AI experts). We did not rely on these sources as evidence to support our findings but rather referenced them to illustrate certain practices and provide context, where relevant.
20. Upon completion of the evidence-gathering phase of our investigation, the Offices issued a Preliminary Report of Investigation ("the Preliminary Report"), which set out the rationale for our preliminary findings, identified a number of recommendations to bring OpenAI into compliance with the Acts, and invited OpenAI to respond. We also met on various occasions with OpenAI to provide an opportunity for the company to ask any questions it may have on the Preliminary Report and discuss any perceived challenges in responding to our recommendations.
21. In its final written response, OpenAI argued that it was, through a combination of existing practices and associated communications, compliant with the Acts in most respects. OpenAI also provided new information and explanations about measures that it has recently implemented in

relation to the development and deployment of ChatGPT. These measures were not applied to GPT-3.5 and 4, but rather exclusively to subsequent versions of the models.

Analysis

Jurisdiction

22. As indicated previously, the Respondent is incorporated in the U.S. That said, in the course of its commercial activities, OpenAI collects, uses, and discloses personal information of individuals who use ChatGPT across Canada, including of users located in the provinces of Alberta, British Columbia, and Quebec, as explained in the next section of this report.
23. Nevertheless, OpenAI challenged the Offices' jurisdiction (both in whole and in part) on the grounds that ChatGPT was not released in Canada until November 30, 2022, and that OpenAI had no establishment or employees in Canada prior to that date. OpenAI also took the position that the Acts do not apply to outputs generated and used by ChatGPT end users, where these outputs are created for personal and domestic purposes. Finally, OpenAI specifically challenged the jurisdiction of the OIPC-BC.
24. The Acts under which this investigation was conducted apply to organizations that, in the course of a commercial activity, collect, use, and disclose the personal information of individuals within each province. As such, each office undertaking this investigation has determined that it has the jurisdiction to investigate and make recommendations or orders²⁰ related to OpenAI's handling of personal information within their respective jurisdiction of responsibility, whether provincial or federal.²¹
25. In addition, PIPEDA applies to organizations outside of Canada where a "real and substantial connection" to Canada exists.²² In our view, the circumstances in this matter clearly demonstrate that a real and substantial connection to Canada exists. In coming to this conclusion, we considered the following factors:
 - i. OpenAI offers its services in Canada and has monthly active ChatGPT users, including paid subscribers (i.e., ChatGPT Plus) in Canada.

²⁰ Unlike its provincial counterparts, the OPC does not have order-making powers; it can make non-binding recommendations only.

²¹ The Offices' jurisdictional analysis is consistent with the approach taken by the Ontario Superior Court in *Toronto Star Newspapers Limited v. OpenAI Inc.*, [2025 ONSC 6217 \(CanLII\)](#).

²² *Lawson v. Accusearch Inc.*, 2007 FC 125, paras. 38-51; *A.T. v. Globe24h.com*, 2017 FC 114 (CanLII), [2017] 4 FCR 310, paras 50-64, citing *Society of Composers, Authors and Music Publishers of Canada v. Canadian Assn. of Internet Providers*, 2004 SCC 45, [2004] 2 SCR 427 at paras 54-63.

- ii. OpenAI's Terms of Service for ChatGPT are applicable to users in Canada and address consent to the collection, use, and disclosure of personal information, and associated rights of access and correction.
- iii. Users located in Canada can visit the OpenAI website and use ChatGPT.
- iv. We note that OpenAI's activities take place exclusively through a website or app. As referenced in paragraph 54 of *A.T. v. Globe24h.com*²³, a physical presence in Canada is not required to establish a real and substantial connection when considering websites under PIPEDA, as telecommunications occur "both here and there."
- v. OpenAI's operations necessitate the transmission and receipt of personal information between Canada and the United States, both when collecting information and disclosing it through ChatGPT.
- vi. OpenAI has collected, used and disclosed the personal information of users in Canada, or derived from Canadian sources, to develop and deploy ChatGPT.

26. Similarly, as cited in the LifeLabs Investigation Report²⁴, Alberta's Privacy Commissioner has jurisdiction to conduct compliance investigations under PIPA-AB. Furthermore, an organization that collects, uses or discloses personal information in Alberta must comply with Alberta privacy legislation, and this includes all aspects of compliance, as provided by section 36(1)(a) of PIPA-AB. If an organization collects, uses or discloses personal information in Alberta, practices throughout the organization must comply with PIPA-AB.²⁵ Finally, as cited in the Clearview decision²⁶, Alberta's Privacy Commissioner has jurisdiction over Clearview because Clearview chose to do business in Alberta and collects, uses, and discloses personal information of Albertans, some of which is hosted on websites with servers in Alberta.

27. Like the other Offices, the CAI does not accept the argument raised by OpenAI regarding jurisdiction and endorses the reasons set out in paragraph 25 of this report.

28. More specifically, section 1 of Quebec's Private Sector Act establishes the basis for the application of the Act and the jurisdiction of the CAI. This section specifies that the purpose of Quebec's Private Sector Act is to establish, for the exercise of the rights conferred by articles 35 to 40 of the Civil Code of Quebec²⁷ concerning the protection of personal information, particular rules with respect to personal information relating to other persons which a person collects, uses, or communicates to third persons in the course of carrying on an enterprise within the meaning of section 1525 of the Civil Code of Quebec²⁸.

²³ *A.T. v. Globe24h.com*, 2017 FC 114 (CanLII), [2017] 4 FCR 310.

²⁴ See LifeLabs, [P021-IR-04](#), paragraph 20.

²⁵ See LifeLabs, [P021-IR-04](#), paragraph 40.

²⁶ See *Clearview AI Inc v. Alberta (Information and Privacy Commissioner)*, [2025 ABKB 287](#). Paragraph 61.

²⁷ [CQLR c CCQ-1991](#).

²⁸ *Ibid.*

29. Open AI collected and used personal information concerning people in Quebec to train its GPT-3.5 and 4 models, which is a real and important link with Quebec²⁹.
30. Justice Abella stated in *Google Inc. v. Equustek Solutions Inc.* that “[t]he Internet has no borders — its natural habitat is global”. Therefore, considering the nature of OpenAI’s activities, the CAI finds that the absence of any establishment or employees in Quebec does not in itself preclude the application of Quebec’s Private Sector Act.
31. The invasion of privacy that may result from the collection and use of the personal information of Canadians, and more specifically people in Quebec, occurs at the place of residence of the individuals concerned by this information, and the place of residence constitutes a sufficient connecting factor in this case³⁰.
32. In addition, the CAI considers that, for the operation of its business, and more specifically for the purposes of supporting the functionality of its GPT-3.5 and 4 models, OpenAI continues to hold and use personal information concerning Quebec residents.
33. Ultimately, considering the nature of the products or services offered by OpenAI and the context in which the collection, use, and disclosure of personal information have been and continue to be carried out, the CAI considers that OpenAI is subject to Quebec’s Private Sector Act regarding the collection, use, and disclosure of personal information concerning Quebec residents.

OpenAI’s challenge to OIPC-BC’s jurisdiction

34. In its response to the Preliminary Report, OpenAI argued that OIPC-BC and the OPC cannot both have jurisdiction over the subject matter of this investigation. In support of this argument, OpenAI raised s. 3(2)(c) of PIPA-BC, which states:

(2) This Act does not apply to the following:

... (c) the collection, use or disclosure of personal information, if the federal Act applies to the collection, use or disclosure of the personal information...³¹

35. This argument has been thoroughly addressed in previous joint investigations that included the OIPC-BC.³²

²⁹ *Libman v. The Queen* [1985] [2 S.C.R. 178](#).

³⁰ *Moran v. Pyle National (Canada) Ltd* [1975] [1 S.C.R. 393](#), p.405 and *Doan c. Clearview AI Inc.* [2024 QCCS 3968](#) (in French only), paras. 85 and 96.

³¹ Section 1 of PIPA-BC defines the “federal Act” to mean PIPEDA.

³² For example, see: [Joint Investigation of TikTok Pte Ltd.](#), Investigation Report 25-02, 2025 BCIPC 31 at paras 20-24, and [Joint investigation of Facebook, Inc. by the Privacy Commissioner of Canada and the Information and Privacy Commissioner for British Columbia](#), PIPEDA-036162 / OIPC P17-72561, at para 47, which was confirmed in a *de novo* hearing in *Facebook, Inc. v. Canada (Privacy Commissioner)*, 2023 FC 534 (CanLII), at [paras 86-87](#) (Facebook, Inc. was granted leave to appeal by the SCC, and the OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending).

36. Exemption Order SOR/2004-220, issued under PIPEDA, clearly exempts an organization from the application of Part 1 of PIPEDA to that organization’s collection, use, or disclosure of personal information if that activity occurs within British Columbia.³³ Therefore, as an organization, OpenAI’s collection, use, or disclosure of personal information is subject to PIPA-BC instead of Part 1 of PIPEDA to the extent that such activity occurs within British Columbia.
37. Privacy regulation is a matter of concurrent jurisdiction and an exercise of cooperative federalism.³⁴ Cooperative federalism is a core principle of this modern division of powers, and jurisprudence reflects the concurrent operation of statutes enacted by the federal and provincial levels of government.³⁵ PIPA-BC is “designed to dovetail with federal laws” in its protection of quasi-constitutional privacy rights of individuals in British Columbia.³⁶
38. The legislative history of the enactment of PIPEDA,³⁷ PIPA-BC,³⁸ and their interlocking structure support the interpretation that PIPEDA and PIPA-BC operate together seamlessly. Moreover, the Court of Appeal for British Columbia recently confirmed the OIPC-BC’s jurisdiction in the context of a joint investigation of an organization operating across provincial and international borders.³⁹
39. This investigation entails a single organization operating across both jurisdictions with a complex collection, use, and disclosure of personal information. In our view, an interpretation of s. 3(2)(c) that removes the authority of the OIPC-BC in any circumstance where the OPC also exercises authority would be inconsistent with the current approach to privacy regulation in Canada and would frustrate the principle of cooperative federalism.
40. The OIPC-BC therefore concludes that s. 3(2)(c) of PIPA-BC does not preclude the application of PIPA-BC to OpenAI’s collection, use, or disclosure of personal information in British Columbia, nor does it limit the jurisdiction of the OIPC-BC to participate in this investigation in any way.

³³ Organizations in the Province of British Columbia Exemption Order, [SOR/2004-220, at s 1](#).

³⁴ Provincial legislatures legislate under their authority over property and civil rights found in s.92(13) of the *Constitution Act, 1867* and the federal Parliament under its trade and commerce power under s.91(2) and its authority over federal works, undertaking, and businesses. See Reference re Subsection 18.3(1) of the *Federal Courts Act*, 2021 FC 723, para 47; See also Michel Bastarache, *The Constitutionality of PIPEDA: A Re-consideration in the Wake of the Supreme Court of Canada’s Reference re Securities Act* (June 2012).

³⁵ *Liberal Party of Canada v. The Complainants*, [2024 BCSC 814](#) at para 119, citing *Quebec (Attorney General) v. Canada (Attorney General)*, [2015 SCC 14](#) at paras 17–19; *Reference re Impact Assessment Act*, [2023 SCC 23](#) at paras 116, 122, and 216; *Rogers Communications Inc. v. Châteauguay (City)*, [2016 SCC 23](#) at para 38.

³⁶ *Liberal Party of Canada v. The Complainants*, *ibid.*, at para 201.

³⁷ See for example: Parliament of Canada, Official Record of Debates (Hansard), No. 137 at 1215 (19 Oct 1998, Hon. John Manley); Parliament of Canada, Standing Committee on Industry, Minutes of Proceeding at 1545-50 (1 December 1998, Hon John Manley).

³⁸ See for example: British Columbia, Legislative Assembly, Official Report of Debates (Hansard) Vol. 14 No. 12 at 6351-6352 (30 Apr. 2003, Hon. S. Santori) and 6415-6416 (1 May 2003, Hon. S. Santori).

³⁹ *Clearview AI Inc. v Information and Privacy Commissioner for British Columbia*, [2026 BCCA 67](#), leave to appeal to SCC pending.

OpenAI’s challenge to the Offices’ jurisdiction prior to the release of ChatGPT

41. As mentioned above, in its response to the Preliminary Report, OpenAI made representations setting out its view that in the absence of any establishment, employees or other factors giving rise to a real and substantial connection to Canada prior to the release of ChatGPT on November 30, 2022, the Offices cannot have jurisdiction over, and the Acts cannot apply to, OpenAI’s activities before that date.
42. After carefully considering these arguments, we disagree with OpenAI’s assertion and find that several factors confirm that there was a real and substantial connection to Canada prior to November 30, 2022. Specifically:
- i. As acknowledged by OpenAI, the development of ChatGPT prior to its launch involved and relied, in part, on the collection of the personal information of individuals in Canada or derived from Canadian sources (e.g., Canadian-based online platforms). This development, which took place well in advance of the release of ChatGPT in Canada is inextricably connected to the commercial activity of deploying this service.
 - ii. OpenAI has retained the datasets including this personal information and continues to use them for the purpose of training its AI models.

The use of ChatGPT for personal or domestic purposes

43. OpenAI represented that “the Acts do not apply to outputs generated and used by ChatGPT end users where such outputs are created for personal or domestic purposes.” The Offices acknowledge that some of the Acts contain exemptions related to collection, use, or disclosure of personal information for personal or domestic purposes.⁴⁰ However, these provisions operate to exempt “individuals” engaged in personal or domestic use, not “organizations” engaged in commercial activities. Furthermore, for these exemptions to apply, the purpose of the collection, use and disclosure must be *exclusively* personal or domestic.⁴¹ As is clear from paragraphs 12 to 13 of this report, OpenAI’s operation of ChatGPT cannot be credibly characterized as exclusively personal or domestic.

The Objectives of the Acts⁴²

44. Generative AI applications such as ChatGPT have the potential to implicate not only privacy rights but also the right to freedom of expression, which is protected under s. 2(b) of the *Canadian Charter of Rights and Freedoms* (“the Charter”). Subject to limited exceptions, this

⁴⁰ See s. 4(2)(b) of PIPEDA; s. 3(2)(a) of PIPA-BC; and s. 4(3)(a) of PIPA-AB.

⁴¹ *Ibid.*

⁴² This section does not apply to Quebec’s Private Sector Law. The CAI will address arguments on the balancing of rights guaranteed by the Charter in the sections devoted to this topic.

Charter right can apply to any activity that “conveys or attempts to convey meaning”.⁴³ In addition to potentially involving personal information, ChatGPT prompts and generated outputs may be an example of expressive content. Accordingly, any limitations on the operation of ChatGPT may in turn place limits on the Charter-protected right to freedom of expression, including, for example, as it relates to ChatGPT users.

45. In its submissions to the Offices, OpenAI highlighted that ChatGPT can promote education, research, creativity and innovation. OpenAI further submitted that the Offices’ analysis should be informed by Charter values, which the OPC, OIPC-BC and OIPC-AB do not disagree with. In addition to balancing the needs of organizations and individual privacy rights consistent with the statutory purposes of some of the Acts,⁴⁴ the forthcoming analysis is also informed by the Charter values of freedom of expression and privacy. In this regard, the OPC, OIPC-BC and OIPC-AB have adopted the approach taken by the OPC in its recent Report of Findings related to Google search results.⁴⁵ In that report, drawing on Supreme Court of Canada jurisprudence, the OPC noted that when Charter protections are engaged, statutory decision-makers are required to balance the statutory objectives of the legislation they oversee with relevant Charter protections (in particular freedom of expression and privacy).⁴⁶ In this case, the Offices’ assessment of appropriate purposes under the Acts (issue 1) has been informed by Charter values.
46. In addition to Charter values, consistent with the Federal Court of Appeal decision in *Englander v. TELUS Communications Inc.*, the OPC, OIPC-BC and OIPC-AB interpret their respective acts with “flexibility, common sense, and pragmatism”.⁴⁷ This is consistent with the modern approach to statutory interpretation, which states that “the words of an Act are to be read in their entire context and in their grammatical and ordinary sense harmoniously with the scheme of the Act, the object of the Act, and the intention of Parliament”.⁴⁸ The Supreme Court’s recent decision in *Telus Communications Inc. v. Federation of Canadian Municipalities* clarified that the modern approach to statutory interpretation entails interpreting legislation in a “dynamic” manner, as “capable of applying to new circumstances including new technology ... consistent with the legislature’s purpose.”⁴⁹ Finally, the Courts have also recognized PIPEDA as quasi-constitutional legislation and “part of an international movement towards giving individuals better control over

⁴³ See *Irwin Toy Ltd. v. Quebec (Attorney General)*, [1989] 1 SCR 927, at p. 969; *Ford v. Quebec (Attorney General)*, [1988] 2 SCR 712, at para. 59 (“Ford”), cited in : [PIPEDA Findings # 2025-002](#), August 27, 2025.

⁴⁴ See s. 3 of PIPEDA; s. 3 of PIPA-AB; and s. 2 of PIPA-BC. See also *Canada (Privacy Commissioner) v. Facebook, Inc.*, [2024 FCA 140](#) (CanLII) at para. 62 (Facebook, Inc. was granted leave to appeal by the SCC, and the OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending).

⁴⁵ [PIPEDA Findings #2025-002](#).

⁴⁶ See *Doré v. Barreau du Québec*, 2012 SCC 12, (*Doré*) at [paras. 55-58](#). See also *R. v. Clarke*, [2014 SCC 28 \(CanLII\)](#) at para. 16.

⁴⁷ *Englander v. TELUS Communications Inc.*, 2004 FCA 387 (CanLII), at paragraph 46.

⁴⁸ E.A. Driedger, *Construction of Statutes* (2nd ed. 1983), at p. 87, as cited in *Rizzo & Rizzo Shoes Ltd. (Re)*, 1998 CanLII (SCC), [1998] 1 S.C.R. 27, at paragraph 21; see also, Ruth Sullivan, “The Construction of Statutes, 7th Ed. §9.01), and s. 12 of the *Interpretation Act*, [RSC 1985, c I-21](#).

⁴⁹ [2025 SCC 15 \(CanLII\)](#) at para. 36.

their personal information” which is “intimately connected to their individual autonomy, dignity and privacy.” Courts have further acknowledged the fundamental role that privacy plays in the preservation of a free and democratic society.⁵⁰

Technical background

Model training:

47. The operation of an LLM relies on a large number of numerical “weights”, which represent the statistical relationship between different words (or portions of words, which are converted to numerical strings called “tokens”) in different contexts.⁵¹ These weights are determined based on the LLM's processing of training data and may be updated as the LLM is subject to further training.
48. To train the GPT-3.5 and GPT-4 models underlying ChatGPT, OpenAI employed a two-stage process, described below in simplified terms:⁵²
 - i. **Pre-training** (also called “self-supervised learning”):
 - a. During this initial phase, the model gains a general understanding⁵³ of language (or more technically, a general functionality in natural language processing) by analyzing vast amounts of unstructured and tokenized text data, which may include personal information.
 - b. As discussed further in paragraph 51, pre-training datasets are generally comprised of (i) information collected from publicly accessible websites⁵⁴ (noting that OpenAI removes certain limited categories of websites, redundant content, and content that violates its policies) and (ii) data licensed from third parties. In response to our Preliminary Report, OpenAI represented that for both (i) and (ii), it now also uses a tool to detect and mask identifying information about private individuals from pre-training data (see below).

⁵⁰ [Alberta \(Information and Privacy Commissioner\) v. United Food and Commercial Workers, Local 401, 2013 SCC 62 \(CanLII\), \[2013\] 3 SCR 733, 2013 SCC 62 at para. 13 and 19.](#) See also [Reference re Subsection 18.3\(1\) of the Federal Courts Act, \[2021\] 3 FCR 503, 2021 FC 723, \[2021\] at para 39](#), aff'd [Google LLC v Canada \(Privacy Commissioner\), 2023 FCA 200 \(CanLII\), \[2023\] FCJ No 1411, 2023 FCA 200.](#)

⁵¹ Tokens may also represent full words or characters. Tokens are represented by token IDs, which are series of numbers.

⁵² The investigative team, including technical analysts, received highly technical explanations in the form of written submissions and through answers in interviews. We have summarized and simplified those explanations for the purposes of this report.

⁵³ Although we use the term “understand” as a shorthand, we are mindful of the risks of anthropomorphizing AI.

⁵⁴ As discussed further at paragraph 161 and following, the common understanding of “publicly accessible” information is distinct from the definition of “publicly available” information, which is provided by each Act’s regulations, with the exception of Quebec’s Private Sector Act. In this report, when we refer to information from publicly accessible websites or sources, we mean information which is freely accessible on the Internet (e.g., not behind paywalls).

- c. Based on this pre-training data, the model builds a representation of the general statistical relationship between tokens, and from this, learns to continuously predict the next token in a sentence.
- ii. **Fine-tuning** (or Alignment): this phase seeks to further improve the model's performance on specific tasks and domains (e.g., translation, summarization, conversation) by refining its behaviour and the statistical correlations it establishes between tokens. This helps improve the model's ability to answer questions in a way that people find useful, as well as prevent the model from returning a response that may be used in harmful ways (e.g., a response that would constitute hate speech, or one that would contain personal information of a private individual). Fine-tuning involves the use of a subset of data that is gathered through individuals' interactions with ChatGPT, as well as information provided by human trainers (see paragraph 51). Fine-tuning is further divided into substages, including:
 - d. **Supervised learning** – the model is trained on examples of “ideal” behaviour that is written by human trainers to demonstrate the type of responses that the model should provide in response to various prompts it might receive.
 - e. **Reinforcement learning** from human feedback (“RLHF”) – the model is “rewarded” if it responds ethically and appropriately to users' prompts (i.e., with relevant, safe, accurate, unbiased responses) and receives less reinforcement or a “lower reward” if it does not. Concretely, human trainers review multiple LLM responses to the same prompt and rank them in order of most to least appropriate or ethical. As the model “learns” from this feedback, the statistical weights between words are updated.⁵⁵

⁵⁵ For GPT-4, OpenAI used a rule-based reward model to augment the reinforcement learning phase.

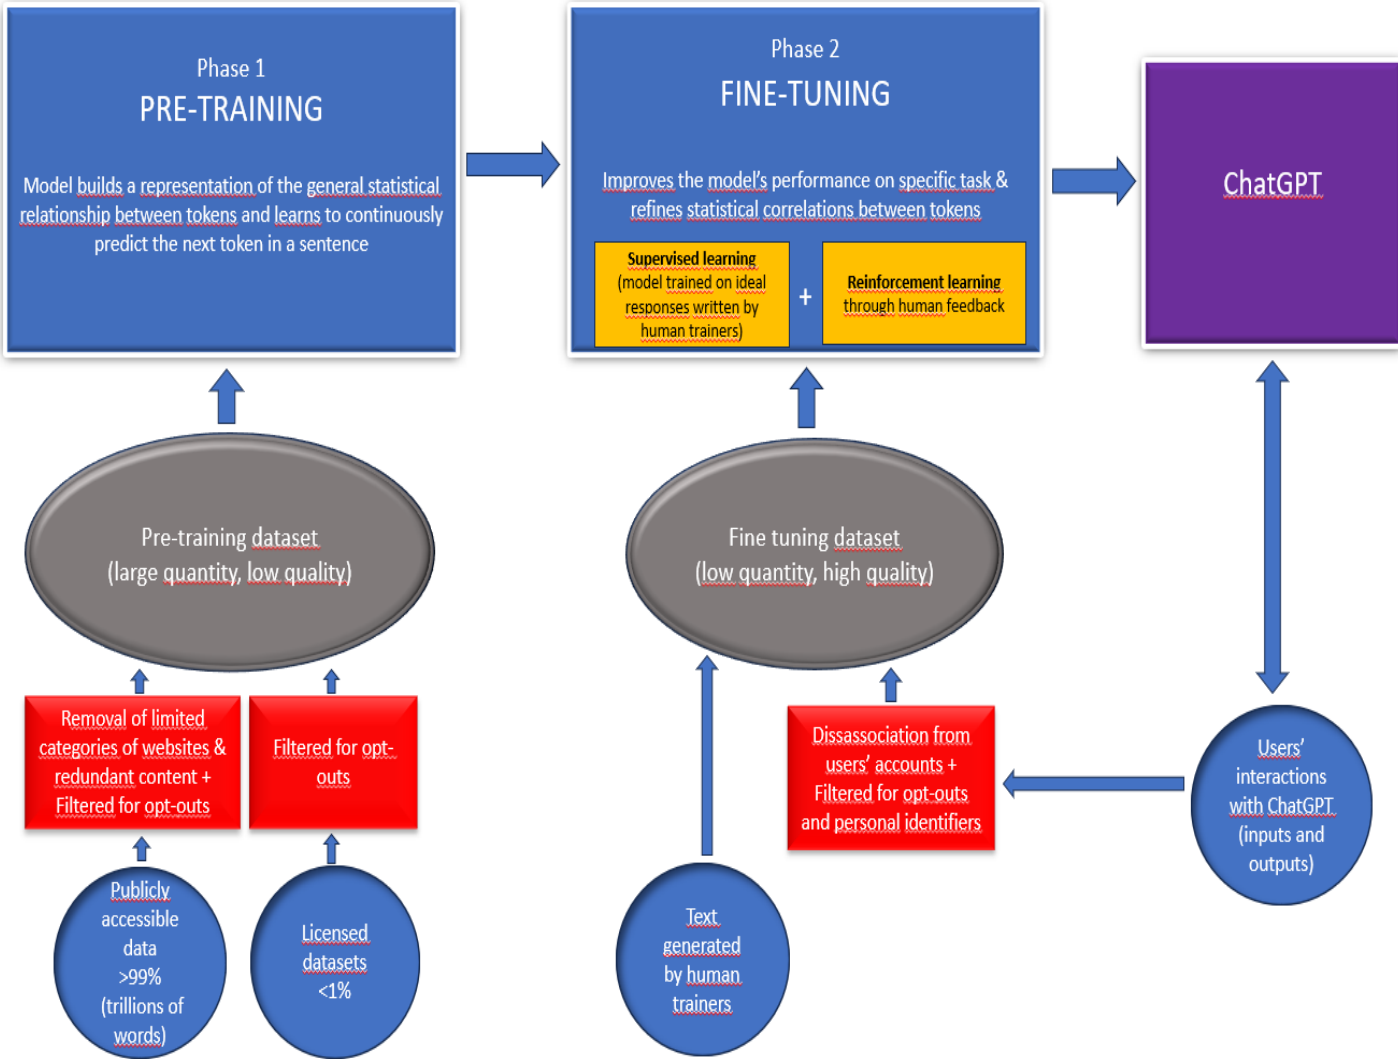


Figure 1 – Model training (GPT-3.5 and 4)

ChatGPT's text generation:

49. The following is a high-level description of the technical process by which ChatGPT generates a response to a user's prompt:

- i. ChatGPT receives text input from users in the form of a "prompt" (defined in paragraph 7).
- ii. The LLM (e.g., GPT-3.5 or ChatGPT-4) breaks down the text input into "tokens".
- iii. To create a response, the input tokens are run through the LLM, which generates a first output token (often, but not always, the statistically most likely). Subsequent output tokens are generated by the LLM, taking into account both the input tokens and the already generated output tokens.

- iv. The series of tokens generated by the LLM is then converted back into human-readable text, which is provided as the output or response to the user's prompt.

OpenAI's collection, use and disclosure of personal information

50. Citing security, confidentiality and operational considerations, OpenAI did not grant our request to access and examine its systems. However, in its representations to the Offices, OpenAI acknowledged that, in the course of its activities, it collects, uses and discloses personal information. The Offices therefore consider that this acknowledgement constitutes evidence to substantiate that the information collected, used and disclosed by OpenAI contains personal information within the meaning of the Acts.

51. Specifically, OpenAI represented that in order to develop or train its models and facilitate users' interactions with ChatGPT, it collects data from four primary sources of information. Each source may include personal information:⁵⁶

- i. **Information from "publicly available Internet sources,"** which, according to OpenAI, currently represents the vast majority of its training datasets.⁵⁷ OpenAI collects this information either:
 - a. from third parties such as Common Crawl⁵⁸ or Wikipedia, which have already gathered and made that information available. OpenAI represented that it does not circumvent paywalls or account-protected websites when collecting this data; or
 - b. via its GPTBot. This web crawling tool crawls and scrapes the content of websites across the Internet. Website owners can, however, choose to restrict or limit GPTBot access.⁵⁹

⁵⁶ OpenAI represented that it also collects (i) account information (e.g., first and last name, contact details); (ii) communication information (e.g., contact details provided by an individual when contacting OpenAI); (iii) individual's social media handle (i.e., when an individual voluntarily interacts with OpenAI on social media); and (iv) log data, usage data and device information, which is collected automatically when a user uses OpenAI's products (e.g., IP address, browser type and settings, time zone, usage analytics). However, OpenAI noted that these categories of information are not used to develop or train its models.

⁵⁷ The media has reported that several groups of copyright owners including newspapers, writers, visual artists and music publishers have sued OpenAI and other major tech companies over the alleged misuse of their work to train generative AI systems. OpenAI contests these allegations. This, however, is not a privacy issue under the Acts and therefore falls outside the scope of this investigation.

⁵⁸ [Common Crawl](#) is a non-profit founded in 2007 that crawls the web and "make(s) wholesale extraction, transformation and analysis of open web data accessible to researchers".

⁵⁹ OpenAI started preparing datasets to train GPT-3.5 months before it launched ChatGPT. The GPT bot was launched in August 2023, so the ability for websites to opt-out of being scraped by the GPT bot was not in place when OpenAI collected the majority of the publicly accessible information used to train GPT-3.5 and GPT-4.

- ii. **Information that it licenses from third parties**, including from various media outlets, a large stock image vendor, and other sources of specialized knowledge.⁶⁰ OpenAI represented that it partners with these content providers to ensure the inclusion of high-quality content in its training datasets, including on specialized topics such as science and mathematics.
- iii. **User interactions with ChatGPT** (i.e., model input and output, image and file uploads, feedback provided by the user to OpenAI regarding whether ChatGPT’s response was helpful). Users can choose not to have this data used for model training, as further explained in paragraph 303.⁶¹
- iv. **Conversations generated by human AI trainers** (both OpenAI’s employees and contractors). As discussed in paragraph 48, these trainers create conversations by writing queries and ideal responses to fine-tune the model. They also evaluate and rank different model-generated responses based on their “quality, safety, and relevance”.

52. OpenAI stated that the inclusion of personal information in its training datasets is incidental to the broader goal of obtaining a large and varied body of text necessary to effectively train its models. It further submitted that it has mitigation measures in place to limit, to the extent possible, the presence of personal information in the training datasets and model outputs and to minimize associated privacy risks. In response to our Preliminary Report, OpenAI also represented that it continually researches and develops safety improvements and iterates on technical and process enhancements for training AI models. This includes privacy-enhancing techniques that reduce the processing of personal information by detecting and filtering personal identifiers in training datasets, thereby enabling models to learn about language and develop intelligence without learning from the specific masked information (these various mitigation measures are described below in relevant sections of the report). At the same time, OpenAI stated that it is not feasible, nor desirable, to completely remove personal information from training text, as models need to learn how such information fits into a sentence, to be able to respond to users’ prompts.

⁶⁰ OpenAI has made some of these licensing agreements public; e.g., with The Associated Press (see [ChatGPT-maker OpenAI signs deal with AP to license news stories](#), AP, July 13, 2023), with International news organizations *Le Monde* and *Prisa Media* (see [Global news partnerships: Le Monde and Prisa Media](#), OpenAI, March 13, 2024), Reddit (see [OpenAI and Reddit Partnership](#), OpenAI, May 16, 2024), with News Corp (see [A landmark multi-year global partnership with News Corp](#), OpenAI, May 22, 2024) and with *The Guardian* (see [Guardian Media Group announces strategic partnership with OpenAI](#), *The Guardian*, February 14, 2025).

⁶¹ This applies to individual customers. By default, ChatGPT Enterprise/Team/Edu and API customer data is not used to train models.

53. We do not accept OpenAI’s assertion that its collection of personal information is merely “incidental”. Rather, we find that OpenAI collects significant amounts of personal information for the purpose of training its AI models. We note that this is consistent with findings made by some other data protection authorities around the world.⁶² That said, we accept that OpenAI does not exclusively target personal information when collecting information for the purpose of building its training datasets.

54. Based on all of the above, we conclude that OpenAI collects, uses and discloses personal information via and in relation to ChatGPT.

Purposes for collection, use and disclosure

55. For each of the above categories of information, OpenAI identified specific purposes for its collection and processing, which are listed in the table below. However, taken together, we consider OpenAI’s collection, use and disclosure of this personal information to be for the purpose of developing, implementing, continuing to advance and operating ChatGPT (hereinafter referred to as “development and deployment”).

Category of information	Main purposes of collection and processing (as identified by OpenAI)
User interaction data	<p>To provide, administer, maintain, and/or analyze OpenAI’s services</p> <p>To improve OpenAI’s service, develop new services and conduct research (unless a user has opted out)⁶³</p> <p>To carry out business transfers (i.e., user interaction data may be analyzed to demonstrate the performance, utility and usage patterns of OpenAI’s products and services in the event of a business transfer)⁶⁴</p>
<p>Information collected from publicly accessible Internet sources</p> <p>Information which it licenses from third parties</p> <p>Conversations generated by human AI trainers</p>	To train AI models, which may include to provide or improve OpenAI’s products and services and to develop new programs and services.

⁶² See for example the European Data Protection Board (EDPB)’s [preliminary assessment of OpenAI’s practices against requirements under the General Data Protection Regular \(GDPR\)](#), issued on May 23, 2024. OpenAI noted that data protection authorities have arrived at a range of outcomes following their reviews.

⁶³ As explained later in the report, this includes the use of user interaction data for model training (unless a user has opted out).

⁶⁴ We do not examine this specific purpose further in the report.

56. Finally, though OpenAI does use personal information for research purposes, it has not claimed, and we have no evidence to suggest, that its collection of personal information from publicly accessible Internet sources and licensed datasets is carried out **solely** for research purposes.⁶⁵ We accept OpenAI's position that while reliance on the research consent exception requires a case-by-case assessment of the specific context, nature of the processing and applicable statutory conditions, OpenAI may, in appropriate circumstances, be able to rely on that exception where those conditions are met.

Issue 1: Did OpenAI collect, use and disclose personal information for an appropriate purpose?

57. As explained below, we accept that OpenAI's purposes for developing and deploying ChatGPT, as listed in paragraph 55 above, are appropriate⁶⁶. We also accept that OpenAI's practices with respect to personal information collected directly from users via their interactions with ChatGPT are necessary and proportional. However, we find that the manner in which OpenAI initially collected personal information from Internet sources and third parties for the purpose of training its GPT-3.5 and 4 models, as well as the scale and nature of the personal information collected and used from those sources, was overbroad and therefore inappropriate, in contravention of the Acts.

58. The OPC's [Guidance on inappropriate data practices: Interpretation and application of subsection 5\(3\)](#) provides that, in interpreting and applying subsection 5(3) of PIPEDA, the OPC considers certain factors set out by the courts,⁶⁷ meant to assist in determining whether a reasonable person would find that an organization's collection, use, and disclosure of personal information is for an appropriate purpose in the circumstances. As noted above, these factors are to be applied in a contextual manner, which suggests flexibility and variability in accordance

⁶⁵ Under Quebec's Private Sector Act, the research context is an exception to the general consent regime, which allows for personal information to be disclosed without the consent of the individual concerned provided that a privacy assessment was carried out and that the assessment's findings comply with section 21 of the Act. OpenAI has not established that the training of its models consisted of research and that the disclosure of personal information used to train its models complied with the requirements of sections 21, 21.01 and 21.02 of the Quebec Act. PIPEDA, PIPA-BC and PIPA-AB also have exceptions for research purposes, although each law has different restrictions (see sections 7(2)(c) and 7(3)(f) of PIPEDA, section 21(1) of PIPA-BC and section 20 of PIPA-AB, as well as section 14 (3) of Alberta's Personal Information Protection Act Regulation).

⁶⁶ When we refer to an appropriate purpose, this includes reasonable purposes under PIPA-AB and PIPA-BC and legitimate, real and important purposes under Quebec's Private Sector Act.

⁶⁷ The degree of sensitivity of the personal information at issue; whether the organization's purpose represents a legitimate need / *bona fide* business interest; whether the collection, use, and disclosure would be effective in meeting the organization's need; whether there are less privacy invasive means of achieving the same ends at comparable cost and with comparable benefits; and whether the loss of privacy is proportional to the benefits.

with the circumstances.⁶⁸ In applying subsection 5(3), the courts have determined that the OPC is required to engage in a balancing between the individual's right to privacy and the commercial needs of the organization concerned.⁶⁹ This balancing must be "viewed through the eyes of a reasonable person."⁷⁰ More recent jurisprudence has reaffirmed that PIPEDA does not require a balance between competing rights, but rather, between an individual's right to privacy and an organization's need to collect personal information.⁷¹

59. Both PIPA-AB and PIPA-BC provide that an organization may collect, use, or disclose personal information only for purposes that a reasonable person would consider appropriate in the circumstances.⁷² Orders issued by the OIPC-AB have also identified a number of questions for determining whether the collection of personal information in an instance was for a reasonable purpose, including whether a legitimate issue exists to be addressed through the collection of personal information.⁷³ OIPC-BC also considers factors similar to those considered by the OIPC-AB in determining whether the purpose is reasonable in the circumstances.⁷⁴

60. In order to determine whether the purposes for which the personal information was collected by the company are serious and legitimate within the meaning of section 4 of Quebec's Private Sector Act, the CAI takes into account the lawfulness of the purpose and its compliance with the principles of law, justice and fairness.⁷⁵ More specifically, section 4 of Quebec's Private Sector Act requires a company that collects personal information to determine the purposes for the collection before engaging in that collection. In addition, the necessity test developed under section 5 of Quebec's Private Sector Act requires that the purposes justifying the collection be legitimate, important and real, and that the resulting invasion of privacy be proportional to the importance of those purposes. The burden of establishing the seriousness and legitimacy of the

⁶⁸ [Eastmond v. Canadian Pacific Railway](#), 2004 FC 852, para 131.

⁶⁹ [Turner v. Telus Communications Inc.](#), 2005 FC 1601, aff'd 2007 FCA 21; see also [Canada \(Privacy Commissioner\) v. Facebook, Inc.](#), 2024 FCA 140, paras. 62 and 121 (Facebook, Inc. was granted leave to appeal by the SCC, and the OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending).

⁷⁰ *Ibid.* [[Turner v. Telus Communications Inc.](#), 2005 FC 1601, aff'd 2007 FCA 21]; see also [Canada \(Privacy Commissioner\) v. Facebook, Inc.](#), 2024 FCA 140, paras. 61-67 (Facebook Inc. was granted leave to appeal by the SCC, and the OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending).

⁷¹ See [Canada \(Privacy Commissioner\) v. Facebook, Inc.](#), 2024 FCA 140 (Facebook, Inc. was granted leave to appeal by the SCC, and OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending) at footnotes 69-70 *supra*.

⁷² Sections 2, 11, 14 and 17 of PIPA-BC and sections 2, 3, 11, 16 and 19 of PIPA-AB.

⁷³ [Order P2006-011](#) – The OIPC-AB set out a number of questions for determining whether the collection of personal information was for a reasonable purpose, as follows: 1) Does a legitimate issue exist to be addressed through the collection of personal information? 2) Is the collection of personal information likely to be effective in addressing the legitimate issue? 3) Is the collection of personal information carried out in a reasonable manner?

⁷⁴ See, for example: [OIPC-BC Order P05-01 \(2005 BCIPC No. 18\)](#); [P12-01 \(2012 BCIPC No. 25\)](#); [Order P13-02 \(2013 BCIPC No. 24\)](#) and [Order P20-04 \(2020 BCIPC No. 24\)](#).

⁷⁵ [Institut généalogique Drouin Inc.](#), CAI 091570, decision by D. Poitras, February 6, 2015 (in French only).

purposes and the necessity of collecting the personal information rests with the company collecting the information.

61. The principles of appropriate purposes, necessity and proportionality are also discussed and described in the [Principles for responsible, trustworthy and privacy-protective generative AI technologies](#) (“Generative AI Principles”), adopted by the federal, provincial and territorial privacy authorities on December 7, 2023.⁷⁶
62. In the following sections, we evaluate OpenAI’s purposes for developing and deploying ChatGPT, taking into consideration the various factors and questions mentioned above.

Sensitivity of the personal information

63. As further detailed in paragraphs 133 and 296, based on the limited privacy-protective measures in place at the time of developing its GPT-3.5 and GPT-4 models (in particular at the data collection and pre-training stages), we find that OpenAI’s training datasets necessarily included significant amounts of personal information of varying levels of sensitivity, such as medical information, individuals’ opinions on sensitive or controversial topics including about other individuals, and information relating to children.

Legitimate need or issue

64. OpenAI’s stated intention for developing and deploying ChatGPT is to “provide benefits to humanity”, such as assistance with individuals’ everyday tasks, conducting scientific research or inspiring creativity. We accept that this purpose represents a legitimate need or issue for OpenAI.⁷⁷
65. More generally, we appreciate that there are many potentially beneficial applications for the implementation of safe, trustworthy and privacy-protective generative AI technologies, and LLMs in particular in today’s society. Indeed, many research publications and media reports have commented on the concrete benefits of generative AI. For example, the Organisation for Economic Co-operation and Development (OECD) highlighted potential benefits in relation to language translation and interpretation, coding and content creation, and even healthcare.⁷⁸ *The New York Times* described “35 ways real people are using A.I. right now,”⁷⁹ including to write a

⁷⁶ The Generative AI Principles themselves are not legally binding but are derived from foundational privacy principles found in Canadian privacy legislation. They represent the Offices’ expectations with respect to Generative AI, arising from statutory requirements. Furthermore, obligations under privacy legislation in Canada will vary by nature of the organization (such as whether it is in the private, health, or public sector) as well as the activities it undertakes. As such, while the federal, provincial and territorial authorities use “should” throughout the document, many of the considerations listed will be required for an organization to comply with applicable privacy law.

⁷⁷ When we refer to a legitimate need or issue, this includes serious and legitimate reasons under section 4 of Quebec’s Private Sector Act.

⁷⁸ [Generative AI – Benefits](#), OECD.ai, Policy Advisory

⁷⁹ [35 Ways Real People Are Using A.I. Right Now](#), *The New York Times*, April 24, 2023.

speech, learn languages or write Excel formulas. Other articles emphasized AI's benefits for business (e.g., ability to enhance efficiency and productivity, improve customer experience or optimize business operations)⁸⁰ or its potential to accelerate and improve research, resulting in groundbreaking ideas that push the limits of current possibilities.⁸¹

66. We also recognize that there are current and potential future risks associated with the use of LLMs that are released without adequate protections. These risks include the potential for the models to disseminate false information (particularly when used to make or support non-automated decisions about an individual), harm individuals' privacy and reputation (including youth and other vulnerable groups), or assist malicious actors in conducting cybersecurity attacks.
67. As described further below (see paragraphs 128, 294, and 378), OpenAI has explained to the Offices the risk mitigation measures that it has implemented at various phases of its AI models' development and deployment, including to protect against inappropriate, unauthorized uses of ChatGPT, and to mitigate the risks outlined in the above paragraph.
68. While, as explained in paragraph 17, the scope of this investigation did not include consideration of all the possible ways in which its clients (e.g., API customers, developers of GPTs, individual users) may use ChatGPT, we strongly encourage OpenAI to ensure that it has robust measures in place to ensure that ChatGPT is not used in violation of OpenAI's Usage policies or for otherwise inappropriate purposes. In particular, we encourage OpenAI to ensure that ChatGPT is not used for purposes that fall under the No-Go Zones identified in the OPC's [Guidance on inappropriate data practices: Interpretation and application of subsection 5\(3\)](#).⁸²

Determination of the legitimate, real and important nature of the purposes, under Quebec's Private Sector Act

69. As required under Quebec's Private Sector Act, the CAI examined the specific purposes identified by OpenAI (see paragraph 55), to determine if they are legitimate, real and important.

Are the company's purposes for collecting personal information from public sources accessible on the Internet legitimate, real and important?

70. OpenAI stated to the Offices that its purpose for collecting information from public sources accessible on the Internet, which includes personal information, is to train its AI models – this training includes testing its products and services and developing new programs and services.

⁸⁰ [5 Key Benefits of Integrating AI into Your Business](#), Harvard Business School Online, August 1, 2024

⁸¹ See [Generative AI Can Supercharge Your Academic Research](#), Harvard Business Publishing, David Maslach, December 14, 2023. Also, [Generative AI in Academic Research: Perspectives and Cultural Norms](#), Cornell University Task Force, December 15, 2023.

⁸² Also see the [Principles for responsible, trustworthy and privacy-protective generative AI technologies](#) adopted by the federal, provincial and territorial privacy authorities on December 7, 2023.

71. In this regard, the CAI considers that the collection of personal information to train the ChatGPT AI model, which involves making a chatbot available to the public, testing this model and developing new models, can be considered to be for a legitimate purpose.
72. As to whether OpenAI's purposes are real and important, the CAI acknowledges that training an AI model requires a significant quantity of data. The CAI also recognizes that mass collection from publicly accessible sources, particularly through web scraping, may result in the collection of personal information without such information having been specifically targeted for collection.
73. However, as explained above, the Offices do not accept OpenAI's assertion that its collection of personal information is merely incidental to the collection of vast amounts of training data accessible on the web and rather considers that a significant volume of personal information is collected.
74. Despite this consideration, and subject to the question of the sufficiency of OpenAI's privacy risk mitigation measures (to minimize the inclusion of personal information and in particular sensitive information), which will be addressed later in this report, the CAI nonetheless accepts OpenAI's argument that it is necessary for its AI model to learn how personal information fits into a sentence so that it can respond to various user prompts. The CAI further accepts that to accomplish this, the model must have been trained based on a significant quantity of relevant scenarios.
75. The CAI is also aware of and stresses the importance of ensuring that AI models of this kind are sufficiently tested before being made available to the public. Properly training, evaluating and testing models is indeed key to ensuring that they are sufficiently accurate, coherent, fair and safe to use. For example, OpenAI has represented that its models must learn and understand how different types of personal information fit into language, so that it can respond to user prompts in a way that respects OpenAI's policies, including rejecting requests for private or sensitive information about individuals. In this regard, the CAI considers that training the ChatGPT model, testing it and developing new programs and services related to the overall objective of making a chatbot available to the public are for real and important purposes.
76. Finally, considering that OpenAI's overall purpose of developing and deploying ChatGPT, can serve the public interest, the CAI finds (again, subject to the question of mitigation, which will be addressed later) that it is useful and important that the models have access to multiple examples of texts that may contain personal information, to learn how such information fits into the structure of sentences, for the purposes of being able to respond to user prompts correctly and efficiently.

Are the company's purposes for collecting personal information relating to users' interactions with ChatGPT legitimate, real and important?

77. When it comes to the collection of personal information from users' interactions, the CAI accepts that it is legitimate, real and important for OpenAI to want to provide, administer, maintain, analyze and test services that are related to providing a chatbot service.
78. Similarly, the CAI considers that the research and development purposes alleged by the company are legitimate, real and important, provided that this research and development is linked to the overall objective of making a chatbot available to the public.

Effectiveness

79. Based on our testing and open-source research⁸³, the Offices accept that LLMs such as ChatGPT are generally effective at generating and simulating natural language and performing other natural language processing tasks, such as text summarization or language translation.
80. However, as explained in the Accuracy section below (Issue 4), we find that at the time of developing and deploying GPT-3.5 and 4, OpenAI did not comply with the accuracy requirements under the Acts.

Less-privacy invasive means and proportionality

81. As explained above, we accept that: (i) OpenAI's general purpose for collecting, using and disclosing personal information – i.e., to develop and deploy its LLMs – represents a legitimate need; and (ii) that ChatGPT is generally effective at generating natural conversational language, subject to the accuracy concerns identified in the above paragraphs. However, we must also consider whether OpenAI could have developed and deployed its GPT-3.5 and 4 models through less privacy-invasive means and whether the harms to privacy resulting from OpenAI's practices were proportional to the potential benefits of these LLMs.
82. To this end, we examine OpenAI's collection, use and disclosure of personal information from both: (i) publicly accessible websites and licensed datasets; and (ii) users' interactions with ChatGPT.

⁸³ e.g., [A Comprehensive Study of ChatGPT: Advancements, Limitations, and Ethical Considerations in Natural Language Processing and Cybersecurity](#), by Moatsum Alawida (Abu Dhabi University), Sami Mejri (Khalifa University), Abid Mehmood (Abu Dhabi University), Belkacem Chikhaoui (TELUQ University, Montreal) and Oludare Isaac Abiodun (University of Abuja), August 16, 2023; [Large language models' ability to generate text also lets them plan and reason](#), *The Economist*, April 19, 2023.

Collection, use and disclosure of personal information from publicly accessible websites and licensed datasets

83. As explained below, we find that when developing and deploying its GPT-3.5 and 4 models, OpenAI did not appropriately minimize the invasion of privacy resulting from its collection, use and disclosure of personal information, as required under the Acts. We also find that the harms to privacy resulting from this development and deployment were not proportional to the potential benefits of ChatGPT. While we acknowledge that OpenAI may have faced technological hurdles in attempting to train its GPT-3.5 and 4 models in a less privacy-invasive manner, we note that in response to our Preliminary Report, OpenAI represented that it has recently implemented new mitigation measures that significantly reduce the privacy risks associated with the development of Generative AI models. This demonstrates that, with innovation and forethought, less privacy-invasive means of training GPT-3.5 and 4, at a comparable cost and with comparable effectiveness, would have been available.
84. As discussed further at paragraph 127 and following, while the exact size of OpenAI's training datasets has not been made public, OpenAI represented that it includes trillions of words. Furthermore, our investigation revealed that, while OpenAI does not exclusively seek to collect personal information when collecting publicly accessible data online, this immense dataset includes significant amounts of personal information of varying levels of sensitivity.
85. OpenAI asserted that the nature of its processing is not intrusive given that it relies on unstructured datasets that are not indexed or organized by reference to an identifier, uses tokenization (i.e., raw text is transformed into numerical representations and is therefore not used in its original format) and does not process personal information in a targeted manner (i.e., to gain specific knowledge about private individuals or generate profiles). OpenAI further stated that the processing aims at using personal information to teach AI models the concept and meaning of personal information in a general manner.
86. That said, OpenAI also acknowledged, and we agree, that there are risks associated with this practice, including in relation to privacy. With this in mind, the company explained that it has implemented measures to decrease the presence of personal information in the datasets.
87. In particular, OpenAI submitted that when collecting training text for GPT-3⁸⁴, it took steps to avoid pirated content and to remove duplicative and/or harmful content (e.g., Child Sexual Abuse Material, hate speech, erotic content, spam) in the training text. For the training of GPT-4 onwards, OpenAI explained that it took additional steps to identify and remove certain sites that were designed as an index or collection of personal information from the training text.⁸⁵ Finally,

⁸⁴ An earlier version of GPT-3.5 released in 2020.

⁸⁵ e.g., [checkpeople.com](https://www.checkpeople.com), [familytreenow.com](https://www.familytreenow.com), and [infotracer.com](https://www.infotracer.com)

OpenAI stated that it does not circumvent paywalls or account-protected websites or obtain information from the dark web.

88. OpenAI represented that it only removes a small portion of website categories from the data it includes in its training datasets. It also confirmed that it does not screen out social media websites, websites aimed at children, or websites that may contain sensitive information regarding other vulnerable groups. Further, given that OpenAI only removes limited website categories from training data, information from discussion forums would likely be included in that dataset.
89. Data from sources such as social media and discussion forums contain vast amounts of personal information (including that of children), some of which will be sensitive, and much of which will reflect the subjective and potentially inaccurate views and opinions of the individuals who post this information.
90. OpenAI explained to the Offices that the collection of data from publicly accessible, non-gated websites is necessary to teach its models about language, and that the goal in showing the model this content is to teach the model how such content is expressed, not to endorse its truth. More specifically, OpenAI stated that in order to develop highly capable general-purpose AI models (“GPAI models”), the models need to be trained on large and diverse datasets, which necessarily includes informal, real-world exchanges. These allow models to learn how language is used organically in interactions between individuals, particularly in informal, spontaneous contexts outside of structured or edited writing, such as casual conversations and everyday exchanges.
91. Moreover, OpenAI submitted that complete anonymization of training data for the development of GPAI models remains technically impossible as it would compromise the effectiveness of the models and their ability to serve broadly beneficial societal purposes. OpenAI further represented that despite efforts to employ and expand the use of innovative privacy protecting measures (such as synthetic data), the current state-of-the-art does not offer less intrusive means for developing highly capable AI models.⁸⁶
92. In response to these representations, the Offices acknowledge that there may be benefits in collecting personal information in this context and that a large and varied body of training text may assist ChatGPT in understanding and responding effectively to users’ queries.

⁸⁶ To corroborate its claims, OpenAI referenced a study [published](#) by Google on September 12, 2025, on its attempt to train a model using differential privacy (DP). Google explained that the resulting model “yields utility comparable to non-private models from roughly five years ago” (analogous to GPT-2) and it required significantly more compute to train. That said, we note that Google concluded that this model “represents a significant step forward in the journey towards building AI that is both powerful and private by design”, and that “while a utility gap still exists between DP-trained and non-DP-trained models, we believe this gap can be systematically narrowed with more research on mechanism design for DP training.”

93. However, we do not accept OpenAI's assertion that there were no less privacy-protective ways to develop GPT-3.5 and 4 at the same (or comparable) cost and with the same effectiveness. Indeed, as explained below, OpenAI represented in response to our Preliminary Report that it has recently implemented new mitigation measures that significantly reduce the privacy risks associated with the development of Generative AI models. In our view, this demonstrates that, with innovation, less privacy-invasive means of training GPT-3.5 and 4 would have been available.
94. We find that, absent these mitigation measures, OpenAI's development and deployment of GPT-3.5 and 4 resulted in a privacy-invasive collection of significant amounts of personal information, which necessarily increased the risk of privacy harms, such as those resulting from inadvertent disclosure of private information in model outputs, breaches of training data and more generally, individuals' loss of control over their personal information. Furthermore, learning about such privacy risks or incidents may have negatively impacted individuals' willingness to engage openly in digital society.
95. Finally, the fact that OpenAI's training datasets included personal information contained in sources such as social media and discussion forums—which may often be inaccurate, for example where they include opinions that are not rooted in fact and/or are biased— would also have amplified the risks of inaccurate personal information appearing in model outputs.
96. OpenAI submitted that there is no concrete evidence of a systemic issue of inaccurate personal information appearing in ChatGPT's outputs. The company further indicated that it takes concrete steps to improve accuracy and mitigate privacy risks in model outputs, such as by training its models to refuse to provide private or sensitive information, even if it is publicly accessible (for more detail, see Issue 4 – Accuracy).
97. With respect to the potential harms that could result from the disclosure of personal information via ChatGPT responses, OpenAI represented that personal information appearing in its model outputs would most likely be included because it is widely accessible on the Internet (as opposed to, for example, being found in a single source).
98. While we accept that this may mitigate, to a certain extent, the risk that personal information occasionally posted on the Internet may be disclosed by ChatGPT, the fact that personal information is widely accessible on the Internet does not mean that it is necessarily accurate and unbiased. This is especially true at a time when misinformation and disinformation can spread across the Internet at an unprecedented speed. More importantly, as further explained in Issue 2, the fact that personal information is accessible does not represent a *carte blanche* to collect and use it without limits.
99. Therefore, we find that OpenAI's mitigation measures in place at the time of training its GPT-3.5 and 4 models were not sufficient to limit the scope of its collection, use and disclosure of personal information to that which was necessary and proportional to effectively train its

models.⁸⁷ For these reasons, we find that at that time, the benefits of that practice did not outweigh the risks of privacy harms.^{88 89}

Collection, use and disclosure of personal information included in users' interactions with ChatGPT

100. As a preliminary matter, the Offices recognize that OpenAI's collection of user interactions is necessary to efficiently respond to user queries. Therefore, this section focuses on OpenAI's use and potential disclosure of personal information included in these user interactions for the purpose of developing and deploying its AI models.
101. We acknowledge that the use of a certain level of user interaction data may be beneficial and necessary to properly train the models, especially during the fine-tuning phase. As mentioned in paragraph 48, OpenAI explained that fine-tuning involves (among other steps) the use of a subset of users' interactions to improve the model's ability to answer user queries in a way that people find useful, that is, in a more relevant, safe, accurate, and unbiased way.
102. Furthermore, OpenAI represented that, when training its GPT-3.5 and 4 models, it put in place certain measures to mitigate the risks associated with utilizing users' interactions for training purposes. As explained in more detail at paragraph 294, these measures included disassociating interactions from user accounts, using a third party filtering tool to remove personal identifiers, allowing users who have an account to choose whether their interactions with ChatGPT would be used for model training, and informing users (albeit not adequately, as discussed at paragraph 293) not to include sensitive information in their interactions with the tool. The company also instructed human model trainers to exclude from the fine-tuning datasets any information that could constitute personal information. Finally, OpenAI explained that it only used a small subset of the user interactions it collected to train its models.

⁸⁷ In [Canada \(Privacy Commissioner\) v. Facebook, Inc.](#), 2024 FCA 140, paras. [114-115](#), the Federal Court of Appeal considered the argument that it was "practically impossible" for an organization to ensure privacy compliance. The Court ultimately held that an organizations' claim of impossibility stemming from its chosen business model cannot limit the scope of its responsibilities under PIPEDA (Facebook, Inc. was granted leave to appeal by the SCC, and the OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending).

⁸⁸ As part of the necessity test for the collection of personal information (Quebec's Private Sector Act), the burden of establishing that the benefits of collecting such information outweigh the invasion of privacy it represents rests with the company.

⁸⁹ We note that the EDPB considered similar factors to those we have considered in our analysis in this section. See the EDPB's [preliminary assessment of OpenAI's practices against requirements under the General Data Protection Regular \(GDPR\)](#), issued on May 23, 2024; and its [broader position paper outlining general considerations for assessing the development and deployment of LLMs against certain GDPR provisions](#), issued on December 17, 2024.

103. As we note in paragraph 296 of this report, the third-party filtering tool that OpenAI used at the time of training GPT-3.5 and 4 removed only a subset of the information that would constitute personal information as defined under the Acts, such that sensitive information, like opinions, could still be included in the user interaction data used for training (and potentially disclosed in model outputs). However, we accept that the combination of the various measures outlined above significantly mitigated the risk of privacy harm associated with training the model using personal information included in user interactions.

104. In light of the necessity to train the model using user interactions, and of the associated benefits highlighted above (i.e., to provide more effective and efficient responses), we accept that the benefits of this practice were **proportional** to the residual risk of privacy harm, taking into consideration the mitigation measures implemented by OpenAI.

Findings related to GPT-3.5 and 4

105. As mentioned above, we find that the nature and scale of OpenAI's collection and use of personal information from publicly accessible websites and licensed datasets, at the time of training its GPT-3.5 and 4 models, was overbroad and therefore not necessary and proportional. Consequently, we find⁹⁰ that OpenAI contravened subsection 5(3) of PIPEDA, sections 2, 11, 14 and 17 of PIPA-BC, sections 11, 16 and 19 of PIPA-AB and section 5 of Quebec's Private Sector Act.

106. Furthermore, we accept that the collection, use and disclosure of personal information from users' interactions with ChatGPT were effective in advancing OpenAI's legitimate need to develop and deploy ChatGPT – in particular, to improve model outputs in responses to user prompts – and that the benefits of this practice were proportional to the residual risk of privacy harm, taking into consideration the mitigation measures implemented by OpenAI. Consequently, we find this aspect of the complaint to be **not-well founded**.

Recent developments and conclusion under PIPEDA

107. In its response to the Offices' Preliminary Report, OpenAI emphasized that it was not aware of a case where Canadian privacy commissioners ruled against an organization regarding the processing of publicly accessible information, where the purpose of collecting the information was not also found to be inappropriate. We would highlight the case of RateMDs, where the OPC has held that certain data practices can be inappropriate within the meaning of subsection 5(3) of PIPEDA even if the overarching purpose is not itself inappropriate.⁹¹ Furthermore, the jurisprudence on subsection 5(3) of PIPEDA emphasizes the importance of a

⁹⁰ Where we speak of the findings of more than one authority, we mean the findings of each relevant authority under their respective legislation.

⁹¹ See the OPC's investigation into RateMDs ([PIPEDA Findings #2020-002](#) at paras. 75-83) where the OPC held that the collection, use, and disclosure of review information related to health practitioners was not inherently inappropriate, but the practice of "pay-for-takedown" for user-generated reviews offended s. 5(3) of PIPEDA.

contextual, case-by-case assessment, rather than one that is focused exclusively on the overarching purpose for the information.⁹²

108. In response to our Preliminary Report, OpenAI also informed the Offices that it has recently developed a tool that can detect and mask identifying information about private individuals in publicly accessible Internet data and in licensed datasets used to pre-train OpenAI's models. OpenAI further explained that it now also uses this tool (in lieu of the previous third-party filtering tool) to redact personal identifiers from users' interactions used to fine-tune the models.
109. According to OpenAI, this new tool can identify a wide range of personal information about private individuals in the training datasets (e.g., names, phone numbers) and mask it prior to it being used for training, so the models do not learn from it. OpenAI indicated that the tool can also detect other categories of personal information that are similarly private or personal but which it has never been trained to recognize. Accordingly, to the extent that a broader range of personal information, such as an individual's opinions or characteristics, are included in the datasets, the tool can detect and redact identifiers which would link such information to an identifiable individual.
110. OpenAI further stated that the tool uses context to detect whether information is private or personal in nature and should be masked. More specifically, the company noted that the tool can distinguish between personal information about private individuals, personal information about public figures, and information about fictional characters. As a result, OpenAI stated that it is able to mask personal information about private individuals, as well as to determine when to mask private information about public figures (e.g., their personal address or personal phone number) and when to maintain information about such public figures which may be of interest to the public (e.g., their business address or business phone number). Finally, OpenAI provided the Offices with the results of recent internal evaluations showing the effectiveness of the tool at detecting different types of personal information.
111. More specifically, OpenAI submitted that it conducted evaluations against other filtering tools using the open-source "PII Masking 300k benchmark".⁹³ OpenAI explained that when fine-tuned on a small subset of the benchmark, its new filtering tool reached 98–99% recall (i.e., proportion of true instances of personal information that a system correctly identifies) with a 3–

⁹² See in particular *Eastmond v. Canadian Pacific Railway*, [2004 FC 852](#) (CanLII) at para. 131, where the Court emphasized the contextual and flexible nature of the s. 5(3) PIPEDA analysis and noted that "the appropriate purposes for collection may be different than the appropriate purposes for use and the appropriate purposes for disclosure of collected information".

⁹³ "[PII Masking 300k benchmark](#)," is a synthetic dataset of roughly 300,000 annotated text samples, developed by AI4Privacy and publicly available on Hugging Face, which is used to evaluate and train models to remove identifying information from text, especially in the context of AI assistants and LLMs.

6% false positive rate (i.e., the proportion of instances of text incorrectly flagged as personal information by the system). Furthermore, OpenAI stated that it ran additional evaluations using 80,000 synthetic chat excerpts labeled by professional data annotators. Comparing the tool's predictions with these human labels, OpenAI indicated that it found substantial alignment with human judgment, far surpassing the third-party filtering tool it previously used.⁹⁴

112. The OPC accepts that this new tool – combined with OpenAI's other mitigation measures in place at the various stages of development and deployment of ChatGPT⁹⁵ – can significantly reduce the risk that the personal information of private individuals, and sensitive information more specifically, will be included in the datasets used to train OpenAI's future models. Similarly, the OPC accepts that this will also reduce the risk of such information being disclosed in model outputs.

113. In making this determination, the OPC also considered the additional transparency measures that OpenAI has committed to implementing. In particular, as further discussed in other sections of this report⁹⁶, OpenAI has agreed to publish a Canadian blog post on its website explaining its privacy practices and take measures to promote the post and its contents in the Canadian media. The blog post will inform individuals that user interactions may be reviewed and used to train its models, advise users not to share sensitive information via their interactions with ChatGPT, and provide information about the categories of content used to train its models. We find that these transparency measures will enhance public awareness of OpenAI's privacy practices, thereby further limiting individuals' sharing, and OpenAI's collection of, sensitive information.

114. Finally, OpenAI informed the Offices that it has deprecated (i.e., retired) its GPT-3.5 and 4 models and confirmed that the new mitigation measures, including the above-mentioned filtering tool, have been used throughout the training of its current models powering ChatGPT.⁹⁷

115. Therefore, with a view to reflecting an appropriate balancing of freedom of expression and privacy, the OPC finds the aspect of the complaint related to the collection, use and disclosure of personal information from publicly accessible websites and licensed datasets to be **well-founded and conditionally resolved** under PIPEDA.⁹⁸

⁹⁴ For example, the new tool was able to detect 25% more personal identifiers and reduce false positives by more than 96% compared to the third-party filtering tool used previously.

⁹⁵ These measures are detailed in various sections of this report and are summarized in Appendix A.

⁹⁶ In particular, see section relating to OpenAI's response to our Preliminary Report.

⁹⁷ Furthermore, OpenAI confirmed that it did not use GPT-3.5 and 4 as base models for the training of the current models.

⁹⁸ The CAI assesses, based on the necessity test set out in section 5 of Quebec's Private Sector Act, that the new mitigation measures put in place could sufficiently minimize the impact of OpenAI's collection of personal information from publicly available Internet sources and third parties holding licensed data, thereby making such collection proportional to the invasion of privacy it may represent. Consequently, the CAI agrees with the conclusion that this aspect of the complaint is well-founded and conditionally resolved.

116. This conclusion is based on OpenAI’s representations and our understanding, as well as our expectation that OpenAI will continue to effectively implement and improve these mitigation measures and develop further innovative privacy-protective techniques in the future.

Issue 2: Did OpenAI obtain valid consent and meet its obligation to inform individuals with respect to its collection, use and disclosure of personal information?

117. For the reasons outlined below, we find that OpenAI did not obtain valid consent for its collection, use and disclosure of personal information for the purpose of developing and deploying its GPT-3.5 and 4 models.

118. The following section (Issue 2A) assesses whether OpenAI’s collection and use of personal information from publicly accessible websites or licensed third-party sources was compliant with the Acts’ consent provisions. We then consider the Respondent’s collection and use of personal information from users via their interactions with ChatGPT (Issue 2B).⁹⁹ Finally, we examine OpenAI’s disclosure of personal information from these various sources (Issue 2C).

Issue 2A: Did OpenAI obtain valid consent for the collection and use of personal information from publicly accessible websites and licensed third-party sources?

Analysis under PIPEDA, PIPA-BC and PIPA-AB

119. We find that OpenAI did not have implied consent¹⁰⁰ for its collection and use of individuals’ personal information from publicly accessible websites and licensed third-party sources for the purpose of training its GPT-3.5 and 4 models.¹⁰¹

120. Pursuant to sections 5(1), 6.1 and 7, as well as Principle 4.3 of Schedule 1 of PIPEDA, sections 6-8 of PIPA-BC and sections 7-8 of PIPA-AB, the consent of individuals is required for the collection, use or disclosure of their personal information, unless an exception applies. The type of consent required will vary depending on the circumstances and the type of information involved.

121. The [Guidelines for obtaining meaningful consent](#) (“the Consent Guidelines”) jointly issued by the OPC, OIPC-AB and OIPC-BC provide that “organizations must generally obtain *express consent*” when: (i) the information being collected, used or disclosed is sensitive; (ii) the

⁹⁹ Our analysis will not examine consent for the personal information that might be included in conversations created by OpenAI’s human trainers. Indeed, this information is not collected from publicly accessible sources, licensed third parties or automatically generated by ChatGPT, and therefore, it falls outside of the scope of this investigation.

¹⁰⁰ Where we refer to “implied consent”, this is considered inclusive of “implicit consent” under PIPA-BC and “deemed or notice consent” under PIPA-AB.

¹⁰¹ The application of the Quebec’s Private Sector Act requires certain distinctions that will be addressed at the end of Issue 2A.

collection, use or disclosure is outside of the reasonable expectations of the individual; and/or (iii) the collection, use or disclosure creates a meaningful residual risk of significant harm.¹⁰²

122. OpenAI represented that it relies on individuals' implied consent to collect and process the personal information found on publicly accessible websites and in licensed datasets that it uses to pre-train its models. OpenAI argued that the Acts may not be designed to address the complex challenges and nuances associated with innovative technologies and business models where there is no direct relationship between the parties involved. OpenAI submitted that the Offices should apply a contextual and balanced approach based on "flexibility, common sense, and pragmatism", and it justified its reliance on implied consent based on the following factors:

- i. the pressing and substantial benefits of training AI models, which OpenAI states have already provided "dramatic benefits to humanity";
- ii. the necessity of this information for the processing, given that models must train on a large amount of text to develop an understanding of how language works;
- iii. the impracticability of direct notification to individual persons given the impossibility of identifying and locating them based on the information found in the unstructured raw datasets;
- iv. OpenAI's reasonable efforts to be transparent about its information handling practices related to the development and training of its models, and the public notice that it provides through readily available means such as its Privacy Policy or Terms of Use;
- v. OpenAI's use of de-identification measures and other risk mitigation measures, such as the unstructured nature of training datasets, the use of filters to exclude certain sites and content from training data, or the fact that personal information is not processed in a targeted manner to build profiles of, or gain knowledge about, specific individuals (mitigation measures that are discussed throughout this report); and
- vi. OpenAI's view that the balance of interests favours an opt-out form of consent. OpenAI further states that it continually analyzes and seeks to optimize the balance of risks and benefits involved in training and making the models available to the public, taking into account the various measures it has implemented to reduce the processing of personal information and mitigate potential risks.

¹⁰² The Consent Guidelines are based on the underlying principles set out in the Acts and incorporates Supreme Court of Canada jurisprudence, notably *Royal Bank of Canada v. Trang*, 2016 SCC 50 at paragraphs 23, 24, 34-36, 39 (sensitivity) and 43-45 (reasonable expectations). For a more recent assessment of consent by the Court under PIPEDA, see [Canada \(Privacy Commissioner\) v. Facebook, Inc.](#), 2024 FCA 140, paras 61-63, 67, 70-73, 120, 123-124 and 132 (Facebook, Inc. was granted leave to appeal by the SCC, and the OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending).

123. Consent is a core requirement of PIPEDA, PIPA-AB and PIPA-BC, limited only by carefully defined legislative exceptions expressly set out in the respective legislations. This has been confirmed in Supreme Court of Canada jurisprudence¹⁰³, as well as in federal¹⁰⁴ and provincial¹⁰⁵ case law.
124. We acknowledge that the development of new technologies, such as AI, may raise new challenges for organizations when it comes to complying with existing privacy laws – in particular, with respect to consent. However, we note that the Acts are technology-neutral, and we are required to assess OpenAI’s practices against the existing applicable legal frameworks. The authority to enact new laws or amend current laws, whether broad in scope or specifically dealing with generative AI or other emerging technologies, remains with Parliament and the legislatures.
125. Consistent with the modern approach to statutory interpretation, the Offices have interpreted the Acts with “flexibility, common sense, and pragmatism”. As noted above, we have relied on this approach to assess OpenAI’s practices against the Acts, taking into consideration the various factors listed above at paragraph 122.

Sensitivity

126. Information that will generally be considered sensitive, and thus require a higher degree of protection, includes health and financial data, ethnic and racial origins, political opinions, genetic and biometric data, an individual’s sex life or sexual orientation, religious or philosophical beliefs, and young people’s personal information.¹⁰⁶
127. While the exact size of the datasets that OpenAI collects directly from publicly accessible sources and indirectly from licensed third-party sources has not been publicly disclosed nor confirmed to the Offices during the course of the investigation, OpenAI represented that it includes trillions of words. Indeed, the Common Crawl database alone – one of the sources which OpenAI relies on to build its datasets – contains petabytes (i.e., millions of gigabytes) of

¹⁰³ *Royal Bank of Canada v. Trang*, [2016 SCC 50, at paragraphs 23-24](#).

¹⁰⁴ [Canada \(Privacy Commissioner\) v. Facebook, Inc.](#), 2024 FCA 140, para 72 (Facebook, Inc. was granted leave to appeal by the SCC, and the OIPC-BC received leave to intervene. The SCC heard the matter on March 19, 2026. Judgment from the SCC is pending).

¹⁰⁵ *Cran v British Columbia (Information and Privacy Commissioner)*, [2024 BCSC 1130, at paragraph 76](#).

¹⁰⁶ See OPC’s [Interpretation Bulletin on Sensitive Information](#), which specifies what factors are relevant when determining whether personal information is sensitive, and the Canadian regulators’ resolution on [‘Putting best interests of young people at the forefront of privacy and access to personal information’](#) (adopted on October 4-5, 2023).

data regularly collected since 2008 (i.e., over 250 billion pages spanning 17 years, with 3–5 billion new pages added each month).¹⁰⁷

128. OpenAI maintained that, as part of its mitigation measures aimed at reducing the presence of personal information in the final pretraining datasets used to train GPT-3.5 and 4, it removed certain categories of websites from the raw data collected from publicly accessible websites (i.e., log-in gated websites, websites with pirated or harmful content, adult websites, and for GPT-4 specifically, websites that aggregate personal information about individuals¹⁰⁸) and “deduplicated” (i.e., removes redundant) content.
129. Regarding data licensed from third parties, OpenAI submitted that it selected datasets that do not contain extensive personal information, while recognizing that they may contain personal information incidentally (e.g., a licensed encyclopedia may contain an entry concerning a public figure who is still alive).
130. In any event, OpenAI stated that the amount of data alone is not determinative of the sensitivity of the information, particularly given the non-access-gated, publicly accessible data and the inherently non-intrusive nature and purpose of the processing. OpenAI also represented that potential privacy risks are further mitigated by the unstructured nature of the pretraining datasets, which are not indexed or organized by reference to individuals, and the fact that information is tokenized (see paragraph 85).
131. OpenAI did not grant our request to access and review their systems and, consequently, we were unable to directly assess the effectiveness of these mitigation measures. As mentioned above, the categories of websites that OpenAI removed from the GPT-3.5 and 4 pretraining datasets were very limited. In particular, OpenAI confirmed that it did not screen out social media websites, websites aimed at children, or websites that may contain information regarding other vulnerable groups whose information is more likely to be considered sensitive.
132. OpenAI’s pretraining datasets primarily consisted of data that was publicly accessible on the Internet, such as forum posts, product reviews, user comments, social media, essays, articles or books. A *New York Times* article reported that OpenAI, and other AI companies, also transcribed one million hours of YouTube videos to harvest text for their AI models.¹⁰⁹
133. In that context, and given the absence of specific mitigation measures aimed at detecting and masking private identifying information in the GPT-3.5 and 4 pretraining datasets, we find that these datasets necessarily included sensitive information such as financial or

¹⁰⁷ See [Common Crawl’s website](#). We note that OpenAI represented that it does not use the entirety of Common Crawl’s database but rather uses selects data from its industry-standard learning datasets. It also indicated that it uses its GPTBot to crawl and scrape publicly accessible websites.

¹⁰⁸ e.g., checkpeople.com or familytreenow.com

¹⁰⁹ [How Tech Giants Cut Corners to Harvest Data for A.I.](#), *The New York Times*, April 6, 2024. The Offices did not independently verify this claim.

medical information, information about religious or political beliefs, opinions about sensitive or controversial topics, and information relating to children; some of which will have been posted by third parties (i.e., not the individual themselves).

134. While licensed datasets represented a much smaller subset of OpenAI's pretraining data, we find that they might also have included sensitive personal information. For example, news articles about an individual's past criminal offences, including mere suspected offences, may reveal sensitive information such as their ethnic origin or health information.¹¹⁰
135. We find that OpenAI could not rely on implied consent for the collection of such sensitive information for the purpose of training its GPT-3.5 and 4 models.

Reasonable expectations

136. We further find that at the time OpenAI trained its GPT-3.5 and 4 models, individuals could not have reasonably expected that their personal information – even where it was publicly accessible on the Internet after they had posted it themselves – would be used to train OpenAI's models.
137. OpenAI stated that the assessment of reasonable expectations must be contextual and reflect the realities of online activity, including the longstanding practice of web crawling for various purposes (including by search engines, archives or academic researchers) and the well-established principle that information posted on the open Internet without access restrictions carries diminished privacy expectations.
138. To support this argument, OpenAI cited (among other cases) the recent Alberta Court of King's Bench decision in *Clearview AI Inc v Alberta (Information and Privacy Commissioner)*, in which the Court held, in the context of Internet search engines, that information posted on the open Internet without access restrictions carries diminished privacy expectations. The Court concluded that "[a] reasonable person posting images and information to a website or social media platform subject to terms of service but without using privacy settings expects that such images and information will be indexed and retrieved by Internet search engines; indeed, that is sometimes the point of posting images and information to the Internet without using privacy settings."¹¹¹
139. While in the case of adult individuals who chose to make their personal information publicly accessible online, reasonable expectations of privacy may be diminished for this information, we do not find that they are altogether extinguished. This is especially true where information has been posted by a third party without the individual's knowledge and consent.

¹¹⁰ This is consistent with the OPC's recent Report of Finding in Google where the OPC found that the continued display of news articles about a complainant on the Internet qualified as sensitive information ([PIPEDA Findings #2025-002](#) at para. 101, 137).

¹¹¹ *Clearview AI Inc v Alberta (Information and Privacy Commissioner)*, [2025 ABKB 287](#), at paragraph 137.

Moreover, as further explained below, the Acts carefully define what constitutes “publicly available information” and continue to apply privacy protections to personal information that does not fall within that exemption.

140. While OpenAI acknowledged that search engines are not the subject of this investigation, it stated that the Alberta Court’s reasoning is directly relevant and instructive, as both search engines and AI models rely on publicly accessible Internet data to serve socially beneficial functions.
141. While we agree that search engines and generative AI technologies are increasingly being integrated, this was not the case when ChatGPT was launched in November 2022. At the time, this service was new and not widely known to the public. It also differed fundamentally from traditional search engines, which simply enable users to navigate and find information on the Internet. By contrast, ChatGPT offered a novel service, which aimed to create new content in response to users’ prompts, some of which could contain plausible but inaccurate or fabricated information¹¹² (including personal information).
142. We are of the view that, in these circumstances, individuals’ reasonable expectations regarding these two technologies – a novel and pivotal one as opposed to a longstanding and well established one – cannot be equated. We accept that individuals’ reasonable expectations about web crawling for the purpose of indexing websites on search engines may have evolved over time, as people became familiar with, and more educated about, these tools. However, we believe that at the time GPT-3.5 and 4 were trained, individuals would not reasonably have expected, and in fact could not have expected, such crawling (and scraping) to be conducted for the purpose of developing a technology that they were unaware of or unfamiliar with.
143. Some of the information collected and used to train OpenAI’s models was posted several years, if not decades, ago. At the time, individuals could not have reasonably expected that this information would eventually be used for the training of a technology that had not yet been released.
144. Similarly, back when GPT-3.5 and 4 models were being trained, generative AI had not yet gained widespread popularity. While the underlying technology had been in development for years, it was not yet popular amongst the general population. Therefore, regardless of whether the option to opt out of this processing was already available, individuals would not have had a reason to consider doing so, as they were likely unaware that the training was ongoing. Of note, jurisprudence on consent supports the proposition that a lack of awareness about the availability of an opt-out can vitiate consent.¹¹³

¹¹² i.e., hallucinations, as discussed later in this report.

¹¹³ *Englander v. TELUS Communications Inc.* (FCA), [2004 FCA 387 \(CanLII\)](#) at para. 67.

145. Furthermore, we do not accept that, even after ChatGPT was released, an individual who posted content to a website subject to terms of service, without using privacy settings, would have reasonably expected such content to be scraped and used for the purpose of training Generative AI models. Indeed, several studies have shown that most individuals either do not read or have difficulty understanding of websites' terms of service and privacy policies, which are often long, legalistic and complex.¹¹⁴ Even for users who intend to adjust their settings, the design of websites' privacy settings can sometimes be confusing, if not deceptive, rendering it more difficult for users to control their information. This often results in individuals being overwhelmed and having significant difficulty effectively managing their privacy settings to reflect their intended choices.
146. We find that there is generally no obvious connection between an individual's posting of personal information online for a specific purpose (e.g., to connect with friends on social media, write a product review, publish a YouTube video, or participate in a forum discussion) and the subsequent scraping and use of that personal information to train AI models developed by an organization with which the individual has no relationship.
147. In many instances, such as where an individual's personal information was posted by another person, the subject may not even have been aware of the existence of their personal information online, let alone that it could be used to train an AI model. In these circumstances, it cannot be assumed that all posted personal information about an individual was provided by them, or with their knowledge and consent.
148. Even in instances where an individual may have intentionally made their information more widely available on the Internet, some websites, particularly social media, are dynamic in nature, whereby individuals can, for example, edit or remove content from their own publicly accessible profiles. Where information is scraped from such a website, the individual loses that control over their personal information.
149. OpenAI argued that given the substantial public attention resulting from coverage that OpenAI and ChatGPT received in news reports around the globe, some individuals whose personal information was incidentally included in its training datasets may have visited OpenAI's website and learned more about its information handling practices.
150. We find this to be speculative and that it puts too much onus on individuals. In any event, even if an individual were to visit OpenAI's website for the purpose of understanding OpenAI's privacy practices:

¹¹⁴ See for example, [Attitudes towards Data Privacy and Transparency](#), Canadian Marketing Association, January 8, 2018; [How Americans View Data Privacy](#), Pew Research Center, October 18, 2023 and OPC's [Sweep Report on Deceptive Design Patterns](#), July 9, 2024.

- i. It is unlikely that they would understand the extent to which their information could be included in the organization’s training datasets. As discussed below, we find that OpenAI is not sufficiently transparent about the nature and categories of publicly accessible websites or licensed datasets that it uses to train its models; and
- ii. OpenAI would, in fact, have already collected the individual’s personal information before they chose to learn more about OpenAI.¹¹⁵ As mentioned above, where the information was posted on the Internet before ChatGPT’s launch in November 2022, the individual would likely have been unaware of ChatGPT and would have had no reason to suspect that their information might be used to train that LLM.

151. With respect to information posted by third parties, OpenAI submitted that this should not materially affect the reasonable expectations analysis. OpenAI further stated that where information is made publicly accessible without restrictions, it is reasonable for OpenAI to proceed on the basis that all such information was lawfully and appropriately published, while acknowledging that remedies exist under applicable law for addressing instances of unlawful disclosure. According to OpenAI, this approach is particularly appropriate given the non-intrusive nature of its processing and the mitigation measures implemented at the various phases of its models’ development and deployment.

152. As mentioned above, we find that it would not be reasonable to assume that all information made publicly accessible without restrictions was provided with the knowledge and consent of the individual to whom the information relates. Furthermore, we do not accept that at the time of training GPT-3.5 and 4 models, the mitigation measures implemented by OpenAI (especially at the pre-training stage) were sufficient to effectively minimize the presence of personal information in the training datasets, whether posted by the individuals to whom it related or by third parties.

153. Therefore, we find that OpenAI’s collection and use of personal information obtained from publicly accessible sources and third-party licensed datasets, for the purpose of training its GPT-3.5 and 4 models, was outside of individuals’ reasonable expectations, such that it could not rely on implied consent for that practice.

Choice

154. Furthermore, the Consent Guidelines explain that individuals cannot be required to consent to the collection, use or disclosure of their personal information unless it is integral to the provision of the product or service – they must be given a choice, and that choice must be clearly explained and easily accessible.

155. OpenAI’s use of any specific individuals’ personal information collected from publicly accessible websites for training purposes is not integral to the provision of its services, even

¹¹⁵ As indicated above, OpenAI started preparing datasets to train GPT-3.5 months before it launched ChatGPT.

where it may be useful to OpenAI. Therefore, OpenAI must offer individuals a choice about whether to participate or not in this practice.

156. OpenAI represented that it provides individuals with the ability to request that their verified personal information be filtered out from future model training runs, and, where their information appears in model outputs, to request removal of verified personal information from these outputs. As discussed further in Issue 5, this is subject to certain conditions (in particular, OpenAI must be able to directly and uniquely associate the requester to the information in question).
157. Similarly, where OpenAI collects personal information from licensed third-party sources, it would generally be for a purpose unrelated to that for which it was originally collected. In such cases, OpenAI would be required to ensure that relevant individuals have been given a choice before it collects that information from the third party.
158. OpenAI submitted that its current data partnership practices¹¹⁶ comply with this requirement for various reasons. In addition to the ability to request removal of verified personal information (see paragraph 156), these include the fact that:
- i. As OpenAI noted, when introducing its OpenAI Data Partnerships,¹¹⁷ is “not seeking datasets with sensitive or personal information, or information that belongs to a third party” and can work with them to remove this information.
 - ii. OpenAI has also been requiring certain contractual assurances from its partners. For example, at the time of training GPT-3.5 and 4, parties to the data partnership agreement had to warrant that they had taken all necessary steps to achieve compliance with Data Protection Laws. OpenAI explained that contractual provisions now include further assurances that any personal information that may be contained in data provided to OpenAI via the data partnership has been publicly published, and that any necessary notices and consents have been obtained, where applicable (e.g., in arrangements with stock photo vendors who provide access to datasets containing photos of consenting models).
159. As mentioned above, OpenAI submitted that at the time of training GPT-3.5 and 4, licensed datasets represented less than 1% of its pre-training datasets. While we were not able to verify this statistic, the Offices acknowledge OpenAI’s representations on this matter and expect it to continue to make reasonable efforts (such as through written agreements with the third parties and associated monitoring to ensure compliance with such agreements) to ensure that the third parties have obtained the personal information in a lawful manner and have the appropriate consent from their users, or that they can rely on a valid exception to the

¹¹⁶ See paragraph 51(ii) for more information about these data partnerships.

¹¹⁷ See [OpenAI Data Partnerships](#), OpenAI, November 9, 2023.

requirement for consent, to share those users' personal information with OpenAI for the purpose of training its models.

Exemption for publicly available information

160. While not specifically claimed by the Respondent, for the reasons below, we note that OpenAI would not be able to rely on the exemption to consent for publicly available information, in respect of the wide array of personal information that it collects from the Internet.

161. PIPEDA, PIPA-BC and PIPA-AB have exceptions to the requirement for consent where the personal information at issue is publicly available as set out in section 7(1)(d) of PIPEDA, sections 12(1)(e), 15(1)(e) and 18(1)(e) of PIPA-BC, and sections 14(e), 17(e) and 20(j) of PIPA-AB.¹¹⁸ The definition of "publicly available" is provided by each Act's regulations¹¹⁹ and is distinct from a common understanding of "publicly accessible" information.

162. Information from sources such as social media or professional profiles, collected from public websites, does not fall under the "publicly available" information exception in PIPEDA.¹²⁰ Similarly, PIPA-BC prescribes sources of public information that include directories, registries, and publications, and social media websites and search engines are not listed as prescribed sources of publicly available information under this Act. Section 7(e)(ii) of the PIPA AB regulation requires that "it is reasonable to assume that the individual that the information is about provided that information." The foregoing evidence above suggests that the indiscriminate scraping of personal information from websites, including social media, by OpenAI will not meet this part of PIPA-AB's requirement in section 7(e) of its regulation given that it is likely that personal information collected from these sites includes that posted by third parties.¹²¹ Therefore, under the regulations, collection from these sources would only be authorized with valid consent and only if the purposes for that collection are what a reasonable person would consider appropriate.¹²²

163. More broadly, given the diverse sources from which OpenAI collects personal information and the fact that such information is used for a purpose unrelated to that for which

¹¹⁸ [Clearview AI v Alberta \(Information and Privacy Commissioner\)](#), 2025 ABKB 287.

¹¹⁹ Section 1 of PIPEDA's [Regulations Specifying Publicly Available Information](#); Section 6 of [PIPA-BC Regulations](#), Prescribed source of public information and Section 7 of [PIPA-AB Regulations](#), Publicly available information.

¹²⁰ See [Company's re-use of millions of Canadian Facebook user profiles violated privacy law](#), OPC, paras 112-113.

¹²¹ In [Clearview AI v Alberta \(Information and Privacy Commissioner\)](#), 2025 ABKB 287 the Court struck the words 'including, but not limited to, a magazine, book or newspaper' from section 7(e) of the Regulation, but did not change the requirements of the PIPA-AB Regulation, sections 7(e)(i) and (ii). See also AB Order P2008-010 at para 60.

¹²² Section 2 of [PIPA-AB](#); Section 6 of [PIPA-BC Regulations](#).

it was initially posted, it would generally not constitute publicly available information as defined in the Acts.¹²³

164. In response to our Preliminary Report, OpenAI stated that the exemption for publicly available information may apply in respect of certain datasets it obtains from data partnerships that may contain personal information (e.g., datasets from news outlets such as the Guardian Media Group or News Corp, see Footnote 60). OpenAI further indicated that this analysis may extend to other comparable datasets that meet the definition of “publicly available information” under the Act.

165. As noted earlier, OpenAI represented that licensed datasets represent less than 1% of the information that OpenAI collects for the purpose of training its models. Given the diversity of the sources from which OpenAI collects personal information, we maintain that such information would generally not constitute publicly available information as defined under the Acts.

Findings related to GPT-3.5 and 4

166. We find that OpenAI did not obtain valid consent for its collection (e.g., via its web crawler GPTBot or from public repositories such as Common Crawl) and use of personal information from publicly accessible websites, to train its GPT-3.5 and 4 models.

167. Consequently, the OPC, OIPC-AB and OIPC-BC find that OpenAI contravened section 6.1 as well as Principle 4.3 of Schedule 1 of PIPEDA, sections 7 and 8 of PIPA-AB, and sections 6-8 of PIPA-BC.

Recent developments and conclusion under PIPEDA

168. As mentioned above, in response to our Preliminary Report, OpenAI informed us that it has recently implemented a tool which it represented would significantly reduce the processing of personal information included in training data collected from publicly accessible Internet data

¹²³ We note that the interpretation of publicly available personal information exceptions in some of the provincial Acts is presently before the courts: [Clearview AI Inc v Alberta \(Information and Privacy Commissioner\)](#), 2025 ABKB 287, appeal to ABCA pending; [Clearview AI Inc. v. British Columbia \(Information and Privacy Commissioner\)](#), 2026 BCCA 67, leave to appeal to SCC pending; [Clearview AI Inc. c. Commission d'accès à l'information du Québec](#), 2025 QCCQ 982.

While the final decisions in these cases will likely impact the scope of information that qualifies as publicly available personal information, this line of jurisprudence will not impact the overall conclusion in this Report that a significant portion of the information OpenAI uses to train its model does not qualify as publicly available personal information.

and data from partnerships. More specifically, OpenAI stated that this tool can detect and mask a wide range of identifying information about individuals (such as names and phone numbers), thereby ensuring that masked information is not used to train the models. OpenAI further submitted that the tool uses context to detect whether information is private or personal in nature, and can distinguish between information about private individuals, public figures or fictional characters. OpenAI also indicated that the tool can detect categories of personal information that are private or personal but that it has not been trained to recognize.

169. As mentioned at paragraph 110, OpenAI provided the Offices with the results of their internal evaluations, which demonstrate the effectiveness of the tool in correctly identifying instances of personal information and aligning with human judgment.

170. The OPC accepts that this new tool, combined with OpenAI's other mitigation measures implemented at the various stages of development and deployment of ChatGPT (detailed in various sections of this report and listed in Appendix A) can significantly reduce the risk that the personal information of private individuals (including information not posted by the individuals themselves), and sensitive information more specifically, will be included in the datasets used to train OpenAI's AI models moving forward.

171. In making this determination, as further discussed in other sections of this report¹²⁴, the OPC also considered OpenAI's additional commitments with respect to openness and model transparency (including the publication and promotion of a Canadian blog post explaining its privacy practices) and its decision to deprecate GPT-3.5 and 4 and to fully train the current models powering ChatGPT with the new mitigation measures.

172. The OPC also acknowledge that the context surrounding generative AI has significantly evolved since the release of ChatGPT in November 2022, with a rapid increase in usage and broader consumer adoption of ChatGPT, and LLMs in general. While this was not the case when ChatGPT was first released, search engines and LLMs are now becoming increasingly integrated, with AI overviews regularly associated with search engines responses. Therefore, individuals are now more likely to have heard of, if not used, this technology and gained a basic knowledge of how AI models are trained.

173. As noted in a recent Supreme Court of Canada decision, statutes should not be interpreted as "frozen in time" but rather as capable of "evolv[ing] with technology" consistent with the purpose of legislation. The Court referred to this mode of statutory interpretation as "dynamic interpretation."¹²⁵

¹²⁴ In particular, see section relating to OpenAI's response to our Preliminary Report.

¹²⁵ *Telus Communications Inc. v. Federation of Canadian Municipalities*, [2025 SCC 15 \(CanLII\)](#) at paras. 32-36, 155. This decision is conceptually consistent with the Supreme Court of Canada's decision in *R. v. Tessling*, [2004 SCC 67](#), (CanLII) where the Court held that privacy is a "protean concept", and, relatedly, that reasonable expectations can fluctuate along with technological change, (paras. 25-29).

174. Given that context, together with the implementation of OpenAI's existing and new mitigation measures (including those set out in paragraph 156), which appear to materially reduce the impact on private individuals' privacy, we accept that individuals are now more likely to expect that OpenAI's future models will be trained using publicly accessible information.
175. The OPC also recognizes that despite the implementation of these combined mitigation measures, OpenAI's future training datasets will still likely include personal information of private individuals, some of which may have been posted years ago and/or by third parties without the knowledge of concerned individuals. However, the OPC accepts that this may represent a small subset of information. Furthermore, in line with a pragmatic and flexible interpretation of the Acts (discussed at paragraph 46), and the necessity to balance the privacy rights of individuals with the right to freedom of expression and the need to facilitate the use of personal information for appropriate commercial purposes, the OPC accepts that where the risk to privacy is significantly and meaningfully mitigated (including by training models to refuse to provide private or sensitive information in their outputs), OpenAI may rely on implied consent in this context.
176. Consequently, subject to the above, the OPC finds this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.
177. This conclusion is based on our expectation that OpenAI will continue to effectively implement and improve its mitigation measures as described to the Offices and develop further innovative privacy-protecting techniques to effectively address evolving privacy risks posed by the products it offers in future.

Conclusion under PIPA-BC and PIPA-AB

178. Turning now to our findings under PIPA-BC and PIPA-AB, additional commentary and analysis are warranted. Both PIPA legislations are deemed substantially similar to Part 1 of PIPEDA by an Order of the Governor General in Council.¹²⁶ The effect of this designation and status is that PIPA-BC and PIPA-AB, but not PIPEDA, will apply to the collection, use, and disclosure of personal information within BC and Alberta if the relevant organization is not a federal work, undertaking, or business.¹²⁷ However, the "substantially similar" status does not require that every provision of each respective Act be interpreted in an identical way, nor does it mean that there are no meaningful differences between the Acts.
179. One meaningful difference between both PIPA-BC and PIPA-AB, and PIPEDA relates to situations where an organization collects, uses, and discloses personal information without express consent from the relevant individual. If the organization does not have a separate legal

¹²⁶ See Organizations in the Province of British Columbia Exemption Order, SOR/2004-220 and Organizations in the Province of Alberta Exemption Order SOR/2004-219.

¹²⁷ This status is discussed in more detail at the section of this report that responds to OpenAI's jurisdictional arguments.

authority to engage in that activity without consent, then the organization must establish that it has another valid form of consent despite the lack of express consent.

180. In the analysis above, the OPC determines that, in the context of dataset training activity, OpenAI may rely on “implied consent” for the collection and use of personal information from publicly accessible websites and licensed third-party sources where the risk to privacy is significantly and meaningfully mitigated (including by training models to refuse to provide private or sensitive information in their outputs). This determination supports the OPC’s finding that the issue is conditionally resolved under PIPEDA.

181. Unlike PIPEDA, neither PIPA-BC nor PIPA-AB contain the term “implied consent”. Instead, s. 8 of PIPA-BC uses the term “implicit consent” and s. 8(2) of PIPA-AB uses the term “deemed consent” or consent by “notice” to refer to valid consent that is not expressly given. These provisions also impose the following discrete requirements for organizations seeking to rely on implicit or deemed or notice consent when collecting, using, and disclosing personal information:¹²⁸

Implicit consent (PIPA-BC)

- 8 (1) An individual is deemed to consent to the collection, use or disclosure of personal information by an organization for a purpose if
- (a) at the time the consent is deemed to be given, the purpose would be considered to be obvious to a reasonable person, and
 - (b) the individual voluntarily provides the personal information to the organization for that purpose.
- (2) An individual is deemed to consent to the collection, use or disclosure of personal information for the purpose of his or her enrolment or coverage under an insurance, pension, benefit or similar plan, policy or contract if he or she
- (a) is a beneficiary or has an interest as an insured under the plan, policy or contract, and
 - (b) is not the applicant for the plan, policy or contract.
- (3) An organization may collect, use or disclose personal information about an individual for specified purposes if
- (a) the organization provides the individual with a notice, in a form the individual can reasonably be considered to understand, that it intends to collect, use or disclose the individual's personal information for those purposes,

¹²⁸ *Personal Information Protection Act*, SBC 2003, c 63, s 8 and *Personal Information Protection Act*, SA 2003, c P-6.5, s.8.

(b) the organization gives the individual a reasonable opportunity to decline within a reasonable time to have his or her personal information collected, used or disclosed for those purposes,

(c) the individual does not decline, within the time allowed under paragraph (b), the proposed collection, use or disclosure, and

(d) the collection, use or disclosure of personal information is reasonable having regard to the sensitivity of the personal information in the circumstances.

(4) Subsection (1) does not authorize an organization to collect, use or disclose personal information for a different purpose than the purpose to which that subsection applies.

Form of consent (PIPA-AB)

8 (1) An individual may give his or her consent in writing or orally to the collection, use or disclosure of personal information about the individual

(2) An individual is deemed to consent to the collection, use or disclosure of personal information about the individuals by an organization for a particular purpose if

(a) The individuals, without actually giving consent referred to subsection (1), voluntarily provides the information to the organization for that purpose, and

(b) It is reasonable that a person would voluntarily provide that information

(2.1) If an individual consents to the disclosure of personal information about the individual by one organization to another organization for a particular purpose, the individual is deemed to consent to the collection, use or disclosure of the personal information for the particular purpose by that other organization.

(2.2) An individual is deemed to consent to the collection, use or disclosure of personal information about the individual by an organization for the purpose of the individual's enrolment in or coverage under an insurance policy, plan or contract that provides for a similar type of coverage or benefit if the individual

(a) has an interest in or derives a benefit from that policy, plan or contract, and

(b) is not the applicant for the policy, plan or contract.

(3) Notwithstanding section 7(1), an organization may collect, use or disclose personal information about an individual for particular purposes if

(a) the organization

(i) provides the individual with a notice, in a form that the individual can reasonably be expected to understand, that the organization intends to collect, use or disclose personal information about the individual for those purposes, and

(ii) with respect to that notice, gives the individual a reasonable opportunity to decline or object to having his or her personal information collected, used or disclosed for those purposes,

(b) the individual does not, within a reasonable time, give to the organization a response that notice declining or objecting to the proposed collection, use or disclosure, and

(c) having regard to the level of sensitivity, if any, of the information in the circumstances, it is reasonable to collect, use or disclose the information as permitted under clauses (a) and (b).

182. In considering OpenAI's collection and use of personal information from publicly accessible websites and licensed third-party sources, it is not apparent to the OIPC-BC or the OPIC-AB that OpenAI has met any of the requirements for implicit consent under s. 8 of PIPA-BC or deemed or notice consent under s. 8 of PIPA-AB:

- i. OpenAI has not established that the personal information from publicly accessible websites and licensed third-party sources was provided to OpenAI by the relevant individuals, nor that the individuals provided their personal information for the purpose of training OpenAI's ChatGPT datasets, as required by s. 8(1)(b) of PIPA-BC and s. 8(2)(a) and (b) of PIPA-AB;
- ii. OpenAI has not, and does not propose to, collect, use, and disclose individuals' personal information for the purpose of those individuals' enrolment or coverage under an insurance, pension, benefit or similar plan, policy or contract under s. 8(2) of PIPA-BC and s. 8(2.2) of PIPA-AB; and
- iii. OpenAI has not established that the relevant individuals received a notice from OpenAI that OpenAI intended to collect, use, or disclose those individuals' personal information for specified purposes, as required by s. 8(3)(a) of PIPA-BC and s. 8(3) of PIPA-AB.

183. PIPA-BC and PIPA-AB do not enable an organization to establish implicit or deemed or notice consent, respectively, on the basis of shifting societal expectations about artificial intelligence and the organization's steps to significantly and meaningfully mitigate risks to privacy. Therefore, while OpenAI's actions to address risks to personal privacy supported the OPC's determinations with respect to implied consent, those actions are not sufficient to establish implicit consent under PIPA-BC or deemed or notice consent under PIPA-AB.

184. In light of the specific requirements to establish implicit consent under PIPA-BC and deemed or notice consent under PIPA-AB, and OpenAI's failure to establish that its activities

meet those criteria, OpenAI may not rely on implicit consent under PIPA-BC or deemed or notice consent under PIPA-AB to collect and use personal information from publicly accessible websites and licensed third-party sources while training its models, where such training is done in the same manner that OpenAI trained the GPT-3.5 and 4 models.

185. Consequently, the OIPC-BC and OIPC-AB find this aspect of the complaint to be **well-founded and unresolved** under PIPA-BC and PIPA-AB.

186. Regarding future and unexamined ChatGPT models, the OIPC-BC and OIPC-AB both decline to draw a conclusion about the future collection and use of personal information and whether such activities meet the requirements of implicit consent under PIPA-BC or deemed or notice consent under PIPA-AB. The OIPC-BC directs OpenAI to the specific requirements of implicit consent under s. 8 of PIPA-BC, which apply to any collection or use of personal information on the basis of implicit consent. Similarly, the OIPC-AB directs OpenAI to the specific requirements of deemed consent under s. 8 (2) and notice consent under s. 8(3) of PIPA-AB which apply to any collection or use of personal information on the basis of deemed or notice consent.

Analysis under Quebec’s Private Sector Act

187. The consent rules set out in Quebec’s Private Sector Act differ from the analytical framework used by the OPC, the OIPC-AB and OIPC-BC, particularly in that Quebec’s Private Sector Act does not specifically provide for a criterion equivalent to “for purposes that a reasonable person would consider are appropriate in the circumstances”.

188. In particular, under Quebec’s Private Sector Act, an enterprise that collects personal information directly from an individual 14 years of age or over is subject to a duty to inform.¹²⁹

189. However, the concept of consent remains central to the Act, because in practice it is consent that allows individuals to exercise control over the use and communication of their personal information.¹³⁰

190. For example, when information is collected from a third party rather than from the individual concerned, or from an individual under 14 years of age, that person’s consent is generally required.¹³¹ Specifically, for an individual under 14, consent must be obtained from the person having parental authority or of the tutor.

¹²⁹ Section 8 of Quebec’s Private Sector Act.

¹³⁰ Section 8 subsection 4 of Quebec’s Private Sector Act.

¹³¹ Section 6 of Quebec’s Private Sector Act.

191. The collection of personal information on the web does not occur in a single context and can involve multiple variables that must be analyzed and taken into account in order to determine the validity of the consent given by the person concerned in a given context.

192. Similarly, certain secondary uses and certain communications to third parties may require consent.¹³²

193. A company that collects, uses, or communicates personal information should be able to sufficiently document the context in which the information was collected to ensure, depending on the situation, that the duty to inform has been properly met or that valid consent has been obtained.

194. In its submissions, OpenAI stated that it collected publicly accessible information, which incidentally could include personal information, for the purpose of training the GPT-3.5 and 4 models.

195. Whether or not the collection was incidental has no bearing on an enterprise's obligation to comply with Quebec's Private Sector Act, since personal information was in fact collected.

The exception relating to journalistic, historical, and genealogical material provided for in section 1, paragraph 4 of Quebec's Private Sector Act

196. In the context of its observations to the Offices, OpenAI raised an alternative argument that, generally speaking, the personal information it collected from publicly accessible websites and through data exchange agreements would fall under the exception provided for in paragraph 4 of section 1 of Quebec's Private Sector Act, namely the exception relating to journalistic, historical or genealogical material for the legitimate information of the public.

197. Although this is an alternative argument, since OpenAI alleges that this exception applies generally to the material it collects, uses and discloses, the CAI will first analyze this exception before addressing the rules relating to the duty to inform and consent.

198. OpenAI states that this exception would apply to the collection and use of publicly accessible informational content for the purposes of research, training, and development of general-purpose artificial intelligence ("GPAI") models.¹³³

¹³² Sections 12 and 13 of Quebec's Private Sector Act.

¹³³ OpenAI's response dated September 2, 2025, to the preliminary investigation report (PIPEDA-045141; P23-93177; 030562; 1032001-S) at p. 23 [free translation].

199. OpenAI also states that:
- i. Its training activities involving the collection and processing of such information enable its models to contextualize user prompts, understand common references in human communication, answer questions, participate in discussions, provide relevant responses to prompts, and are consistent with the scope and purpose of this exemption.
 - ii. These activities are expressive in nature, promote access to knowledge, support the legitimate informing of the public in areas such as current events, history, science, and culture, and facilitate a broader public understanding of important social issues.
 - iii. This exception should be interpreted in a manner that reconciles the Act's objective of protecting personal information with the constitutional protection of freedom of expression.
 - iv. This exception should be interpreted broadly and liberally to encompass activities such as training GPAI models.¹³⁴
200. Section 1, paragraph 4 of Quebec's Private Sector Act provides that the collection, retention, use, or communication of journalistic, historical, or genealogical material for the legitimate information of the public is excluded from the scope of the Act.
201. The CAI agrees from the outset that, depending on the context, some of the information collected and used by OpenAI to train its models could fall under the scope of this exception.
202. For example, the CAI recognizes that this exception may apply to datasets obtained through partnership agreements with recognized news and newspaper publishers.
203. However, the CAI does not share OpenAI's position that all of its activities related to training its models are consistent with the scope and purpose of this exception.¹³⁵
204. In its position, OpenAI refers to the three-part test developed in *Institut généalogique Drouin c. Commission d'accès à l'information du Québec*¹³⁶ and reiterated in *Morin-Lachance c. La Presse inc.*¹³⁷
205. According to this test, to be covered by the exception in section 1, paragraph 4 of Quebec's Private Sector Act, the activity must consist of:
- i. The collection, retention, use or communication;
 - ii. Of journalistic, historical, or genealogical material;

¹³⁴ *Ibid.*

¹³⁵ *Ibid.*

¹³⁶ 2021 QCCQ 557 et 2021 QCCS 5806.

¹³⁷ 2025 QCCA, 132.

iii. For the legitimate information of the public.

206. However, the CAI notes that the context of this case differs from the contexts in which this exception was analyzed in *Institut généalogique Drouin* and *Morin-Lachance*.

207. In *Institut généalogique Drouin*, the genealogical nature of the material was not in question, and the analysis focused more specifically on the concept of legitimate public information.¹³⁸ In *Morin-Lachance*, the journalistic nature of the material was clear, whereas in this case, it is precisely this concept of journalistic material that must be analyzed.

208. The concept of journalistic material is not defined in the Act.

209. OpenAI argues that the purpose of this exception is to ensure that any information serving the legitimate purpose of informing the public remains accessible and calls for a broad interpretation of the scope of this exception.

210. On the other hand, the CAI considers that OpenAI's interpretation is too broad and that it would significantly restrict the protections afforded by the Act.

211. The CAI believes instead that each situation must be interpreted in light of a specific context and that this exception cannot be applied in a "general" manner, since the material in question must be able to qualify, on a case-by-case basis, as journalistic, historical, or genealogical material.

212. Regarding journalistic material, the criteria developed in *A.T. v Globe24h.com*¹³⁹ and reiterated in other Canadian decision¹⁴⁰ define what may constitute journalistic material.

213. Under this test, in order to qualify as journalistic material, the activity in question must:

- i. be intended to inform the community on issues the community values;
- ii. involve an element of original production; and
- iii. be guided by a self-conscious discipline calculated to provide an accurate and fair description of facts, opinion and debate at play within a situation.

¹³⁸ 2021 QCCQ 557, para. 50 and 51, 2021 QCCS 5806, para. 43 to 46.

¹³⁹ [2017] 4 F.C.R.310

¹⁴⁰ Google LLC v. Canada (Privacy Commissioner), [2023 FCA 200](#), para. 90, and [2021 FC 723](#), para. 83., *Luminos Consulting & Production Inc. (Re)*, 2021 [CanLII 88596](#) (AB OIPC), para. 23.

214. The CAI believes that not all information found on the web can qualify as journalistic material under this test.

215. In this regard, the CAI agrees with the Alberta Court of Appeal when it states, in its analysis of a similar exception, that it is unreasonable to think that the Legislature intended this exception to be so wide as to encompass everything that could qualify as freedom of opinion and expression, and that not every piece of information posted on the Internet qualifies as journalism¹⁴¹.

216. Furthermore, the CAI understands from OpenAI's arguments that OpenAI considers the final result of the use of its GPT-3.5 and 4 models to constitute journalistic material in itself. However, the CAI does not share this opinion.

217. It is true that ChatGPT's features allow users to learn about topics that are of interest to them.

218. However, the approach taken by these models, which is based on a statistical analysis of a large volume of data, without any focus or consideration of the issues that really matter to communities, is far removed from the journalistic approach.

219. These tools are designed to answer questions that individuals ask themselves on an individual basis, rather than to inform a community about issues the community values.

220. To this end, the activities of the GPT-3.5 and 4 models do not meet the first part of the test.

221. Regarding the second criterion, although the GPT-3.5 and 4 models can, to a certain extent, create content from training data, the CAI finds that these results cannot be considered an original production in the journalistic sense of the term.

222. In fact, these results come from statistics that models construct from training data in order to determine which words are most likely to appear in a given sentence, for the purpose of responding to user prompts.

¹⁴¹ *United Food and Commercial Workers, Local 401 v Alberta (Attorney General)*, [2012 ABCA 130](#) (CanLII), 522 A.R. 197, para. 56 and 59, and *A.T. v. Globe24h.com* [2017] 4 F.C.R. 310, para. 69.

223. This process is based on a statistical approach to language and does not involve the type of editorial judgment that usually guides the original production of journalistic content.¹⁴²
224. In any event, the probabilistic nature of these models and the issues surrounding the accuracy of the answer they provide, as presented in issue 4 of this report, mean that the information provided by these models cannot be classified as journalistic material under the third criterion of the test developed in *Globe24h*.¹⁴³
225. The methodology used to train these models does not reflect the methodology of journalism, which tends to produce an accurate and fair description of the facts, opinions, and debates that a situation may bring.
226. Although the training of the GPT-3.5 and 4 models differs from the referencing techniques used by Google, among others, the CAI nevertheless considers that OpenAI does not have complete control over the content of the results displayed by its models.¹⁴⁴
227. Although this content may be influenced by the data and training methods, the fact remains that this process involves an element of randomness over which OpenAI has no real control.
228. Regarding OpenAI's argument that this exception under Quebec's Private Sector Act should be interpreted in light of the constitutional values that support freedom of expression and the right of access to information, the CAI considers that the interpretation of the application of this exception provided for in Quebec's Private Sector Act is not a matter of discretion requiring a balancing of competing values, but rather a matter of interpretation of the Act.¹⁴⁵
229. The interpretation exercise required by this exception consists of verifying whether the factual circumstances necessary for the application of the standard are present and do not leave the decision maker with a choice as to the suitability of the measures to be taken.¹⁴⁶ Consequently, the power exercised by the CAI in interpreting this exception cannot be characterized as discretionary.

¹⁴²Ethics Advisory Committee of the Canadian Association of Journalists: *What is journalism?*, Update October 2021 and approved by the CAJ board December 2021: [caj_whatismjournalism_october2021](#)

¹⁴³ *Globe24h.com* [2017] 4 F.C.R. 310, para. 68.

¹⁴⁴ *Google LLC v. Canada (Privacy Commissioner)*, 2021 FC 723 at para. 82.

¹⁴⁵ *Clearview AI Inc. v. Information and Privacy Commissioner for British Columbia*, 2024 BCSC 2311 (CanLII), para. 229 to 239, aff'd 2026 BCCA 67, leave to appeal to SCC pending, and *Ontario Nurses' Association v. 10 Community Care Access Centres*, 2021 ONSC 5348 (CanLII), para. 105 and 106, and *Doré v. Barreau du Québec*, 2012 SCC 12 (CanLII), [2012] 1 SCR 395, para. 60.

¹⁴⁶ *Shiler c. Bousquet*, 2017 QCCA 276, para. 35 and 36, *Baker v. Canada (Minister of Citizenship and Immigration)*, 1999 CanLII 699 (SCC), [1999] 2 SCR 817, para. 52.

230. Ultimately, considering that not all of the information collected, used, and disclosed by OpenAI is excluded from the scope of Quebec's Private Sector Act, with the exception of material collected that actually qualifies as journalistic material, OpenAI must comply with the rules relating to the duty to inform and consent provided for in the Act.

Rules relating to the duty to inform that apply in the context of using personal information for primary purposes

231. Section 6 of Quebec's Private Sector Act provides, subject to certain exceptions,¹⁴⁷ that as a general rule, any person who collects personal information about another person should collect it from the person concerned, unless that person consents to it being collected by a third party.¹⁴⁸

232. When the collection is made directly from the person concerned, section 8 of Quebec's Private Sector Act provides for the obligation to inform that person of the following:

- i. The purposes for which the information is being collected;
- ii. The means by which the information will be collected;
- iii. The rights of access and rectification provided for by law;
- iv. Their right to withdraw their consent to the disclosure or use of the information collected.

233. Similarly, paragraph 2 of section 8 provides that, at the time of collection, the person concerned must be informed of the third persons or categories of third persons to whom it is necessary to communicate the personal information collected for the purposes disclosed at the time of collection.

234. The last paragraph of section 8 states that information, including the purposes for which their personal information is collected, must be provided to the person concerned in simple and clear language.

235. Section 8.3 provides that any person who provides their personal information, after having been informed in accordance with section 8, consents to the use of that personal information and its communication for the purposes disclosed to them.

236. These rules, found in Division II of Quebec's Private Sector Act entitled "Collection of personal information" are intended for a "primary" purpose that justifies the use or communication of the personal information collected.

¹⁴⁷ CQLR, chapter P-39.1, s. 6 para. 2 and 3.

¹⁴⁸ Pursuant to section 4.1, this general rule does not apply when the individual concerned is under 14 years of age.

237. Regarding the use of personal information within the company, section 12 of Quebec's Private Sector Act provides that personal information may only be used within the company for the purposes for which it was collected, unless the person concerned has given their consent.

Rules relating to consent that apply to the use of personal information for secondary purposes

238. Section 12 of Quebec's Private Sector Act provides for specific situations in which personal information may be used for "secondary" purposes without the consent of the person concerned. However, the CAI considers that the exceptions provided for in this section do not apply in this case.

239. Regarding use for a "secondary" purpose, section 13 of Quebec's Private Sector Act provides that no person may communicate to a third person the personal information he holds on another person, unless the person concerned consents to, or this Act provides for, such communication.

240. Similarly, this section specifies that in this context, consent must be given expressly when it concerns sensitive personal information.

241. Finally, section 14 states that consent under this Act must be clear, free and informed and be given for specific purposes. It must be requested in clear and simple language and, if made in writing, it must be presented separately from any other information.

The rules of consent that apply to personal information concerning individuals under 14 years of age

242. Section 4.1 of Quebec's Private Sector Act provides that personal information concerning a minor under 14 years of age may not be collected from him directly without the consent of the person having parental authority or of the tutor, unless collecting the information is clearly for the minor's benefit.

243. Section 14 of Quebec's Private Sector Act states that the consent of a minor 14 years of age or over is given by the minor, by the person having parental authority or by the tutor.

244. Finally, as stipulated in the guidelines regarding the interests of young people in matters of privacy and access to personal information, the CAI, like many other Offices, consider that the personal information of children is particularly sensitive.¹⁴⁹

Information OpenAI collects from public sources

245. As specified in paragraphs 51 and following, it is first established that OpenAI used information collected through data scraping for the purpose of training its GPT-3.5 and 4 models,

¹⁴⁹ [Putting best interests of young people at the forefront of privacy and access to personal information](#), OPC et al., 2023.

including data scraping using the GPTBot tool or information made accessible on open information sources such as Common Crawl and Wikipedia.

246. Regarding this collection, OpenAI specified that:

- i. When collecting publicly accessible data, it did not bypass paywalls or password-protected accounts.
- ii. It did not collect information from the “dark web” or closed discussion groups.

247. In addition, OpenAI stated that it considered that:

- i. In situations where individuals concerned had made personal information publicly accessible without restriction, the expectation of privacy regarding that information was diminished and, in such circumstances, tacit or implied consent could be inferred for subsequent access and reuse by a third party for legitimate purposes.
- ii. In this context, information made publicly accessible did not usually meet the definition of sensitive personal information and, as a result, such information generally carried a reduced expectation of privacy.

248. On this point, Quebec’s Private Sector Act does not confer public status on personal information solely because it is published on the web.¹⁵⁰ Paragraph 5 of section 1 of Quebec’s Private Sector Act provides for an exception concerning personal information that is public **under the Act**, but this exception does not apply in this case.

249. The CAI believes that the appropriate approach in this context should be more nuanced and that it should not necessarily be inferred from the fact that personal information is published on the web without restriction that the person concerned has, depending on the context, been properly informed of the use or communication that could be made of it or that consent has been duly obtained.

250. In the context of data scraping on the Internet, it is only in rare exceptions that the collection can be considered to have been carried out directly from the person concerned. This situation may occur, in particular, when the harvesting is carried out directly on a Web page that actually belongs to the person concerned or when the terms of service specifically provide for this.

251. When it comes to social networks, in general and subject to terms of service that would have the opposite effect, the information published is subject to licenses granted by users to allow these social networks to publish this information. Generally, in this context, this information cannot be considered to have been collected directly from the individual, and the rules concerning collection from third parties usually apply.

¹⁵⁰ *Clearview AI inc.*, CAI file No 1023158-S, December 14, 2021, para. 85 and 86 (appealed to the Court of Quebec).

252. However, it is entirely possible that users, based on the information provided at the time of collection, the terms of service, and the privacy policies in effect, may be duly informed and thereby consent to their personal information being made accessible on the web and, as a result, communicated to third parties and harvested for the purpose of training artificial intelligence models. However, this analysis must be carried out on a case-by-case basis, taking into account all the factors involved.
253. That said, the evidence on file and OpenAI's submissions do not comprehensively establish the Internet sources from which the GPTBot tool collected personal information and on which the Common Crawl data archive is based.
254. For example, Wikipedia's current privacy policy¹⁵¹ specifies that when a person makes a contribution to this site, it creates a permanent public record of each piece of content added, removed or altered by the user.
255. We may ask ourselves whether a disclosure to users that their personal information will be made accessible on the web is sufficiently specific to allow collection via data harvesting and use for training artificial intelligence models.
256. To this end, the CAI believes that the best practice in the circumstances would be to be specific about the consequences of making personal information publicly accessible on the web and to specify that this may include data harvesting for the purpose of training artificial intelligence models.
257. Despite this, the CAI still considers that a reasonable and well-informed person, subject to their age, to the information provided and in particular to its simplicity and clarity, is aware of the consequences of making their personal information accessible on the web and of the fact that, once published, it could be used by third parties for other legitimate purposes.
258. Regarding OpenAI's argument that the issue of consent requires an interpretation that reconciles the values protected by the Charter, including freedom of expression, the CAI considers that the question of the duty to inform under section 8 of Quebec's Private Sector Act does not involve a discretionary decision that leaves the decision maker with a choice of options within the limits imposed by the Act.¹⁵² The question of whether or not the person concerned was properly informed of the purposes for which their personal information was collected is a question of fact and not a question that allows for a balanced interpretation under the values protected by the Charter.

¹⁵¹ [Wikimedia Foundation Privacy Policy - Wikimedia Foundation Governance Wiki](#).

¹⁵² *Baker v. Canada (Minister of Citizenship and Immigration)*, 1999 CanLII 699 (SCC), [1999] [2 SCR 817](#), para. 52.

259. Depending on the context, it is also possible that personal information made accessible on the web, including on Wikipedia, concern an individual under 14 years of age or may have been published by a third party, rather than the person concerned, without consent.
260. On this subject, OpenAI's position is that when personal information is published online by a third party, it would be reasonable in this context to assume that the publication of this information has been authorized by the person concerned.
261. The CAI considers that OpenAI should instead take into account the general context of the publication to ensure that the personal information concerned is not communicated without consent and, in case of doubt, refrain from collecting it, rather than relying on such a presumption.
262. Similarly, the CAI considers that the company should also verify that the publication of such information does not constitute communication without consent in the context of a "secondary" use of personal information initially collected for another purpose by a third party, or the publication of personal information concerning a person under the age of 14 without the consent of the person having parental authority or of the tutor.
263. Regarding the issue of so-called "secondary" purposes, OpenAI argued that section 13 of Quebec's Private Sector Act did not apply to its collection of personal information accessible on the web or to personal information it obtained through data exchange partnership agreements.
264. It alleges that interpreting this section as applicable in this case would be tantamount to confusing the spheres of collection and communication, would render redundant the specific provisions of the Act governing collection, would undermine the consistency of the sections of the Act, and would be contrary to the principles of legislative interpretation.
265. The CAI believes, on the contrary, that section 13 of Quebec's Private Sector Act complements the rules set out in sections 6, 8 and following of the Act by specifying that, in the context of the use of personal information for so-called "secondary" purposes, no person may communicate to a third person information he holds on another person unless the person concerned consents to such communication, unless there is an exception specifically provided for in the Act.
266. To this end, during parliamentary debates on Bill 64, it was clarified that section 13 of Quebec's Private Sector Act does not apply to "first-hand" collection, but rather to situations where a third party wishes to communicate such information to another party.¹⁵³

¹⁵³ Quebec, National Assembly, Commission de la Culture, [Clause-by-clause consideration of Bill 64 in Hansard, 42nd leg., 1st sess., vol. 45, No. 151, May 27, 2021, 11:40 a.m.-12:10 p.m.](#) (in French only).

267. The CAI considers that OpenAI's proposed interpretation, according to which the collection and communication of personal information to third parties are distinct and separate concepts, runs counter to the fact that personal information can have several life cycles with various organizations.
268. In this case, it is not the interpretation of the Act that creates a tautology, but rather a repetition in fact. When personal information is communicated by a third party, that same information is necessarily collected by the party receiving the communication, and as a result, a new life cycle begins with that new organization. The counterpart to the communication of personal information to a third party can only be the collection of that information by that same third party.¹⁵⁴
269. Similarly, sections 7 and 15 of Quebec's Private Sector Act illustrate the connection between provisions concerning the collection and disclosure of personal information¹⁵⁵.
270. Depending on the context, posting personal information on the web could constitute a "secondary" use of information collected for another purpose and constitute disclosure without consent under section 13 of Quebec's Private Sector Act.
271. Similarly, OpenAI's argument ignores the fact that personal information held by OpenAI for training purposes may be communicated to a user by the model.
272. OpenAI was unable to demonstrate that, in the context of the personal information it collected from publicly accessible sources on the web, it had sufficiently documented the context in which the information obligation was fulfilled with the data subject or, where required, that consent was obtained, all for the purpose of ensuring that such collection was in compliance with Quebec's Private Sector Act.
273. The CAI concludes that OpenAI should implement additional checks on the sources from which personal information is collected, in order to, depending on the context:

¹⁵⁴ Quebec, National Assembly, Hansard of the Commission on institutions, 42-1, Thursday May 13, 2021- Vol. 45 No. 149, Clause-by-clause consideration of Bill 64, An Act to modernize legislative provisions as regards the protection of personal information, at 1:10 p.m.: [translation] **Mr. Caire** : *"Yes, it is yes, because every transaction puts us in a situation where there is a source and someone who collects. So even if the information that company A communicates to company B. . . even if it was A who collected it from the initial source, which is the person themselves, it doesn't change the fact that, from the moment I communicate it from A to B, B becomes the person who collects it. . ."*

¹⁵⁵ Section 7 of Quebec's Private Sector Act requires any person collecting personal information from another person carrying on an enterprise to, at the request of the person concerned, inform the latter of the source of the information and section 15 states that consent to the communication of personal information by a third person may be given by the person concerned to the person who collects the information from the third person.

- i. Ensure that data subjects were clearly informed at the time of initial collection that by providing their personal information, it would be made public and could therefore be collected and used by third parties, in particular for the purpose of training artificial intelligence models, and communicated by them; and
- ii. Ensure that the communication of such personal information does not constitute communication by a third party without consent or communication of personal information concerning a person under the age of 14 without the consent of the person having parental authority or of the tutor.

274. In the absence of evidence that such checks were carried out, the CAI concludes that OpenAI has not been able to demonstrate that its practices comply with sections 6, 12(1), 13 and 14 of Quebec's Private Sector Act, relating to the collection of personal information from publicly accessible Internet sources and authorized third parties, the use of such information for the purpose of training the GPT-3.5 and 4 models, and the communication of such personal information to users of these models¹⁵⁶.

Information collected under partnership agreements

275. As already mentioned in paragraph 51, OpenAI also collected information through partnership agreements with content providers (i.e., licensed third-party sources).

276. More specifically, the evidence shows that OpenAI collected training data from, among others, various media outlets, a large stock image vendor, and other sources of specialized knowledge.

277. In this context, OpenAI specifies that it ensures, as part of its data exchange agreements with third parties, that appropriate notice has been provided and consent has been obtained, where applicable.

278. However, in light of the submissions and the copy of the partnership agreement obtained, it is clear that the contractual terms used by OpenAI have evolved and been refined over time.

279. [Text redacted]

280. Regarding its current agreements with third-party providers, OpenAI states that when it enters into data exchange agreements with third parties, it specifies in its agreements that it does not seek data sets that contain sensitive personal information or information belonging to third parties.

¹⁵⁶ To avoid repetition in the reasons, the issue of OpenAI's communication of personal information, which is addressed in issue 2C of this report, has been included in this section.

281. [Text redacted]

282. Finally, OpenAI states that it applies the mitigation measures outlined in the table found in Appendix A of this report to the datasets obtained from its partners.

283. As already mentioned in the analysis of the exception relating to journalistic, genealogical, and historical material, the CAI considers that some of the information collected from these partners could qualify as journalistic, genealogical, or historical material and thus be excluded from the scope of Quebec's Private Sector Act.

284. However, as specified, this analysis must be carried out in a specific manner, and considering that these various partners have not been specifically identified and that the sources of the information they hold and communicate are not known, this exclusion cannot be invoked in a general manner.

285. In addition to the possible application of this exception, regarding personal information collected under partnership agreements, considering the specifications added to the agreements over time, the CAI considers that the contractual measures currently taken by OpenAI and the verifications that it states it performs on the validity of consent can be considered reasonable in the circumstances and encourages OpenAI to implement any other measures that would further ensure compliance with Quebec's Private Sector Act, including the audits recommended in the above section.

286. In conclusion, the CAI considers, subject to evidence to the contrary, that OpenAI's current contractual practices regarding the collection of personal information through data exchange partnership agreements comply with the rules relating to the duty to inform and consent set out in Quebec's Private Sector Act and adopts the recommendation made by the other Offices at paragraph 159.

Issue 2B: Did OpenAI obtain valid consent and meet its obligation to inform users with respect to the collection and use of personal information included in their interactions with ChatGPT?

Analysis under PIPEDA, PIPA-BC and PIPA-AB

287. As explained above, OpenAI collects and uses a subset of users' interactions with ChatGPT to fine-tune its models.

288. For the reasons set out below, we do not accept OpenAI's assertion that it could rely on implied consent for the collection and use of the personal information included in its users' interactions for the training of its GPT-3.5 and 4 models. It should have obtained express consent for that practice, which involved sensitive information and/or which users could not reasonably expect.

289. OpenAI represented that consent for the collection and processing of personal information included in users' interactions with ChatGPT is based on users' positive action – i.e., the voluntary provision of personal information to OpenAI. The organization takes the position that this consent is sufficiently informed given the context in which the information is collected, as well as the detailed explanations provided in the Privacy Policy, Terms of Use, contextual notices and other resources such as the Help Center informing users about the purposes for which their information will be used.

290. Specifically, OpenAI stated that it provides the following notice to users at the onboarding stage: "Don't share sensitive info. Chats may be reviewed and used to train our models – Learn more." Clicking on "Learn more" leads users to a more detailed Help Center article "[How your data is used to improve model performance](#)" (see figure 2).

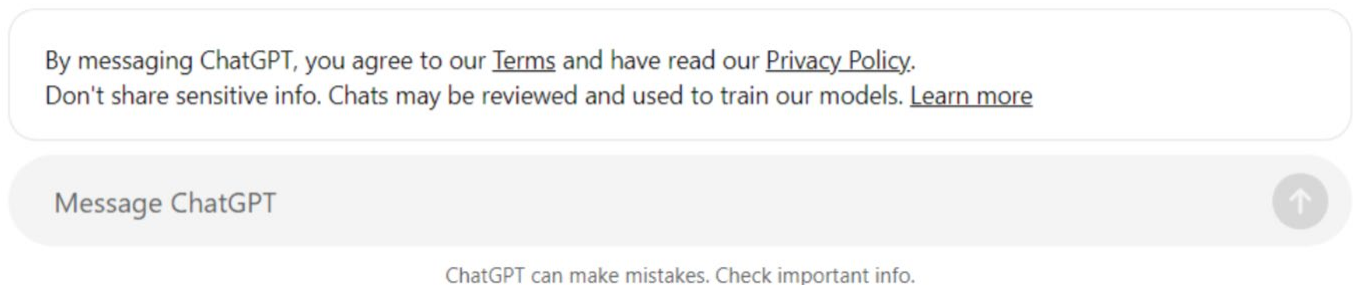


Figure 2. Example of notification about the sharing of sensitive information (web browser)

291. OPC's in-house testing of ChatGPT during the investigation confirmed that this notice is displayed **only once**, when a user first creates an account.¹⁵⁷ For the online free version of ChatGPT, which was made available to users without an account in April 2024, our testing revealed that the notice is also displayed only once, after the user has entered their first prompt and received a first response.

292. As explained in the previous section, the form of consent required will depend on the sensitivity of the information in question and the reasonable expectations of the individual.

Sensitivity

293. We find that OpenAI's notification, which is provided only one time, is not sufficient to ensure that sensitive data will not be included in the user interactions data that the company uses for model training. The notice is particularly inadequate where it is displayed after the user

¹⁵⁷ In April 2024, OpenAI made GPT-3.5 available to users without an account. Our testing showed that a contextual notice about sensitivity is displayed after the user has entered the first prompt and received a first response. The notice is provided once per online session.

has already inputted their first prompt (i.e., in the ‘free version without an account’), or where another user, who did not see the initial notification, uses the service.

294. As noted earlier, OpenAI explained that it has implemented risk mitigation measures to reduce the presence of personal information, including sensitive personal information, in its training datasets and model outputs. With respect to user interactions that may be included in fine-tuning data, OpenAI confirmed that for the fine-tuning of its GPT-3.5 and 4 it:
- i. used only a subset of user interactions for model training (fine-tuning);
 - ii. disassociated the interactions from the user account;
 - iii. used a third-party filtering tool to identify and remove certain specified personal identifiers from the interactions; and
 - iv. instructed the employees and contractors who review the subset of interactions that might be used for fine-tuning purposes to exclude any content that may include personal information.

295. We accept that the above measures may have limited the extent to which certain types of personal information were included in the datasets used to fine-tune GPT-3.5 and 4. Further, given that interactions were disassociated from users’ accounts, the risk of potential inferences being made about individual users was likely limited.

296. However, we find that at the relevant time these mitigation measures did not cover the broad spectrum of personal information captured under the Acts, and that some of the remaining personal information would have been, in some instances, sensitive.¹⁵⁸ While the use of the third-party filtering tool was intended to filter out certain specific personal identifiers, it did not remove all types of personal information, such as an individual’s opinions or characteristics, including where those may be about another identifiable individual – which can include or reveal sensitive personal information (e.g., ethnicity, sexual orientation, health conditions, etc.).

Reasonable expectations

297. We further find that OpenAI’s collection and use of users’ interactions with ChatGPT for the purpose of training its GPT-3.5 and 4 models was outside individuals’ reasonable expectations or what would have been obvious to a reasonable person.¹⁵⁹

298. OpenAI submitted that since ChatGPT was first released, it has consistently been transparent to users about how their data is used to improve its models. For example, at the time of account creation or when users open the mobile ChatGPT application for the first time, they are informed that their “Conversations may be reviewed by our AI trainers to improve our

¹⁵⁸ See [Gordon v. Canada \(Health\)](#), 2008 FC 258.

¹⁵⁹ Principle 4.3.5 of Schedule 1 of PIPEDA, section 8 of PIPA-BC, section 8(1)(b) of PIPA-AB.

[systems/services]” or “train our models.”¹⁶⁰ OpenAI further represented that it informs users of this practice in its Privacy Policy and Help Center articles. Finally, OpenAI stated that all users, other than non-logged-in users accessing ChatGPT via a web browser, must affirmatively click the “Okay, let’s go button (in the web browser version) or the “Continue” button (in the mobile app version) after being presented with the above notice.

299. We find that when ChatGPT was released, many users were likely unaware or lacked basic understanding of the implications of their personal information being used to train OpenAI’s models, including the potential review of their conversations by human trainers.
300. In that context, we do not believe that a one-time notification at account creation or during first use was sufficient to ensure that users would be informed of the nature, purposes and consequences of the processing. While we acknowledge that some of these users may have sought and found additional information about the practice in OpenAI’s Terms of Use, Privacy Policy and Help Center articles, not all users would have taken the time to read these privacy communications.
301. Consequently, we find that OpenAI should have obtained express consent for the collection and use of personal information included in users’ interactions with ChatGPT to train its GPT-3.5 and 4 models.

Choice

302. OpenAI’s use of data stemming from users’ interactions, for training purposes, is not integral to the provision of its services, even where it may be useful to OpenAI. Therefore, OpenAI must offer its users a choice about whether to participate or not in this practice.
303. OpenAI represented that its users can opt out of having their content used to improve model performance, either through their account settings (this option was made available in April 2023) or by submitting a privacy request to OpenAI.¹⁶¹
304. In response to our questions, the Respondent further submitted that as of October 2023 (almost a year after the launch of ChatGPT), only a negligible number of users in Canada had opted out of having their interactions data used for training purposes.
305. For the reasons detailed below, this number may not be reflective of the number of individuals who would have preferred that OpenAI not use their interaction data to train its models:

¹⁶⁰ The language can slightly vary depending on whether the notice is displayed on the web browser or mobile app, etc.

¹⁶¹ Initially, users could reach out to OpenAI support to opt out. In February 2023, OpenAI created an online form for users to opt-out of the use of their content to train and improve OpenAI’s models. In October 2023, OpenAI migrated the online form to a new Privacy Request Portal at privacy.openai.com.

- i. First, as mentioned above, we are of the view that individuals who did not reasonably expect this practice and were not otherwise aware of it would have been less likely to be aware of the availability of an opt-out control, thereby making it less accessible and effective.
- ii. Second, until April 2024, users who wished to opt out of having their interaction data used for training purposes in their account settings were also required to relinquish their chat history. This is an example of a “deceptive design pattern” of “forced action”, where the user was required, unnecessarily, to accept OpenAI’s use of the information to train its models in order to maintain access to their chat history.¹⁶² While users still had the option to reach out to OpenAI’s privacy team or use the online form, this may have dissuaded some users from opting out, rendering the option less accessible. Our testing showed that over the course of the investigation, OpenAI adjusted the account settings options available to its users, who now have the ability to opt out of having their interactions used for training purposes, without having to disable their chat history. OpenAI also introduced the option for users to engage in “temporary chats,” which are not used for model training.¹⁶³ These changes address our concern that setting design may have discouraged users from opting out of training for fear of losing their chat history.
- iii. Finally, we note that until recently the opt-out control was not consistently available to those users who used the free online version of ChatGPT (i.e., without an account) on their cell phones. Therefore, such users could not always opt out, in their settings, of having their interactions used for training, even if that was their preference. OpenAI represented that the option to opt out of training when using the free, signed-out version of ChatGPT launched concurrently with the ability to use ChatGPT while signed out on the web, the iOS app, and the Android app between April and June 2024. However, OpenAI further stated that there was a period of time where the option did not appear in signed-out settings on mobile web. OpenAI confirmed that this issue was fixed on November 4, 2025.

Findings related to GPT-3.5 and 4

306. Given all of the above, we find that OpenAI was required to obtain express consent for its collection and use of personal information included in users’ interactions with ChatGPT for the purpose of training its GPT-3.5 and 4 models. Indeed, this collection and use may have involved sensitive personal information and/or was not within users’ reasonable expectations.

¹⁶² For more information on deceptive design patterns, see for example [Dark Commercial Patterns, OECD Digital Economy Papers](#), October 2022, and [OPC’s Sweep Report on Deceptive Design Patterns](#), July 9, 2024.

¹⁶³ Temporary chats allow users to chat with ChatGPT without it appearing in history, affecting ChatGPT’s ‘memory’ (i.e., ChatGPT’s ability to personalize responses based on context from earlier conversations), or being used to train OpenAI’s models. See OpenAI’s [Temporary Chat FAQ](#).

307. Consequently, the OPC, OIPC-AB and OIPC-BC find that OpenAI did not obtain valid consent for this practice, such that OpenAI contravened section 6.1 as well as Principles 4.3 of Schedule 1 of PIPEDA, sections 7 and 8 of PIPA-AB, and sections 6-8 of PIPA-BC.

Recent developments and conclusion under PIPEDA

308. In response to our Preliminary Report, OpenAI informed the Offices that it has developed a novel tool that can detect and redact additional categories of personal identifiers from disassociated conversation data, which the third-party filtering tool (previously used) could not detect.¹⁶⁴ OpenAI provided the Offices with the results of recent internal evaluations showing the effectiveness of the tool at identifying private or sensitive information.

309. The OPC accepts that this new tool, combined with OpenAI's other mitigation measures implemented at the various stages of development and deployment of ChatGPT (detailed in various sections of this report and listed in Appendix A), can significantly reduce the risk that the personal information of private individuals, and sensitive information more specifically, may be included in the datasets used to fine-tune OpenAI's AI models and/or disclosed in model outputs. In making this determination, as further discussed in other sections of this report¹⁶⁵, the OPC also considered OpenAI's additional commitments with respect to transparency (including the publication and promotion of a Canadian blog post explaining its privacy practices) and its decision to deprecate GPT-3.5 and 4 and fully train the current models powering ChatGPT with the new mitigation measures.

310. We also acknowledge that individuals' use of generative AI has increased significantly since the launch of our investigation. Users of these services are now more likely to have gained a basic understanding of AI model training. Given that context, together with the implementation of OpenAI's existing and new mitigation measures that we expect will materially reduce the risk of privacy harms, we accept that users may reasonably expect that OpenAI's future models will be fine-tuned using a subset of user interactions.

311. Consequently, the OPC accepts that OpenAI may rely on implied consent in this context and finds this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

¹⁶⁴ This tool is the same one that OpenAI is using to detect personal information in the pre-training datasets, although OpenAI started to use it to filter user interactions first.

¹⁶⁵ In particular, see section relating to OpenAI's response to our Preliminary Report.

312. This conclusion is based on OpenAI's representations and our understanding that OpenAI will continue to effectively implement and improve these mitigation measures and develop further innovative privacy-protecting techniques in the future.

Conclusion under PIPA-BC and PIPA-AB

313. The OIPC-BC and OIPC-AB are encouraged by the measures OpenAI has taken but are not making findings about these recent developments, as the fundamental issue (described in Issue 2A analysis) remains with respect to publicly accessible information, which is the primary source of training data for the models.

Analysis under Quebec's Private Sector Act

314. As discussed above in paragraphs 231 to 237 under the heading "Rules relating to the duty to inform that apply in the context of using personal information for primary purposes", the rules relating to consent under Quebec's Private Sector Act differ from those under PIPEDA, PIPA-BC and PIPA-AB. When personal information is collected directly from the individual concerned, as in the case of collecting user chats for the purpose of training GPT-3.5 and 4 models, the enterprise is subject to a duty to inform.

315. This duty to inform means that the company must inform the person, in plain and clear language, of the purposes for which the information is being collected, the means by which it is being collected, the rights of correction provided for in the Act, and the right to withdraw consent to the communication or use of the information collected.¹⁶⁶

316. In the preliminary report, the CAI took the position that the notice given to users when creating an account or for the free online version, which was given after users had first entered a prompt (see screenshot in paragraph 290), was insufficient to inform users that the content of their chat with the GPT-3.5 and 4 models could be collected for the purpose of training these models.

317. Following OpenAI's submissions, the CAI notes that an additional notification window alerting users not to share sensitive information and informing them that their chat history could be reviewed or used to improve services was displayed to users when creating an account on the system's web interface or when opening the application for the first time after downloading it, and again after creating an account on the application.

318. In addition, the CAI notes that this notification window included a link to help center articles that informed users, in particular, that their chats were used by default to train the model and that the content of these chats could be reviewed by authorized personnel for the purposes specified in these articles.

¹⁶⁶ Section 8 of Quebec's Private Sector Act.

319. Considering this additional information provided to users and considering that OpenAI's terms and conditions of use and privacy policy informed users in clear and simple language that their interactions with the GPT-3.5 and 4 models could be reviewed for the purpose of improving services and used for training these models, the CAI agrees to review its preliminary position regarding the information provided to users and the quality of consent to the use of their chat for model training purposes in order to conclude that the information measures put in place for users with an account or who have downloaded the application, were sufficient and that, as a result, the consent obtained from these users was in compliance with Quebec's Private Sector Act.
320. However, the CAI's opinion remains that the information provided to users of the free online version of ChatGPT was not compliant, because the notice regarding the use of chat data for model training was only given to users after an initial collection of information, which could include personal information.¹⁶⁷
321. The CAI considers that the footnote stating that by sending a message to ChatGPT you accept the terms of use and acknowledge that you have read the privacy policy and the hyperlinks referring to these documents were not sufficient in this case to inform individuals, of the purposes for which their personal information was being collected, in a timely manner and in accordance with the law.
322. To this end, the CAI concludes that OpenAI's practices relating to the free web version of GPT-3.5 and 4 concerning the collection of user chats for model training purposes did not comply with section 8 of Quebec's Private Sector Act.
323. It should be noted that, as of March 24, 2026, OpenAI informed the CAI that it was committed to updating the free, account-free online version of ChatGPT so that the notice informing users that their chat logs are collected for model training purposes and the notice advising against sharing sensitive information appear before the first chat with the system.
324. Ultimately, and as will be detailed below, despite the information provided to users regarding the collection and use of their chats by default for training models, section 9.1 of Quebec's Private Sector Act provides for a specific obligation regarding privacy settings, and the CAI considers that OpenAI's practices in this case contravened this section.

Issue of the highest privacy settings by default under section 9.1 of Quebec's Private Sector Act

325. Section 9.1 of Quebec's Private Sector Act stipulates that a company that collects personal information by offering the public a technological service with privacy settings must ensure that those settings provide the highest level of confidentiality by default, without any intervention by the person concerned.

¹⁶⁷ *Guidelines 2023-1- Consent: validity criteria*, Version 1.0, October 31, 2023, title B.9. Time of consent.

326. According to OpenAI's submissions, when ChatGPT was made available to the public in November 2022, users could contact OpenAI's support service to opt out of having their conversations collected for model training purposes.
327. In February 2023, an online form was made available to allow users to opt out of this collection.
328. Then, in April 2023, OpenAI introduced a feature within the data control settings interface to allow users to disable their chat history so that their new conversations could not be used for model training purposes.
329. In October 2023, OpenAI migrated its online form to its new portal, privacy.openai.com.
330. In April 2024, OpenAI modified the data control interface to allow users to opt out of having their conversations collected for model training purposes without losing their conversation histories.
331. During the same period, OpenAI introduced a temporary conversation mode in which user interactions are not collected for model training purposes. Finally, OpenAI allowed users without accounts to opt out of having their keystrokes collected for model training purposes.
332. Ultimately, regardless of the period in question, the system's privacy settings ensured that the option to collect user chats for the purpose of refining model training was enabled by default and that users had to take positive action to opt out of this collection.
333. However, section 9.1 of the Quebec's Private Sector Act, which came into force on September 22, 2023, provides that the highest privacy settings must be enabled without any action on the part of the individual concerned.
334. OpenAI alleges in its comments that section 9.1 of Quebec's Private Sector Act is not applicable in this case.
335. It first argues that the use of user interaction data for training its models is a so-called "secondary" purpose and that this use may benefit from the exceptions to the consent rules provided for in paragraphs 1 and 5 of section 12 of the Quebec's Private Sector Act, namely the exception for purposes compatible with those for which the collection was made or the exception for use necessary for study, research, or statistical purposes, when such information is de-personalized.
336. OpenAI then argues that its users take clear and positive steps to activate the collection settings for their conversations by clicking "continue" during the onboarding process and accepting ChatGPT's privacy policy and terms of use.

337. OpenAI argues that the terms “technology services,” “privacy settings,” and “highest default privacy level” create real ambiguity, such that one must refer to the underlying principles of the Charter to reconcile the objectives of the Act with the protections afforded to other fundamental rights, including freedom of expression.
338. Referring to the regulatory impact analysis produced on July 30, 2020 by the Secrétariat à la réforme des institutions démocratiques, à l'accès à l'information et à la laïcité (secretariat for democratic institutions, access to information and secularism)¹⁶⁸ as part of the modernization of Quebec's Private Sector Act, OpenAI argues that the concept of privacy settings refers more to settings related to the disclosure of information. It notes that this interpretation aligns with the ordinary meaning of the term “confidentiality”, which pertains to restricting access and communication to unauthorized persons, as opposed to the term “privacy”, which in its broad sense could include how information is used within the organization.¹⁶⁹
339. To the end, OpenAI argues that the underlying purpose of section 9.1 is not to regulate the use of information within the company and that the function “improving the model” is not a confidentiality parameter, since this parameter concerns the use of chat logs within the company itself and not the act of communicating or making this information available to third parties.
340. As will be described below, the CAI does not agree with the arguments raised by OpenAI and considers that section 9.1 of Quebec's Private Sector Act applies in this case, for the following reasons.
341. This section aims to provide users of technological services or products with broader protection so that their personal information is protected by default through the systems' initial settings.
342. This protection, which falls under the principle of “privacy by design”,¹⁷⁰ is intended to prevent users from having to configure the systems themselves to ensure the protection of their privacy.
343. The CAI considers that this protection applies to the entire lifecycle of personal information, including its collection, use, disclosure, and retention.

¹⁶⁸ *Analyse réglementaire : Projet de loi modernisant des dispositions législatives en matière de protection des renseignements personnels*, Secrétariat for democratic institutions, access to information and secularism, July 30, 2020.

¹⁶⁹ Letter of March 2, 2026.

¹⁷⁰ Ann Cavoukian, [Privacy by Design: The 7 Foundational Principles](#), Information and Privacy Commissioner of Ontario, 2011.

344. A reading of this provision in light of the context of the Act and evidence regarding Parliament’s intent reveals that parameters relating to the training of artificial intelligence models fall within the scope of this section.
345. To be subject to section 9.1 of Quebec’s Private Sector Act, the company must collect personal information via:
- i. A technological product or service;
 - ii. This product or service must be offered to the public; and
 - iii. It must include privacy settings.
346. Although not defined in Quebec’s Private Sector Act, the terms “technological product or service” have been defined by the legislator in section 3 of the *Act Respecting Health and Social Services Information*,¹⁷¹ and beyond the specific health-related context of that Act, this definition can easily be applied to the context of section:
- “technological product or service” means equipment, an application or a service required to collect, keep, use or communicate information, such as a database or an information system, a telecommunications system, technological infrastructure, software or a computer component of medical equipment
347. To this end, any application, including any web service, that collects, stores, uses or communicates personal information is a technological service within the meaning of the Act.¹⁷²
348. That being the case, there is no doubt that OpenAI, through its ChatGPT 3.5 and 4 services, provided technological services to the public.
349. As shown in the image below, the interfaces of these models included privacy settings regarding the use of users’ chat logs for the purpose of training these models.

¹⁷¹ [Act Respecting Health and Social Services Information, RLRQ, c. R-22.1.](#)

¹⁷²It should be noted that the minister responsible at the time of the adoption of section 63.7 of the *Act respecting access to documents held by public bodies and the protection of personal information*, RLRQ c A-2.1, which is the public counterpart of section 9.1 of Quebec’s Private Sector Act, used these terms to refer to the concept of a technological service [translation] “. . . if you use a system. . . or a website, or an application of any kind. . .” [Journal des débats de la Commission des institutions - Assemblée nationale du Québec](#), Journal des débats de la Commission des institutions, 42nd Legislature, 1st session, Wednesday March 10, 2021 – Vol. 45, No. 123, 12:04 p.m.

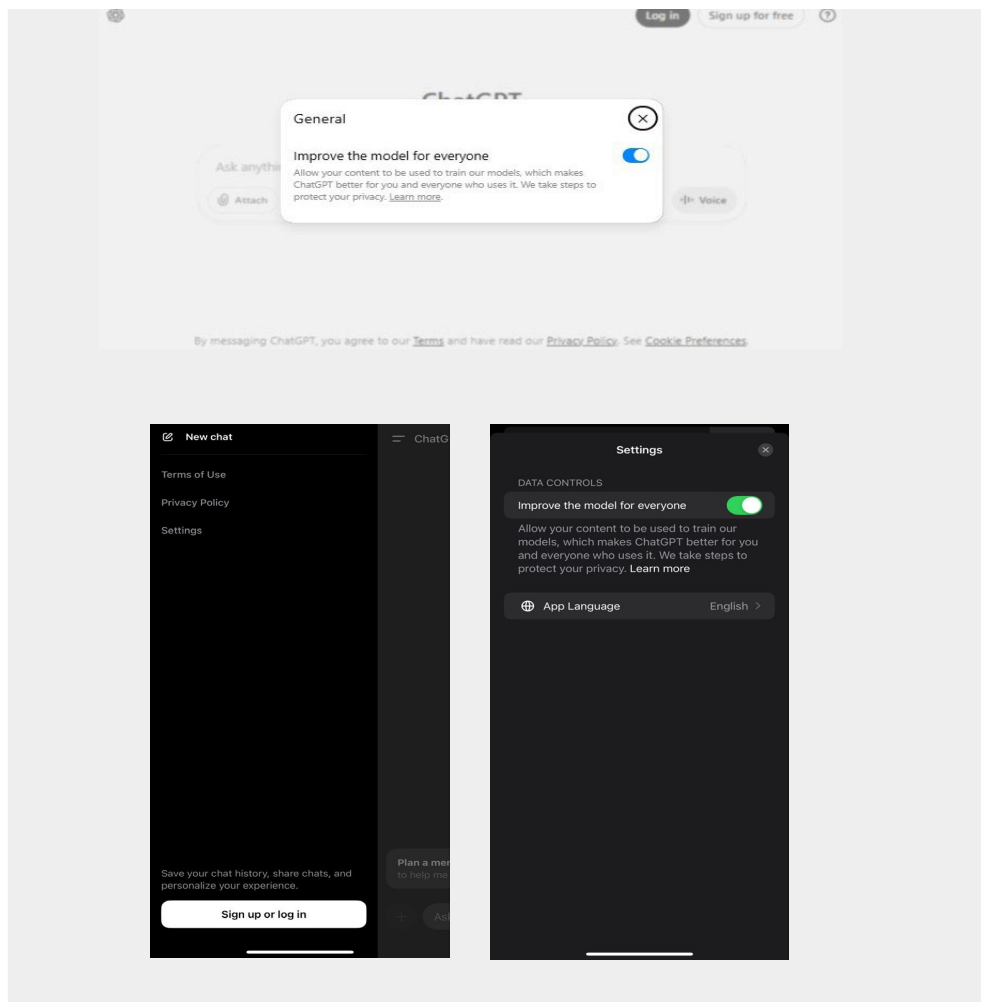


Figure 3. Opt-out mechanism for logged out users

350. To that end, after September 22, 2023, section 9.1 of Quebec's Private Sector Act would apply to these settings.
351. These settings were therefore to be configured in such a way that the highest level of confidentiality is ensured.
352. The ChatGPT 3.5 and 4 systems had two levels of settings, and depending on the selected level, the collection of chat logs was either permitted or not permitted.
353. In this context, the highest level of confidentiality was the level at which the collection of chat logs was not permitted.
354. In fact, the choice that was offered by OpenAI to ChatGPT users has the opposite effect of what is provided for in section 9.1, namely that users must take positive action to disable the collection of their chats for model training purposes.

355. With regard to OpenAI's argument that the use of user chat logs to train its models constitutes a so-called "secondary" use and that such use may qualify for the exceptions to the consent requirements set out in paragraphs 1 and 5 of section 12 of Quebec's Private Sector Act, the Commission considers that section 9.1 applies at the stage of collecting personal information and that this stage is distinct from the stage at which the so-called "secondary" uses referred to in section 12 of Quebec's Private Sector Act apply.
356. Consequently, the application of the exceptions regarding consent set forth in section 12 of Quebec's Private Sector Act cannot affect the application of section 9.1 of Quebec's Private Sector Act.
357. Furthermore, the CAI considers that OpenAI's interpretation of the application of section 12 of Quebec's Private Sector Act would run counter to the very intention of the legislature to take a more protective approach in the context of technological products or services.
358. This interpretation would mean that the use of personal information for a so-called "secondary" purpose under section 12 of Quebec's Private Sector Act could be subject to less restrictive confidentiality conditions than would be the case for use for a so-called "primary" purpose.
359. Regarding the second argument raised, the CAI considers that interpreting the acceptance of privacy and technology service use policies as a positive action by users in accordance with section 9.1 runs counter to the additional protection desired by the legislature when it comes to technology products or services. This interpretation, as OpenAI put forward, would land the user with the choice of whether or not to use the service or product, rather than with the choice desired by the legislature, namely whether or not to allow more intrusive privacy settings.
360. Furthermore, the CAI does not share the position put forward by OpenAI that the terms used in section 9.1 of Quebec's Private Sector Act create a real ambiguity¹⁷³ and finds that this ambiguity stems more from the interpretation proposed by OpenAI.
361. When addressing the issue of "privacy settings," the common understanding of this term is not limited solely to controlling access to and disclosure of information, but normally includes control over the collection, use, and purposes for which the information is collected.
362. When addressing the concept of the highest level of privacy, a user of a technological service expects not only that their information will not be made accessible to or shared with unauthorized third parties, but also that their information will be used solely for the purposes

¹⁷³ Google, LLC v. Canada (Privacy Commissioner), [2023 FCA 200](#), para. 76, *R. v. Jarvis*, [2019 SCC 10](#), at para. [105](#); *Bell ExpressVu Limited Partnership v. Rex*, [2002 SCC 42](#), at paras. [29 and 62](#).

for which it was collected and that such use will be restricted even within the parameters of the system they are using.

363. And even if the literal interpretation of the term “privacy setting” were limited to the concept of access to and disclosure of information, the legislative context in which this term is used and Parliament’s intent regarding the purpose of section 9.1 indicate that the scope of this section extends beyond these two aspects alone.
364. Contextually, one must take into account the fact that this provision is part of legislation aimed at protecting the public and must therefore be interpreted broadly and liberally.¹⁷⁴
365. Minister Lebel’s memorandum to the Cabinet when Bill 64 was tabled, which incorporates the Secrétariat à la réforme des institutions démocratiques, à l’accès à l’information et à la laïcité’s regulatory impact analysis referenced by OpenAI in its comments, sheds light on the terms used in section 9.1 and intentions of the legislature:

[translation]

4.2.5. Privacy by Default

The principle of “privacy by default” or “protection by default” implies that personal data is automatically protected without any additional action being required on the part of an individual.

In other words, these measures ensure that, by default, personal information is not made accessible to an indeterminate number of natural or legal persons without the intervention of the individual concerned.

Thus, when a product or service (app, social media, connected device) offers privacy settings (primarily regarding data sharing), these settings must initially ensure the highest level of privacy. Subsequently, any modification of these settings must require the intervention of the individual concerned. To enshrine this principle in the Private Sector Act, it is proposed to add a provision requiring that the settings of the products or services offered ensure the highest level of privacy without any intervention by the individual concerned¹⁷⁵.

¹⁷⁴ *Conseil de presse du Québec c. Lamoureux-Gaboury*, 2003 CanLII 33002 (QC CQ), paras. 41–53, *Québec (Commission des droits de la personne et des droits de la jeunesse) v. Montréal (City)*; *Québec (Commission des droits de la personne et des droits de la jeunesse) v. Boisbriand (City)*, 2000 SCC 27, paras. 28–30, and *New Brunswick (Human Rights Commission) v. Potash Corporation of Saskatchewan Inc.*, [2008] 2 SCR 604, 2008 SCC 45 (CanLII), paras. 19, 65–67, and section 41 of the *Interpretation Act*, RLRQ c. I-16.

¹⁷⁵ Sonia LEBEL, [Brief to the Cabinet—Act to modernize legislative provisions regarding the protection of personal information \(publicly accessible section\)](#), Government of Quebec, 2020, sect. 4.1.6 (in French only).

366. Section 9.1, as highlighted by Minister Lebel in her memorandum and by the Secrétariat à la réforme des institutions démocratiques, à l'accès à l'information et à la laïcité in its regulatory impact analysis,¹⁷⁶ is grounded in the principle of protection by default, which is a corollary of the principle of privacy by design.

367. The concept of privacy by default is defined as follows by Dr. Ann Cavoukian:

“We can all be certain of one thing — the default rules! Privacy by Design seeks to deliver the maximum degree of privacy by ensuring that personal data are automatically protected **in any given IT system or business practice**. If an individual does nothing, their privacy still remains intact. No action is required on the part of the individual to protect their privacy — it is built into the system, by default.”¹⁷⁷

[Emphasis added.]

368. However, contrary to what OpenAI argues, the protection sought by this privacy by default philosophy goes far beyond simply controlling access to and the disclosure of information.

369. This concept of privacy by default aims to protect privacy in its broadest sense, which includes the collection, use, communication, and retention of personal information. The reference made by Minister Lebel and the Secrétariat à la réforme des institutions démocratiques, à l'accès à l'information et à la laïcité to this concept contradicts the narrow scope supported by OpenAI's proposed interpretation.

370. Similarly, it is important to note the use of the term “primarily” in the Minister's memorandum and the regulatory impact analysis when referring to privacy settings and citing data disclosure as an example. The use of this term indicates that the scope of section 9.1 of Quebec's Private Sector Act is not limited solely to settings that control access to and communication of information.

[translation]

Thus, when a product or service (application, social media, connected device) offers privacy settings (**primarily** data disclosure), these settings must initially ensure the highest level of privacy.¹⁷⁸

¹⁷⁶ Minister Lebel's memorandum, cited in footnote 175, p. 11-12 and regulatory impact analysis cited in footnote 168, p. 10.

¹⁷⁷ [Ann CAVOUKIAN, *Privacy by Design: The 7 Foundational Principles*](#), Ontario Information and Privacy Commissioner, 2011.

¹⁷⁸ Cited above at footnote 175.

[Emphasis added.]

371. It should also be noted that the French term “paramètre de confidentialité” is translated in the English version of this article as “privacy setting,” which also refers to the concept of privacy protection in its broadest sense.

372. The CAI therefore concludes that section 9.1 of Quebec’s Private Sector Act reflects the legislature’s intent to uphold, at the time of collecting personal information, the concept of privacy by design, and that the protection afforded by this section applies to the entire lifecycle of such information and not merely to the aspects concerning access to and disclosure of such information.

373. It should be noted that this interpretation is also shared by the authors of the 3rd annotated edition of the Act respecting the protection of personal information in the private sector:

[translation]

Section 9.1 imposes an obligation on businesses to provide the highest standards regarding the confidentiality of information. This implies that, by default, technology—such as an application—**must collect or share the minimum amount of personal information possible**. Consequently, profiling or location tracking features must be disabled by default when the technologies implementing them fall within the scope of section 9.1.¹⁷⁹

[Emphasis added.]

374. Finally, contrary to OpenAI’s claims, the use of ChatGPT 3.5 and 4 user chat logs was not limited to so-called “internal” use, as OpenAI characterizes it in its arguments. In fact, in certain cases, training the model could lead to the disclosure of certain information to other users and correspond to the more restrictive definition of the term “privacy parameter” proposed by OpenAI.

375. For these reasons, the CAI concluded that OpenAI’s practices after September 23, 2023, did not comply with section 9.1 of Quebec’s Private Sector Act, because the privacy settings of the ChatGPT 3.5 and 4 systems did not provide the highest level of privacy.

¹⁷⁹ 3^e édition annotée de la Loi sur la protection des renseignements personnels dans le secteur privé, les Éditions Yvon Blais, ISBN : 978-2-89730-759-2, P. 118.

Issue 2C: Did OpenAI obtain valid consent and meet its obligation to inform individuals with respect to the disclosure of personal information collected from publicly accessible websites, licensed third-party sources and users' interactions with ChatGPT?

Analysis under PIPEDA, PIPA-BC and PIPA-AB

376. As previously noted, information that is publicly accessible is not “publicly available” as defined under the Acts. For the reasons set out below, OpenAI should have obtained express consent for the disclosure, via or in relation to ChatGPT, of personal information collected from publicly accessible websites, licensed third-party sources and/or users' interactions with ChatGPT, where the information was sensitive and/or the practice was outside the reasonable expectations of the individual.

377. As explained above, at the time of training its GPT-3.5 and 4 models, OpenAI's training datasets may have included significant amounts of personal information of varying levels of sensitivity.

378. OpenAI represented that it implemented various measures to reduce the risk that its GPT-3.5 and 4 models disclose “sensitive or private information” in their outputs. In addition to the measures already discussed at paragraphs 128 and 294 (i.e., removal of certain categories of websites from the training datasets, disassociation and filtration of user interactions, instructions provided to human trainers), OpenAI further represented that:

- i. It trained these models not to return “private or sensitive information” about people (such as personal addresses or other personal identifiers) in response to user requests – even if the information is accessible publicly on the Internet and via search engines, and even if the individual is publicly known – thus limiting the personal information disclosed.¹⁸⁰ OpenAI further represented, albeit without corroboration, that according to an internal evaluation, GPT-3.5 and GPT-4 were respectively 98% and 100% effective in blocking such information in model outputs.
- ii. The information that ChatGPT returns in its responses is generally limited to information about public figures or other individuals who have a significant presence online, and personal information about private individuals is unlikely to appear in a response, given their low statistical occurrence in the training datasets. In addition, Open AI stated that an Internet search for information about a private individual would generally be more likely to yield personal information than a question to ChatGPT about a private individual.

379. We will not comment on the amount and nature of personal information included in search engine responses, as these are returned in a different context and are not the subject of this investigation. While we recognize that this may have already evolved, we do note, however,

¹⁸⁰ OpenAI's models learn to generalize from their training, such that they are able to refuse or provide safe responses to a broader set of requests than the ones they were trained on.

that until recently, ChatGPT fundamentally differed from search engines which simply enabled users to navigate and find information where it is located on the Internet. ChatGPT, on the other hand, can create new content, some of which may contain plausible but inaccurate or fabricated information, including personal information (i.e., “hallucinations,” as discussed later in this report).

380. With regard to the mitigations implemented by OpenAI in the context of GPT-3.5 and 4 (noting that this has since evolved), our investigation revealed that the categories of “sensitive or private information” that OpenAI sought to remove from model outputs were more limited than the broad range of personal information covered by the Acts. For example, OpenAI indicated that it had no metrics to determine whether certain types of personal information, such as opinions or rumours about individuals, would be disclosed.

381. Where OpenAI discloses personal information about individuals that are not public figures, we find that this disclosure would be outside individuals’ reasonable expectations for the same reasons as those discussed in subsections 2A and 2B. Furthermore, depending on the nature of the personal information disclosed and the specific circumstances of the case, this may also be outside of the reasonable expectations of public figures, even where that personal information is accessible on the Internet.

Findings related to GPT-3.5 and 4

382. Given the above, we find that at the time of developing and deploying GPT-3.5 and 4, OpenAI should have obtained express consent for its disclosure of personal information via ChatGPT, where that information was sensitive and/or the disclosure outside individuals’ reasonable expectations.

383. Consequently, the OPC, OIPC-AB and OIPC-BC find that OpenAI did not obtain valid consent for this practice, such that OpenAI contravened section 6.1 as well as Principles 4.3 of Schedule 1 of PIPEDA, sections 7 and 8 of PIPA-AB, and sections 6-8 of PIPA-BC.

Recent developments and conclusion under PIPEDA

384. As noted earlier, in response to our Preliminary Report, OpenAI represented that it has implemented a new tool to detect and mask identifying information about individuals which may be included in training data collected from publicly accessible information, licensed datasets or user interactions.

385. OpenAI stated that this tool can more accurately mask personal information about private individuals, as well as determine when to mask private information about public figures and when to maintain information about such public figures which may be of interest to the public. OpenAI further asserted that to the extent that a broader range of personal information (such as an individual's opinions or characteristics) are included in the training datasets, the tool can detect and redact identifiers which would link such information to an identifiable individual.

386. The OPC accepts that this new tool, combined with OpenAI's other mitigation measures implemented at the various stages of development and deployment of ChatGPT (detailed in various sections of this report and listed in Appendix A), may significantly reduce the risk that the personal information of private individuals, and sensitive information more specifically, will be included in the datasets used to train OpenAI's future models and potentially disclosed in model outputs moving forward. In making this determination, as further discussed in other sections of this report¹⁸¹, the OPC also considered OpenAI's additional commitments with respect to transparency (including the publication and promotion of a Canadian blog post explaining its privacy practices) and its decision to deprecate GPT-3.5 and 4 and fully train the current models powering ChatGPT with the new mitigation measures.

387. Consequently, the OPC finds this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

388. This conclusion is based on OpenAI's representations and our understanding and expectation that OpenAI will continue to effectively implement and improve these mitigation measures and develop further innovative privacy-protecting techniques in future.

Conclusion under PIPA-BC and PIPA-AB

389. Under the Issue 2A analysis above, the OIPC-BC and OIPC-AB determine that OpenAI has not established compliance with the specific requirements for implicit consent under s. 8 of PIPA-BC and deemed or notice consent under s. 8(2) of PIPA-AB when OpenAI collects and uses personal information from publicly accessible websites and licensed third-party sources.

390. The requirements for implicit consent under s. 8 of PIPA-BC and deemed or notice consent under s. 8(2) of PIPA-AB extend to an organization's disclosure of personal information. Therefore, the analysis and conclusions in the Issue 2A discussion above are also applicable to OpenAI's disclosure of personal information collected from publicly accessible websites, licensed third-party sources, and users' interactions with ChatGPT, to the extent that OpenAI does not obtain express consent for that disclosure.

391. Consequently, the OIPC-BC and OIPC-AB find this aspect of the complaint to be **well-founded and unresolved** under PIPA-BC and PIP-AB.

392. Regarding future and unexamined ChatGPT models, the OIPC-BC and OIPC-AB decline to draw a conclusion about future disclosures of personal information and whether such disclosures meet the requirements of implicit consent under PIPA-BC and deemed consent under PIPA-AB. The OIPC-BC directs OpenAI to the specific requirements of implicit consent under s. 8 of PIPA-BC, which apply to any disclosure of personal information on the basis of implicit consent. Similarly, the OIPC-AB directs OpenAI to the specific requirements of deemed

¹⁸¹ In particular, see section relating to OpenAI's response to our Preliminary Report.

or notice consent under s. 8(2) of PIPA-AB, which apply to any disclosure of personal information on the basis of deemed or notice consent.

Analysis under Quebec's Private Sector Act

393. Regarding the question of the communication of personal information that may be carried out by ChatGPT, under Quebec's Private Sector Act, the rules governing consent discussed in detail in issues 2A and 2B apply.
394. Regarding the communication of information from publicly accessible websites and third parties holding data licensed by the system, the CAI reiterates that OpenAI has not been able to demonstrate that it has sufficiently documented the context in which consent was obtained to ensure its validity.
395. To that end, the CAI considers that the reasons, recommendation, and reservation of rights set out in issue 2A apply to the issue of OpenAI's communication of personal information.
396. Regarding the communication of personal information collected for model training purposes and originating from chats used for model training purposes, the CAI reiterates its reasons as well as the recommendation and reservation of rights made in issue 2B.
397. Consequently, and for the reasons set out in issues 2A and 2B, the CAI concludes that OpenAI's practices regarding the disclosure of personal information contravened sections 6, 8, and 12, paragraph 1, 13 and 14 of Quebec's Private Sector Act.

Issue 3: Was OpenAI sufficiently open about its models?

398. We find that OpenAI did not meet the openness and transparency requirements under the Acts.
399. The Acts provide that an organization must be open and transparent about the collection, use and disclosure of individuals' personal information.¹⁸²
400. In the context of generative AI, specific expectations arising from statutory requirements were communicated in the Generative AI Principles published by the federal, provincial and territorial privacy authorities in December 2023.¹⁸³ In particular, these Principles state (not

¹⁸² This is per Principle 4.8 of Schedule 1 of PIPEDA, and section 13 of PIPA-AB. Section 10 of PIPA-BC requires transparency of the purposes for the collection of the personal information. In the case of Quebec's Private Sector Act, the concept of openness differs and refers to the obligation to provide information as set out in section 8 of the Act, which was discussed under Issue 2B, as well as the obligation to implement policies and practices regulating the governance of personal information set out in section 3.2 of the Act and the obligation to disseminate a privacy policy when personal information is collected by technological means, section 8.2 of the Act.

¹⁸³ See paragraph 61.

exhaustively) that developers and providers of generative AI should, in relation to openness, transparency and model explainability, do the following:

- i. inform individuals what personal information is collected, as well as how, when, and why it is collected, used or disclosed throughout any stage of the generative AI system's lifecycle (including development, training and operation) for which the party is responsible;
- ii. maintain and publish documentation about the datasets used to develop or train the generative AI tool, including the sources of the datasets, the legal authority for their collection and use, whether there are any licensing agreements or other restrictions on the acceptable uses of the datasets, and any modification, filtering or other curation practices applied to the datasets;
- iii. take appropriate steps to make the outputs from generative AI systems traceable and explainable. Where a developer or provider is of the opinion that outputs from a generative AI tool are not explainable, this should be made explicit to any organization using or individual interacting with the tool; and
- iv. ensure that all information communicated about a generative AI system is designed to be understandable by the intended audience, and made readily available both before, during and after use of the system.

401. OpenAI represented that it makes reasonable efforts to be transparent about its information handling practices. In addition to its [Privacy Policy](#) and [Terms of Use](#), which users must respectively acknowledge and accept before using ChatGPT on its website or app, OpenAI explained that it provides contextual notices and prompts at relevant points in the user journey, such as during registration and onboarding. It also maintains a [Help Center](#), where individuals can find responses to commonly asked questions about ChatGPT, and a [Research index](#), where research documents of a more technical nature are available.

402. We acknowledge that OpenAI's Privacy Policy and Help Center articles are readily accessible and generally written in plain language. We note, however, that the Privacy Policy, and some important documents, such as Help Desk articles (including the one discussed in the paragraph below), were not originally made available to individuals in Canada in French. OpenAI confirmed that it only made a French version of the Privacy Policy available in May 2024.

403. Furthermore, we find that certain key information is either incomplete, unclear or missing from OpenAI's communications. In particular, OpenAI provides only very high-level information about the datasets used to train its models. Indeed, its Privacy Policy includes a link to a Help Desk article entitled, "[How ChatGPT and our language models are developed](#)"¹⁸⁴,

¹⁸⁴ OpenAI included a reference to this article in its Privacy Policy in June 2023, that is, several months after the launch of ChatGPT. Prior to this date, the privacy policy referred exclusively to OpenAI's collection, use and disclosure of users' personal information for the purpose of providing the service, and provided no information about the collection, use and disclosure of users' and non-users' personal information for model training purposes.

which explains in generic and vague terms that ChatGPT was “developed (1) using information that is publicly available information on the Internet, (2) information that we (OpenAI) license from third parties, and (3) information that our (ChatGPT’s) users or human trainers provide”.

404. We find that OpenAI does not sufficiently explain the categories and sources of the personal information included in its training datasets. For example, an individual reading the above-referenced article would not necessarily understand that information posted by them or about them, sometimes many years ago, on a blog, discussion forum or social media, could be considered publicly accessible and potentially collected and used for the purpose of training OpenAI’s models.

405. OpenAI is no more transparent in this regard in the technical documents published on its website, which are intended for a more tech-savvy audience. For example, in its [GPT-4 Technical Report](#), OpenAI specifically states that “given both the competitive landscape and the safety implications of large-scale models like GPT-4, **this report contains no further details about the architecture (including model size), hardware, training compute, dataset construction, training method, or similar**” (emphasis added).¹⁸⁵

406. Finally, we note that Sam Altman, CEO of OpenAI, publicly stated that ChatGPT is by nature a “black box,” suggesting that the precise reasons why the LLM behaves the way it does, as well as the mechanisms that underpin its behaviour, are not known, even to its creators.¹⁸⁶ Although OpenAI has commented tangentially on this lack of explainability in some research documents published on its website, we find that it does not sufficiently make this explicit to organizations using the tool or individuals interacting with it. However, given the interest in and importance of this rapidly emerging and evolving technology, we would expect that strategies and techniques will emerge to improve the explainability of ChatGPT and other LLMs.

407. While we acknowledge that OpenAI has developed and made available numerous communication materials to explain its privacy practices and how its models work, including its Privacy Policy, relevant Help Center articles and contextual notices, some key information, as explained above, is not provided.

Findings related to GPT-3.5 and 4

408. Consequently, the OPC, OIPC-BC and OIPC-AB find that OpenAI contravened Principle 4.8 of Schedule 1 of PIPEDA, section 10 of PIPA-BC, and section 13 of PIPA-AB. In addition, the CAI

¹⁸⁵ The requirement regarding openness set out in Quebec’s Private Sector Act does not apply to technical documentation.

¹⁸⁶ [Sam Altman Says OpenAI Doesn’t Fully Understand How GPT Works Despite Rapid Progress](#), Observer, May 30, 2024.

strongly encourages OpenAI to implement the recommendations related to openness that are made in this report in relation to its models.

Recent developments and conclusion under PIPEDA

409. In response to our Preliminary Report, OpenAI indicated that it would make relevant Help Center content available in French. We see this as a positive development.
410. With respect to the disclosure of data sources, OpenAI initially submitted that the language that it uses in its communications is aligned with market practices for frontier model training. It further stated that interpreting the Acts as requiring developers of Generative AI models to specify the sources at such a level of granularity would be inconsistent with the statutory objectives of the Acts, which seek to balance the protection of individuals' privacy with the legitimate needs of organizations to conduct research and development. In its view, this would place an undue burden on organizations whereas, by contrast, the incremental benefits for individuals in terms of their ability to exercise privacy rights would be limited.
411. As mentioned above, while the OPC accepts that individuals' reasonable expectations regarding Generative AI have evolved over the past years, these expectations are closely tied to, and depend on, the level of transparency of developers and providers of Generative AI models.
412. While we agree that a full disclosure of every single data source would not be practical or necessary to comply with the Acts, we find that OpenAI's high-level description of the categories of personal information it collects does not meet the level of transparency required by the Acts. We further find that to allow individuals to understand its practices relating to the management of personal information, OpenAI should publish a comprehensive and sufficiently detailed overview of the main categories of content it uses to pretrain and fine-tune its models in a form that is generally understandable.¹⁸⁷
413. Following further discussions with the Offices, OpenAI committed to expanding its "[How ChatGPT and our foundation models are developed](#)" article to include more plain-language explanation about the sources of information used to train its models (including that OpenAI collects and uses publicly accessible content such as blogs or other public posts). It also agreed to adding similar language to the Canadian blog post it will publish on its website and promote in the Canadian media.¹⁸⁸
414. Consequently, we find this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

¹⁸⁷ This approach appears to align with that of other jurisdictions including the EU, where the AI Act (Recital 107) provides that this summary should be generally comprehensive in its scope instead of technically detailed. See Artificial Intelligence Act (Regulation (EU) 2024/1689), Official version of 13 June 2024 – [Recital 107](#).

¹⁸⁸ In particular, see section relating to OpenAI's response to our Preliminary Report.

Issue 4: Did OpenAI take reasonable steps to ensure that the information it generates about individuals is as accurate, complete, and up-to-date as is necessary for the purposes for which it is to be used?

415. For the reasons explained below, the Offices find that OpenAI did not meet the accuracy requirements under the Acts.

416. Principle 4.6 of Schedule 1 of PIPEDA states that “personal information shall be as accurate, complete, and up-to-date as is necessary for the purposes for which it is to be used”. Section 33 of PIPA-AB provides that “an organization must make a reasonable effort to ensure that any personal information collected, used or disclosed by or on behalf of an organization is accurate and complete to the extent that is reasonable for the organization’s purposes in collecting, using or disclosing the information.”

417. Accuracy requirements are defined more narrowly under British Columbia’s and Quebec’s laws. Pursuant to section 33 of PIPA-BC, “An organization must make a reasonable effort to ensure that personal information collected by or on behalf of the organization is accurate and complete, if the personal information (a) is likely to be used by the organization to make a decision that affects the individual to whom the personal information relates, or (b) is likely to be disclosed by the organization to another organization.” Section 35 of PIPA-BC requires that personal information used to make such a decision be retained for at least one year. Finally, section 11 of Quebec’s Private Sector Act states that “every person carrying on an enterprise must ensure that any personal information held on another person is up to date and accurate when used to make a decision in relation to the person concerned. The information used to make such a decision is kept for at least one year following the decision.”

418. In its submissions to the Offices, OpenAI represented that LLMs are designed to produce plausible, sensible text in a conversational style, by predicting the next most likely word in a sentence. OpenAI noted that in some cases, the most likely next word(s) might not be the most factually accurate. OpenAI further explained that it is aware that this is an ongoing issue, and noted that there is active research across the entire AI field to improve the factual accuracy of LLMs.

Level of accuracy

419. In this context, OpenAI has carried out internal evaluations of the accuracy of its models. In particular, these were aimed at assessing the performance of GPT-4 on nine internal accuracy evaluations (i.e., learning, technology, writing, history, math, science, recommendation, code and business), compared to earlier versions of ChatGPT. The results of this evaluation, as documented in its [GPT-4 Technical Report](#),¹⁸⁹ suggest that GPT-4 was, on average, 70 to 80%

¹⁸⁹ See [GPT-4 Technical Report](#), Figure 6, p. 10.

accurate, depending on the categories of questions.¹⁹⁰ In comparison, the report states that GPT-3.5 (the basis for ChatGPT before the launch of GPT-4) was found to be on average 19 percentage points less accurate than GPT-4.¹⁹¹

420. This means that, according to OpenAI's own testing, between 20 and 50% of the answers provided by ChatGPT, depending on the version and subject being tested, were deemed factually inaccurate. Furthermore, GPT-3.5, OpenAI's more widely used model at the time of conducting the evaluation (in 2023), was found by OpenAI to be significantly less accurate than GPT-4.

421. OpenAI submitted that those figures were only established on the basis of nine internal evaluations it conducted and were not intended to represent overall system accuracy, nor to assess the accuracy of outputs containing personal information.

422. We asked OpenAI if it could provide any internal report or statistical analysis regarding the accuracy of personal information included in model outputs, for example, in response to a prompt seeking information about an individual.

423. In response, OpenAI referred the Offices to a third-party study that found that ChatGPT had the lowest hallucination rate amongst leading AI services (i.e., 3%).¹⁹² However, this study has limited relevance to our question, in that it relates exclusively to how often LLMs introduce hallucinations when performing a very specific task, namely summarizing a provided document, and it did not include an assessment of ChatGPT's ability to provide accurate personal information in response to a user's prompt.

424. In response to further questioning, OpenAI confirmed that it had not conducted an assessment that would expressly validate the general accuracy of personal information provided by GPT-3.5 and 4 in response to a request about, or in relation to, an individual.¹⁹³

425. During the course of our investigation, we found several examples in the media of inaccurate and harmful statements made by ChatGPT, including false claims of sexual harassment or bribery, which may have deeply impacted the lives of affected individuals.¹⁹⁴

¹⁹⁰ An accuracy of 100% means that the model's answers are judged to be in agreement with human ideal responses for all the questions in the evaluation.

¹⁹¹ GPT-3.5 is still accessible in ChatGPT's free (with or without an account) and Plus versions.

¹⁹² See [AI Models Ranked By Hallucinations: ChatGPT is Best, Palm-Chat Needs to Sober Up](#), Tom's Hardware, November 14, 2023.

¹⁹³ The only evaluations that it performed with regard to the accuracy of personal information were to address the issue of false criminal claims generated by ChatGPT about certain individuals.

¹⁹⁴ See [Luiza Jarovski's News Letter, 101st edition](#), April 30, 2024, which includes a screenshot of a fake biography that ChatGPT generated about her; [Australian mayor prepares world's first defamation lawsuit over ChatGPT content](#), *The Guardian*, April 6, 2023 (in response to our Preliminary Report, OpenAI clarified that the individual concerned ceased to pursue the claimed); [US law professor claims ChatGPT falsely accused him of sexual assault](#),

Although OpenAI addressed these issues reactively in certain cases, OpenAI has publicly acknowledged that GPT-3.5 and 4 can generate factually inaccurate information.¹⁹⁵ Despite this, OpenAI did not provide the Offices with assurances that inaccurate and harmful disclosures such as those outlined above are unlikely to happen again in the future.

426. While OpenAI stated that there is active research across the entire AI field to improve the factual accuracy of LLMs, it is worth noting that many AI experts and researchers believe that hallucinations and inaccuracies are inherent to LLMs and will be a persistent issue because of the way LLMs are designed and their lack of “cognitive” understanding of the world.¹⁹⁶

427. As detailed above, OpenAI’s GPT-3.5 and 4 models were trained based on data sourced from social media websites and discussion forums, which contain vast amounts of personal information. Much of this information would have been subjective and/or biased, reflecting the views of the individuals who posted the information online, including about other individuals. The inherent risk of inaccurate information contained in these sources is such that they were not well suited to training models which would then be used to provide factual responses to user questions.

428. This is especially true given that OpenAI’s mitigation measures in place at the time of pretraining GPT-3.5 and 4 were limited (as described at paragraphs 128 and 133). Furthermore, while we acknowledge OpenAI’s assertion that it fine-tuned its GPT-3.5 and 4 models so that they would not provide private or sensitive information, we found various examples of wrong statements about private individuals being included in model outputs (see above). More generally, we note that privacy harms may not be limited to the disclosure of private and sensitive information and may also result from the disclosure of inaccurate personal information about public figures, or professionals such as physicians.

Accuracy notices and fact verification

429. The accuracy provisions under the Acts do not require the personal information contained in generative AI outputs to be 100% accurate. In particular, PIPEDA provides that

[says 'cited article was never written'](#), *Business Today*, April 8, 2023; [OpenAI's ChatGPT targeted in Austrian privacy complaint](#), Reuters, April 29, 2024 (this matter has been transferred to OpenAI lead supervisory authority in Ireland under the GDPR’s one-stop-shop); and [ChatGPT hit with privacy complaint over defamatory hallucinations](#), TechCrunch, March 19, 2025.

¹⁹⁵ See for instance OpenAI’s March 14, 2023 [Paper on GPT-4](#).

¹⁹⁶ E.g., [ChatGPT ‘hallucinates.’ Some researchers worry it isn’t fixable.](#), *The Washington Post*, May 30, 2023; [We have to stop ignoring AI’s hallucination problem](#), *The Verge*, May 15, 2024; [Bots like ChatGPT aren't sentient. Why do we insist on making them seem like they are?](#), CBC, March 17, 2023; [The hidden meaning of the errors of ChatGPT \(and friends\)](#), Gerben Wierda, November 1, 2023, and [Rage against the machine](#), Alva Noë, Aeon, October 25, 2024.

personal information shall be as accurate, complete, and up-to-date as is necessary for the purposes for which it is to be used.

430. OpenAI represented that the general purpose for which it uses personal information is to develop and deploy Generative AI models that produce contextual relevant text by predicting the next more likely word in a sentence. That said, OpenAI authorizes the use of ChatGPT for a wide variety of purposes, subject to certain limitations listed in its Usage policies.¹⁹⁷ For example, while at the time of writing this report, OpenAI's policies do not permit the automation of high-stakes decisions in sensitive areas (such as employment, financial activities or essential government services) without human review, they do not prohibit the use of its service to support decision-making in these areas. Similarly, OpenAI's policies allow ChatGPT to be used for the provision of tailored advice that requires a license, such as legal or medical advice, as long as there is appropriate involvement by a licensed professional.
431. In addition to the potential privacy-related harms that could result from inaccurate personal information being generated for a purpose that violates OpenAI's policies, it is important to recognize that reputational or other harms could also occur when the tool is being used for a purpose that is allowed by OpenAI's policies, such as when a user asks for the biography of an individual.
432. Therefore, users need to understand the level of accuracy of the personal information included in ChatGPT's outputs in order to determine if that is sufficiently accurate for their intended purposes. If, for example, ChatGPT is used to write a wedding speech, draft a birthday invitation or other similar purposes, there will be a lower expectation regarding accuracy than if it is used to support the assessment of employment or housing applications; in the latter examples, inaccurate outputs or hallucinations could result in significant harms to an individual whose application is negatively assessed based on incorrect information.
433. Recognizing its models' limitations in terms of factual accuracy, OpenAI represented that it displays prominent notices to users in ChatGPT, in addition to notices in its Terms of Use, Privacy Policy and Help Center, warning not to rely on the factual accuracy of outputs from its models, and that users should verify factual accuracy.
434. However, our testing of GPT-3.5 and 4 revealed that OpenAI: (i) provided insufficient notification to ChatGPT users regarding the potential for response information to be inaccurate; (ii) did not clearly or consistently inform users of the need to verify the accuracy of facts

¹⁹⁷ While the scope of our investigation did not extend to the unlimited potential applications of the tool by OpenAI's client, OpenAI indicated that it has controls in place to enforce this policy. However, it provided limited information regarding the measures it employs to ensure that the actual use of ChatGPT, by individual users or third parties such as API customers, is consistent with its Usage policies.

provided; and (iii) did not consistently provide a viable mechanism for users to effectively or reliably verify those facts.

435. In particular, OpenAI represented that there is a permanent disclaimer about accuracy displayed at the bottom of the ChatGPT interface, which reads “ChatGPT can make mistakes. Check important info.” (see figures 4.1 and 4.2 below). However, our testing showed that it was not displayed prominently (i.e., in small grey font and not conspicuous). Furthermore, we noted that the disclaimer was only displayed at the bottom of the page (under the space where users can input their prompts), as opposed to next to the answer itself. As a result, many users may not see the disclaimer.

ChatGPT 3.5 ▾

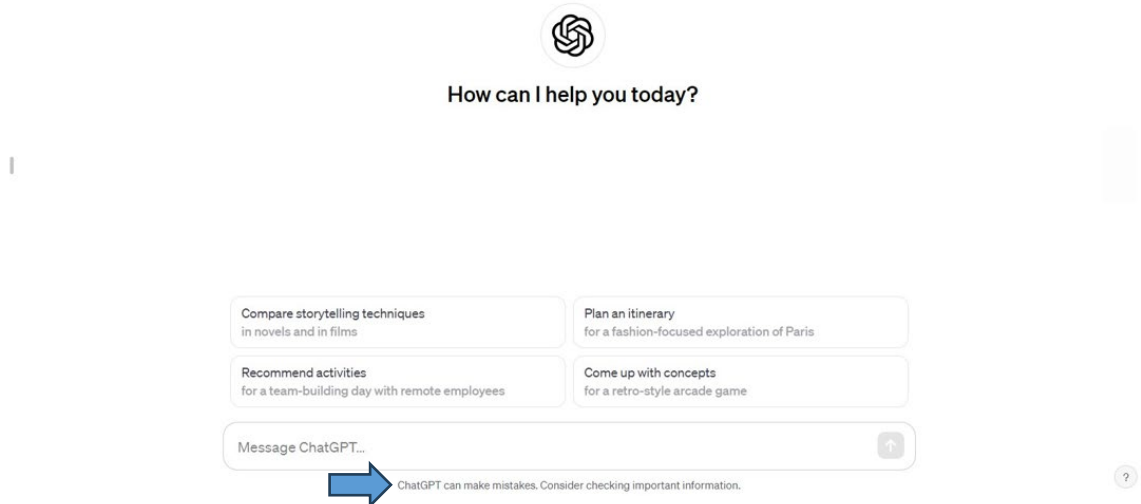


Figure 4.1 ChatGPT interface with accuracy disclaimer at the bottom




Figure 4.2 Accuracy disclaimer

436. During our tests of GPT-3.5 and 4, we further noted that model outputs sometimes included disclaimers such as “as of my last update in...”, “there might have been new developments since then...”, and “I recommend checking news sources...”. Models are indeed

trained on data collected until a certain date (e.g., up to September 2023 for GPT-3.5 and up to October 2023 for the new GPT-4o).

437. However, such notices in model outputs were not provided by default. In some instances, our testing showed that the disclaimers were only provided in the initial responses, but not in subsequent interactions (even where the user had logged out of their account and resumed their conversation with ChatGPT later). As a result, we find that not all users would have received such a notice, and where they did, the notice would not have clearly informed them of the need to verify the facts contained in the response received.
438. Both the authoritative tone in which ChatGPT often presents its answers and the fact that accurate information and hallucinations may be intertwined in model outputs increase the risk that users may overly rely on, and trust, the outputs generated by the models. For example, during our testing, the model inaccurately indicated that an academic had been involved in a specified research project. When prompted to confirm the individual's involvement in the project, ChatGPT confidently maintained its assertion.
439. There are also several well-documented examples of individuals who used ChatGPT in a professional context¹⁹⁸, and then used the information provided to them by the model, without realizing that the responses provided were inaccurate. In certain instances, those individuals asked ChatGPT to confirm if the information was accurate, which it confidently did.¹⁹⁹
440. During our testing of GPT-3.5, OpenAI's most widely used model at the time of the complaint, the sources of information used to generate an output were not provided to the user. The same applied to GPT-4, except in certain circumstances discussed below.
441. Where OpenAI does not include the source(s) behind the personal information provided in ChatGPT outputs, it is neither practical nor feasible for individuals to verify the accuracy of those facts – the information could have been drawn from a wide and unspecified range of websites or datasets from across the Internet, or it could have been 'hallucinated' by the model.
442. Finally, during our testing of GPT-4, we found that the model conducted an Internet search to assist in generating its output (i.e., using a Browser Search feature) in certain circumstances that depended on the nature of the query. In such cases, the links to a limited number of sources used to generate the output were included with the response. The inclusion of sources makes it easier for a user to verify the accuracy of personal information included in the responses, by allowing the user to check those facts against the sources provided.

¹⁹⁸ We recognize that this might have been done in violation of [OpenAI's usages policies](#).

¹⁹⁹ See [B.C. lawyer who used fake, AI-generated cases faces law society probe, possible costs](#), Global News, January 31, 2024, and [ChatGPT: US lawyer admits using AI for case research](#), BBC News, May 27, 2023. See also [Why does ChatGPT generate fake references?](#), Teche, February 20, 2023.

443. At the same time, our testing revealed that, when asked the same questions, GPT-4 disclosed more personal information using the Browser Search feature than GPT-3.5 (which did not have that feature). Furthermore, while OpenAI explained that the feature was automatically triggered when GPT-4 determined that newer or additional information was needed to reply to a user's prompt, the company did not provide details on how the model selected the Internet sources used to generate an output. That said, and as further explained below, we generally see the inclusion of sources as a positive development provided that this feature is implemented in a privacy-protective manner.
444. Ultimately, we find that OpenAI did not provide users with sufficient information for them to be able to determine whether GPT-3.5 and 4's outputs were sufficiently accurate in the context of the specific purposes for which the user intended to use the tool. This included:
- i. not adequately notifying users of the limitations on the accuracy of personal information included in ChatGPT responses – i.e., not prominently notifying them of the general limitations on that accuracy, nor providing specific details on the level of that accuracy;
 - ii. not establishing the level of accuracy of personal information included in its responses, being unable to provide evidence to substantiate that such responses were more accurate than other types of answers, which OpenAI had established to be between 20% and 50% inaccurate;
 - iii. not consistently or clearly informing users of the need to verify the personal information contained in ChatGPT responses for factual accuracy, and not making it obvious to users that this is required; and
 - iv. in the absence of cited sources, like those that were only sometimes provided via GPT-4 and never provided via GPT-3.5, not providing sufficient information, for users to verify the factual accuracy of personal information included in responses.

Findings related to GPT-3.5 and 4

445. Consequently, the OPC and OIPC-AB find that OpenAI contravened principle 4.6 of Schedule 1 of PIPEDA and section 33 of PIPA-AB. Further, the OIPC-BC finds that OpenAI contravened section 33 of PIPA-BC where outputs were used to make a decision that affected the individual to whom the personal information related or were disclosed to another organization.
446. Pursuant to section 11 of Quebec's Private Sector Act, the obligation to keep personal information up to date and accurate applies when the information is used to make a decision in relation to the person concerned and not when the information is collected. The CAI is of the view that the evidence on file does not support a finding that ChatGPT used inaccurate personal information to make decisions in relation to the persons concerned. Despite this, the CAI

encourages OpenAI to implement the recommendations related to accuracy that are made in this report.

Recent developments and conclusion under PIPEDA

447. As outlined below, OpenAI provided additional representations regarding measures it has implemented in relation to accuracy.
448. More specifically, in response to our Preliminary Report, OpenAI represented that it takes concrete steps to improve accuracy and mitigate risks relating to inaccuracies in model outputs. In addition to measures already in place at the time of deploying GPT-3.5 and 4²⁰⁰, OpenAI highlighted the recent development of its filtering tool which detects and masks identifying information in publicly accessible Internet data, licensed datasets and user interactions prior to training, thereby reducing the risk of inaccurate information about individuals being used in training.
449. OpenAI submitted that it has also recently implemented a new web search feature²⁰¹ which, when activated, conducts a real-time web search and references specific web sources for the content output by its models. OpenAI further explained that this web search feature is automatically triggered when ChatGPT's models determine that newer or additional information may be needed to reply to a user's prompt (e.g., when a user is seeking information about recent news and events, time-sensitive facts or highly specific facts or details). Users can also manually activate the feature by clicking the "web search" icon.
450. Since users can consult the sources that the model relied on to generate its response, we accept that this feature will enhance the factual accuracy of model outputs and facilitate users' independent verification of information. That said, and as confirmed by OpenAI, the web search functionality remains a distinct, optional tool. While it will help enhance verifiability in specific use cases, it is not always activated and, as such, does not fully address our concerns regarding the inaccuracy of the model's output
451. Moreover, OpenAI represented that it has started proactively communicating about its assessments of the accuracy of individuals' information found in model outputs through 'model system cards' in a "[Deployment Safety Hub](#)" (formerly "Safety Evaluation Hub"). OpenAI submitted that it initially used the "PersonQA" evaluation which aimed to elicit ChatGPT's hallucinations using a dataset of questions and publicly accessible facts about individuals. This evaluation measured the model's accuracy based on answers delivered by the model when it

²⁰⁰ For example, deduplication and tokenization processes, which OpenAI states lowers the risk of reproduction or repetition of training data, or ex-post measures such as rectification and deletion mechanisms discussed in Issue 5.

²⁰¹ See [Introducing ChatGPT search](#), OpenAI, October 31, 2024. This web search feature is available to both logged-in users and users with no account. Logged-out users can currently perform up to ten queries per day using this feature, whereas logged-in users, including those on the free plan, can perform an unlimited number of search queries.

does not have the ability to browse the web. Based on these results, accuracy levels ranged from 15.5% (for GPT-oss-20b²⁰²) to 70% (for GPT-4.5²⁰³), with a model average of 41% (GPT-3.5 and 4, as well as the more recent GPT-5 were not included in the list of tested models).²⁰⁴ This indicates a relatively low level of accuracy with respect to personal information included in these models' outputs.

452. OpenAI further indicated that, more recently, it has introduced new factuality evaluations for GPT-5, which include asking models open-ended factual questions about people, places or concepts, or prompting them to generate biographical summaries of notable figures.²⁰⁵ OpenAI explained that evaluations were conducted in both “browse-on” and “browse-off” settings, and asserted that their results show that the GPT-5 models have significantly lower hallucination rates across both settings. In particular, OpenAI noted that GPT-5-Thinking makes over five times fewer factual errors than OpenAI o3²⁰⁶, in both browsing settings, across the benchmarks.

GPT-5-Thinking – Hallucination rates:

Evaluation	Example Prompt	Browse-On	Browse-Off
LongFacts Concepts	"What was the relationship like between Thomas Edison and Nicola Tesla?"	0.7%	1.1%
LongFacts Objects	"Who is Xochitl Gomez?"	0.8%	1.4%
FActScore	"Tell me a bio of Samuel Oboh"	1.0%	3.7%

²⁰² This is one of the models developed by OpenAI. See [Introducing gpt-oss](#), OpenAI, August 5, 2025.

²⁰³ *Ibid.* See [Introducing GPT-4.5](#), OpenAI, February 27, 2025.

²⁰⁴ While these results were initially published on the “Safety evaluation hub”, OpenAI later represented that it now includes these evaluations in the models’ System cards. (e.g., OpenAI o3 and o4-mini, as detailed in these models’ [System Card](#), which can be found on OpenAI’s [Deployment Safety Hub](#)).

²⁰⁵ These include the [LongFact](#) and [FActScore](#) benchmarks. LongFact consists of LLM-generated questions asking for detailed responses about either specific objects (e.g., people or places) or broad concepts, while FActScore consists of questions seeking biographies on notable individuals.

²⁰⁶ This is one of the models developed by OpenAI. See [Introducing OpenAI o3 and o4-mini](#), OpenAI, April 16, 2025.

453. This information, which was published in the [GPT-5 System card](#), would appear to corroborate the assertion that OpenAI has reduced the frequency of factual hallucinations, including with respect to personal information.
454. Finally, following discussions with the Offices, OpenAI has committed to linking its updated "[Does ChatGPT tell the truth?](#)" article in the Canadian blog post that it will publish on its website and promote in the Canadian media.²⁰⁷ We find that this will further enhance ChatGPT users' awareness of the potential accuracy limitations of the tool.
455. Consequently, we find this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

Issue 5: Did OpenAI provide individuals with the ability to obtain access to, correct and delete their personal information?

456. For the reasons explained below, we find that OpenAI failed to adequately provide individuals with the ability to access, correct and delete their personal information.
457. OpenAI represented that individuals – both users and other individuals whose information may be collected, used and disclosed via its models – are informed of their privacy rights through OpenAI's Privacy Policy and can exercise certain of those rights through the self-service tools available within their account settings. If an individual is unable to do so, they are invited to submit their requests via OpenAI's Privacy Request Portal or by email.
458. To determine whether OpenAI's practices for responding to such requests are compliant with the Acts, we examine below the different situations in which ChatGPT users, and individuals more generally, may choose to exercise their privacy rights in relation to access, correction and deletion.

Issue 5A: Access to personal information

459. The Acts provide that, upon request, an individual must be informed of the existence, use, and disclosure of their personal information and must be given access to that information. In addition, the organization must provide an account of the use that has been made or is being made of this information and an account of the third parties to which it has been disclosed.²⁰⁸

²⁰⁷ See section relating to OpenAI's response to our Preliminary Report.

²⁰⁸ Principle 4.9 of Schedule 1 of PIPEDA, section 23 of PIPA-BC, sections 24 and 25 of PIPA-AB and section 27 of Quebec's Private Sector Act. Section 27 of Quebec's Private Sector Act stipulates that a company must, at the request of the person concerned, confirm the existence of personal information concerning him or her and provide access to it by allowing the person to obtain a copy. Similarly, computerized personal information disclosed to a requester must be in the form of an intelligible transcript.

460. The Acts further provide that an organization must make specific information about its policies and practices relating to the management of personal information, including the means to gain access to personal information held by the organization, **readily available**. Individuals must be able to acquire this information **without unreasonable effort**, and it must be made available in a form that is **generally understandable**.²⁰⁹
461. The Generative AI Principles provide that developers and providers of generative AI should ensure that procedures exist for individuals to access and correct any information collected about them during their use of the system, and develop processes to permit individuals to exercise their ability to access or correct personal information contained in an AI model, particularly where that information may be included in outputs generated in response to a prompt.
462. OpenAI represented that: (i) in the vast majority of cases, individuals request access to the personal information associated with their ChatGPT user account; and (ii) in exceptional cases, individuals make a specific request for access to their personal information from training data. We examine both types of requests below.

Access to personal information relating to a ChatGPT account

463. OpenAI represented that users who are logged into their ChatGPT account can use its “Export Data” tool to download a copy of their account details, including their account information and the history of their interactions with ChatGPT.
464. As explained below, our analysis of OpenAI’s Export Data tool confirmed that it does not fully satisfy the Access requirements under the Acts.
465. The “Export Data” tool is located on the “Settings > Data Controls” page of the ChatGPT interface. After confirming their request to export their data, users receive an email with a link to their data, which expires after 24 hours. Users are then invited to download a “.zip file”²¹⁰ that includes the history of their interactions with ChatGPT, as well as other data on file with OpenAI (e.g., OpenAI-generated user identifier, email address, and metadata associated with their ChatGPT conversations, including titles and time of creation).
466. In principle, we believe that a well-designed self-service tool can be a helpful initial mechanism for users to access their personal information, and that such a tool might meet the needs of a majority of individuals, even where the tool may not be sufficient to meet all legal

²⁰⁹ Principles 4.8, 4.8.1 and 4.8.2 of Schedule 1 of PIPEDA, section 5 of PIPA-BC, sections 13 and 13.1 of PIPA-AB and sections 3.1 and 3.2 of Quebec’s Private Sector Act.

²¹⁰ A ZIP file is a computer file whose contents of one or more files are compressed for storage or transmission, often carrying the extension .ZIP.

470. However, we observed that while the email sent to users via the Export Data tool mentions the possibility to contact OpenAI through its general “Help Center” page in case of questions, it does not inform the user about the avenues formally available to them if they would like to challenge the completeness, accuracy or nature of the information provided via the tool. As such, users may not be aware of, or think to look for, this alternative mechanism.

Findings related to GPT-3.5 and 4

471. Consequently, we find that OpenAI did not allow individuals to fully exercise their right to access their personal information in relation to a ChatGPT account, in violation of Principle 4.9 of Schedule 1 of PIPEDA, section 24 of PIPA-AB, section 23 of PIPA-BC, and sections 27 and 29 of Quebec’s Private sector Act.

Recent developments and conclusion under PIPEDA

472. In response to our Preliminary Report, OpenAI indicated that it has improved the auto-response email that users receive when they submit an access request to OpenAI by email (at dsar@openai.com). This response now explains how different types of personal information can be accessed. While the email states that most of the user personal data can be obtained directly through their account settings or through the Privacy Portal, it further explains to users that they can respond to the email if they have outstanding questions or requests that need to be addressed.

473. Following discussions with the Offices, as further explained in the section relating to OpenAI’s response to our Preliminary Report (see paragraph 549), OpenAI committed to implementing additional measures with respect to access requests:

- i. It will provide personal information in a more accessible and user-friendly format in its data exports; and
- ii. It will revise the information it shares with users who are seeking a data export, to inform them about the avenues available to them if they would like to challenge the completeness, accuracy, or nature of the information provided.

474. Based on the above, we find this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

Access to personal information contained in OpenAI’s training datasets

475. With regard to personal information contained in its training datasets, OpenAI represented that it only provides access to personal information that is verifiably related directly and uniquely to the requestor, in particular through their email address or telephone number. OpenAI did not provide the Offices with other examples of identifiers that it could rely on to verify the link between the requestor and the information. That said, OpenAI explained that it can conduct limited reviews for a range of identifiers, as long as the identifier information

provided is specific and verifiable so that it can be reliably associated with the requesting individual in order to protect against the disclosure of data to the wrong individual.

476. Open AI explained to the Offices that for personal information that cannot be verified as being associated with the requestor, such as when the requestor has a common name and/or there is no way to otherwise verify attribution of the information to the requestor (e.g., through an email or telephone number, as outlined above), it will only tell the requestor whether the name appears in its training datasets. OpenAI claims that in those instances, given the volume and unstructured nature of the data in its datasets, it is not able to determine with a reasonable degree of certainty and without disproportionate time and effort, whether the information relates to the requestor or to another individual with the same name.

477. Moreover, OpenAI explained that the process of confirming whether verified information is included in its training datasets is in itself an extremely complex and difficult process, because:

- i. Due to the unstructured nature of the training data, there is no index with which to look up and retrieve all data concerning an individual; a search for verified information is therefore not simply a matter of executing pre-existing queries or accessing a centralized database. It requires custom-engineered workflows that must be manually adapted to OpenAI's data architecture.
- ii. Checking for the presence of specific personal information in the training dataset requires significant resources and represents a substantial effort in terms of workload and computing power.
- iii. Any search in the training data set is constrained by the need to avoid disclosing data relating to individuals other than the person who made the request which would affect their privacy rights. As a concrete example, the word following any identifier (that could be included in the unstructured data) may be completely unrelated to that identifier, or even include the personal information of other individuals with the same name.

478. While we recognize that the design of OpenAI's models and the nature of the data it collects to train them creates technical challenges, particularly in the context of authenticating the requestor and identifying information linked to them for purposes of actioning their request, we find that in the context where OpenAI was collecting, using and disclosing vast amount of personal information, including sensitive information, without robust mitigation measures to limit the information collected, it was not doing enough to comply with its access obligations under the Acts.

479. This was even more important in a context where individuals were often unaware that their information had been collected by OpenAI to train its GPT-3.5 and 4 models, let alone that it could be disclosed or relied upon in ChatGPT responses.

Findings related to GPT-3.5 and 4

480. Consequently, we find that OpenAI did not provide individuals with the ability to access their personal information included in the GPT-3.5 and 4 training datasets, pursuant to Principle 4.9 of Schedule 1 of PIPEDA, section 24 of PIPA-AB, section 23 of PIPA-BC, and sections 27 and 29 of Quebec's Private Sector Act.

Recent developments and conclusion under PIPEDA

481. As discussed above, OpenAI explained in response to our Preliminary Report that it has developed an internal filtering tool that detects and masks private information in the training datasets. While we note that OpenAI still collects and retains the raw unfiltered data, we accept that this measure may significantly reduce the use and processing of personal information associated with the training of future models.

482. Furthermore, while OpenAI did not provide the Offices with any recent statistics, it submitted that access requests to personal information from training data are only made on an exceptional basis. While we recognize that, due to the technical challenges described above, OpenAI may not be able to fulfill a subset of this small number of access requests, we also note that where OpenAI cannot uniquely associate that information with an individual, it may not constitute the individual's personal information.²¹² In these circumstances, we accept that the risk of privacy harm resulting from OpenAI's inability to return certain information relating to an individual contained in its unstructured datasets is limited.

483. Therefore, taking into account OpenAI's new mitigation measures, and in line with the pragmatic and flexible approach to the interpretation of PIPEDA (discussed at paragraph 46) and the necessity to balance the privacy rights of individuals with the need for businesses to use personal information for appropriate purposes, the OPC finds this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.²¹³

484. This conclusion is based on our understanding that OpenAI will continue dedicating resources and energy to finding solutions that further mitigate the risk to privacy in areas where access is not practical or feasible.

²¹² Under Quebec's Private Sector Act, information must no longer allow the person concerned to be identified directly or indirectly.

²¹³ The CAI acknowledges the new mitigation measures that have been put in place and the fact that, in certain circumstances, the technical challenge highlighted by OpenAI may constitute a serious practical difficulty within the meaning of paragraph 3 of section 27 of Quebec's Private Sector Act. However, access to personal information held by a company is a matter subject to the dispute review procedures set out in Division V of Quebec's Private Sector Act, and this issue is analyzed based on the facts of each case. Despite this specificity and in order to facilitate the exercise of the right of access to personal information by the individuals concerned, the CAI joins the recommendations of the other Commissioners regarding the issue of access to personal information.

Issue 5B: Correction of personal information

485. The Acts provide that an individual must be able to challenge the accuracy and completeness of their personal information and have it amended as appropriate.²¹⁴
486. Our investigation found that individuals can submit requests to OpenAI to correct their personal information contained in the company's training datasets or ChatGPT outputs, either via OpenAI's Privacy Request Portal or by email.
487. OpenAI represented that if it can verify that the personal information relates to the requestor (as detailed under Issue 5A above), it will use the information provided by the requestor to check whether the model outputs the information mentioned by the individual in their request.
488. If it confirms that the model did output this information, it will then conduct a case-by-case assessment of this information based on various factors (including whether the individual is a public figure) before implementing a corrective measure. OpenAI submitted that it asks individuals to provide sufficient details to identify the inaccuracy, including links to chats and URLs referenced by the chats, as well as sufficient details to determine that the information ChatGPT shares about them in its response is inaccurate. In the situation where multiple people may share the same name, OpenAI takes that into consideration in its case-by-case assessment.
489. If the information is confirmed to be inaccurate, OpenAI stated that it may attempt to fine-tune the models to correct the inaccuracy. However, if it is not possible to correct the inaccuracy due to technical challenges (discussed further below), OpenAI will prevent the personal information in question from appearing in ChatGPT's outputs, thereby ensuring that information that was verified as belonging to the individual requester is no longer generated. To do so, the individual's verified personal information (such as their name) is added to a blacklist. Input and output filters are used to detect when this information is included either in ChatGPT user prompts or in ChatGPT outputs. This prevents ChatGPT from responding to queries asking about the individual's verified personal information (such as their name) or displaying the verified personal information in the output.
490. OpenAI explained that the technical challenges that it would encounter when attempting to correct inaccuracies are related to the complexity of how their models work. We accept that it may be difficult, or almost impossible, for OpenAI to guarantee that ChatGPT will not include in an output, at some point in time, information that has been demonstrated to be inaccurate. This is because, at least in part, ChatGPT responses are not drawn from a database of facts that can be easily corrected; they are based on a vast number of statistical correlations between words (or portions of words), which might, in certain circumstances, yield the

²¹⁴ Principle 4.9 of Schedule 1 of PIPEDA, section 24 of PIPA-BC, section 25 of PIPA-AB and section 28 of Quebec's Private Sector Act.

inaccurate personal information. In those circumstances, blocking the disclosure of personal information relating to the requestor might indeed be the only viable alternative to correction.

491. While we acknowledge OpenAI's efforts in providing a pragmatic solution to respond to correction requests in the face of the technical challenges outlined above, this approach leaves some gaps. For example, where OpenAI is unable to verify that the personal information relates to the requestor, in circumstances like those discussed under the Access section above, OpenAI will not correct or block the information in question. Furthermore, OpenAI provided no evidence regarding the proportion of requests for which it is able to identify the requestor and uniquely link them to the information in question, other than to raise the fact that there are often challenges in doing so.

Findings related to GPT-3.5 and 4

492. Consequently, we find that OpenAI did not provide individuals with the ability to correct their personal information included in the GPT-3.5 and 4 training datasets and outputs, pursuant to Principle 4.9 of Schedule 1 of PIPEDA, section 25 of PIPA-AB, section 24 of PIPA-BC, and sections 28 and 29 of Quebec's Private Sector Act.

Recent developments and conclusion under PIPEDA

493. As discussed above, OpenAI explained in response to our Preliminary Report that it has developed a new internal filtering tool that detects and masks private information in the training datasets. We accept that this may reduce the risk of inaccurate private information being used for model training and potentially disclosed in model outputs.

494. In addition, OpenAI represented that it has recently implemented a new measure that facilitates the processing of correction requests. When an individual submits a correction request, OpenAI can leverage its web search capabilities in response to prompts about that individual to nudge the models to conduct searches in order to retrieve up-to-date publicly accessible information from the Internet about that individual. OpenAI asserted that this effectively addresses the risk that inaccurate information will appear in model outputs.

495. With regard to the technical challenges associated with correction requests, we accept that the introduction of this new measure, together with OpenAI's other mitigation measures,²¹⁵

²¹⁵ These include the possibility of blocking verified personal identifiers from model outputs when such information is flagged as being inaccurate via a correction request, and the training of models to reject prompts seeking private or sensitive information about individuals.

provides individuals with an adequate mechanism to correct their personal information contained in model outputs.²¹⁶

496. OpenAI also explained that it has developed a technical solution to address the situation where there is a public interest in further disclosure of information about a public figure in line with freedom of expression and information rights, but the individual's interests prevail with respect to a specific item of personal information (e.g., because ChatGPT responds inaccurately with respect to that specific item). In that situation, OpenAI can now granularly block specific personal details about a public figure from appearing in an output, while still allowing accurate details about them to be included. We acknowledge that a targeted blocking appears to be a positive development in the circumstances.

497. We therefore find this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

498. This finding is based on our understanding that OpenAI will continue dedicating resources and energy to finding solutions that further mitigate the risk to privacy in areas where correction is not practical or feasible.

Issue 5C: Removal of personal information from OpenAI's models

499. Canadian privacy laws do not grant individuals an explicit right to request the deletion of their personal information.²¹⁷ However, the Acts provide that individuals may withdraw their consent to the collection, use and disclosure of their personal information at any time, subject to legal or contractual restrictions and reasonable notice.²¹⁸ The Acts also provide that an organization must only retain personal information for as long as necessary to fulfill the identified purposes.²¹⁹ As such, where an organization does not have the individual's consent, or another legal basis to retain their personal information, the information should be deleted.

²¹⁶ Regarding Quebec's Private Sector Act, the correction of personal information held by a company is a matter subject to the dispute review procedures set out in Division V of Quebec's Private Sector Act, and this issue is analyzed on a case-by-case basis. In this regard, the CAI takes note of the new measures put in place by OpenAI concerning this issue and encourages OpenAI to implement any other measures that would promote the exercise of this right by the individuals concerned.

²¹⁷ However, section 40 of the Civil Code of Quebec and section 28 of the Quebec's Private Sector Act stipulate that an individual may request to have outdated personal information or information not justified by the purpose of the file deleted.

²¹⁸ Principle 4.3.8 of Schedule 1 of PIPEDA, section 9 of PIPA-BC, sections 9(1) and 9(4) of PIPA-AB. Section 28.1 of Quebec's Private Sector Act stipulates the right to require that a company cease disseminating the personal information if the dissemination of the information contravenes the law or a court order, or if the dissemination of the information causes injury that is greater than the public interest in knowing the information and the cessation of dissemination, re-indexation or de-indexation does not exceed what is necessary for preventing the perpetuation of the injury. In addition, section 40 of the Civil Code of Quebec stipulates the right to have personal information deleted when it is out of date or not justified by the purpose of the file.

²¹⁹ Principle 4.5 of Schedule 1 of PIPEDA, section 35 of PIPA-BC, section 35 of PIPA-AB and section 23 of Quebec's Private Sector Act.

500. OpenAI represented that, although research in this area continues to progress, reverse-training LLMs to account for the deletion of information is not currently feasible because models do not contain or store copies of information that they learned from. OpenAI explained that models are trained through repeated adjustments of billions of weights (parameters) over successive runs of training datasets, and that each step depends on all previous steps. As a result, the influence of any given data point is not preserved in isolation, but is rather diffused and compounded across model weights through subsequent adjustments which cannot be isolated after the fact. OpenAI further stated that, even if a single data point's influence on a parameter could be identified, ChatGPT's models contain no index or records on which data points affected which parameter changes.

501. Therefore, OpenAI explained that it addresses requests for deletion by preventing an individual's verified personal information from appearing in ChatGPT's outputs and filtering out the information from future training runs.

502. Furthermore, OpenAI asserted that, in assessing how to respond to a deletion request, it aims to balance privacy and data protection rights with other public interests (such as public access to information), in accordance with applicable laws. OpenAI also explained that it may refuse to remove personal information from its outputs where the individual in question is someone OpenAI determines to be a public figure (e.g., celebrities or politicians); OpenAI will determine whether an individual is a public figure based on whether they have a significant presence on the Internet (generally, in the form of a Wikipedia page). Where personal information would be part of the public figure's "persona" (e.g., information about which the individual has been vocal), OpenAI noted that it will tend to favour public access to information and refuse to delete. We did not obtain sufficient evidence to evaluate the effectiveness or appropriateness of this process.

503. In any event, as with requests for access and correction, OpenAI confirmed that it will only undertake the actions detailed above when it is able to verify that the personal information uniquely relates to the requestor, which, as discussed above, may often not be possible. This is problematic in the context where OpenAI has collected personal information without a lawful basis.

504. We further note that OpenAI is often not able to delete personal information from existing models. While it may, as an alternative, block personal information from being disclosed in models' outputs in certain circumstances (i.e., using a blacklist as described in paragraph 489, after conducting an output review and a case-by-case assessment), it will not do so where it is unable to verify that the information uniquely relates to the requestor – this will often not be possible.

Findings related to GPT-3.5 and 4

505. Consequently, we find that OpenAI did not provide individuals with the ability to withdraw their consent and have their personal information deleted pursuant to Principle 4.9 of Schedule 1 of PIPEDA, sections 9(1), 9(4) and 35 of PIPA-AB, sections 9 and 35 of PIPA-BC, and sections 28.1 and 29 of Quebec's Private Sector Act.

Recent developments and conclusion under PIPEDA

506. As mentioned above, OpenAI explained in response to our Preliminary Report that it has developed an internal filtering tool that detects and masks private information in the training datasets. We accept that this may significantly reduce the risk of private information being used for model training and potentially disclosed in model outputs, thereby limiting the need for, and reducing the scope of, potential subsequent requests for deletion.

507. OpenAI also indicated that it now has the ability to granularly block specific personal details about a public figure from appearing in outputs, while still allowing relevant details about them to be included. As such, OpenAI states that it can ensure that it continues to provide the public access to information about the public figure that is of interest to them while also ensuring that public figures can avail themselves of their privacy rights. We acknowledge that a more targeted blocking is an improvement.

508. More generally, we accept that given the technical challenges outlined earlier, the combination of OpenAI's various mitigation measures²²⁰ means that individuals now have an adequate mechanism to request deletion of their personal information contained in the outputs and training datasets of models.²²¹

509. The OPC therefore finds this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

510. This conclusion is based on our understanding that OpenAI will continue dedicating resources and energy to finding solutions that further mitigate the risk to privacy in areas where full deletion is not practical or feasible.

²²⁰ These include the possibility of blocking verified personal identifiers from model outputs and filtering them out from future training runs, and the training of models to reject prompts seeking private or sensitive information about individuals.

²²¹ Regarding Quebec's Private Sector Act, the application of section 28.1 of the Act is a matter subject to the review of disagreements provided for in Division V of Quebec's Private Section Act, and this matter is analyzed based on the facts specific to each case. In this regard, the CAI takes note of the new measures put in place by OpenAI to implement any other measures that would promote the exercise of this right by the individuals concerned.

Issue 6: Did OpenAI establish appropriate retention and disposal procedures for the personal information that it collects, uses and discloses?

Analysis under PIPEDA, PIPA-BC and PIPA-AB

511. For the reasons outlined below, we find that OpenAI did not establish appropriate retention and disposal policies and procedures for the personal information that it collected, used and disclosed for the purpose of developing and deploying its GPT-3.5 and 4 models.
512. The Acts provide that organizations must destroy, de-identify or anonymize (depending on the Acts) personal information that is no longer required to fulfill identified purposes.²²²
513. Principle 4.5.2 of schedule 1 of PIPEDA specifically provides that organizations should develop guidelines and implement procedures with respect to the retention of personal information, and that these guidelines should include minimum and maximum retention periods.
514. Similarly, the Generative AI Principles state that developers and providers of generative AI should establish and abide by appropriate retention schedules for personal information, including (as applicable) that contained within training data, system prompts, and outputs. These schedules should both: (i) limit retention for information that is no longer required; and (ii) ensure that information is retained long enough for individuals to exercise their right to access (particularly where a decision has been made about them).
515. Finally, in its [Personal Information Retention and Disposal: Principles and Best Practices](#), the OPC recommends that organizations consider safely disposing of personal information if retaining it any longer would result in a prejudice for the concerned individual, or increase the risk and exposure of potential data breaches.
516. During the evidence-gathering phase of our investigation, OpenAI stated that while it had data management policies relating to sensitive business information, trade secrets or customer data, it was still in the process of developing the specifics of its formal retention and deletion policy for personal information. As a result, we were not able to effectively review its retention practices for all the personal information that it collects, uses and discloses.
517. Nonetheless, OpenAI explained that while its formal retention and deletion policy was still being finalized, it had put in place specific retention periods for various categories of personal information.
518. For example, if a user deletes their account or turns off their chat history, OpenAI noted that it will delete the user's account-tied information or conversations within 30 days, with a few

²²² Principles 4.5, 4.5.2 and 4.5.3 of Schedule 1 of PIPEDA, section 35(2) of PIPA-BC, section 35(2) of PIPA-AB and section 23 of Quebec's Private Sector Act, which sets out the obligation to anonymize.

limited exceptions relating to anti-fraud, legal, or other similar purposes. With respect to users' conversations used to train its models, OpenAI explained that it filters them for personal identifiers, disassociates them from users' accounts and stores them for up to three years. OpenAI also noted that it periodically re-evaluates whether the data is still needed.

519. However, OpenAI represented that it has no retention schedule for the unstructured raw data collected from publicly accessible websites. This data, which is intended to be filtered to create training datasets (i.e., via removal of certain websites and deduplication of websites, and more recently, filtering of personal identifiers, as discussed earlier in this report), is stored "as long as necessary to train successive iterations of OpenAI's models".

520. According to OpenAI, this is critical in the AI context as it allows reproducibility (i.e., executing training runs with known and consistent data sets with different settings and parameters to evaluate the difference in results), auditing (i.e., examining the data used to train the models to determine its impact on the results) and refinement (i.e., refining and improving filtration techniques with each training run), which it represented to be core principles in scientific research.

521. We note that retaining personal information longer than necessary may increase the risk of harm to individuals in case of a privacy breach, especially if the information is sensitive, inaccurate or outdated, and has been collected without appropriate consent. As mentioned above, OpenAI's filtered training data inevitably includes personal information, some of which may be sensitive. Furthermore, given that it has been collected over the course of many years, it is likely to include Web pages that have either been removed from the Internet or updated since they were initially scraped. The raw datasets may include pirated and harmful content, as well as content from adult websites or websites that aggregate personal information about individuals, which OpenAI represented would subsequently be filtered out to create the training datasets.

522. OpenAI and other Generative AI companies hold massive datasets, which may be attractive to potential malicious actors and as such, add to the existing risk of data breaches and other security threats.

523. Moreover, although OpenAI represented that it has already built a process of periodic re-evaluation of its training data retention practices, it did not provide us with any corroborating evidence, such as copies of internal policies, that would allow us to determine whether this process is sufficiently formalized.

Findings related to GPT-3.5 and 4

524. Consequently, we find that OpenAI contravened Principles 4.5, 4.5.2 and 4.5.3 of Schedule 1 of PIPEDA, section 35(2) of PIPA-BC, and section 35(2) of PIPA-AB.

Recent developments and conclusion under PIPEDA

525. In response to our Preliminary Report, OpenAI explained that it has now: (i) developed and implemented a formal “Personal Data Retention and Deletion Policy” that outlines specific rules governing the retention and deletion of personal information processed in connection with ChatGPT; and (ii) adopted a “Customer Data Retention Schedule” that sets defined retention periods for specific categories of personal information collected from users of its platform.

526. With respect to unstructured training data, OpenAI submitted that it has implemented defined retention criteria based on the core principles highlighted above (i.e., reproducibility, auditing and refinement). OpenAI explained that once it has ascertained that a dataset is no longer needed for active training or reproducibility research on the basis of these criteria, it is deprecated. Once deprecated and inactive, OpenAI does not use the dataset in ongoing model development and retains it solely as a historical benchmark for scientific integrity purposes, including to demonstrate the validity of past research. OpenAI further stated that it stores the underlying raw data in a locked-down state in a secure, access-controlled archive, and that access to it is limited to a small group of employees.

527. The OPC accepts that the approach of segregating data, securing and limiting access to that data and using it for scientific integrity purposes only may be reasonable for future datasets collected lawfully, provided that:

- i. Strong protections are in place to ensure that these datasets are used for these purposes only and not for model development purposes.
- ii. The datasets are removed from model development (i.e., segregated) as soon as OpenAI ascertains that they are no longer needed for this purpose.
- iii. Data subject rights continue to apply to these segregated datasets, to the extent that they contain personal information.
- iv. OpenAI regularly re-evaluates whether retention of each dataset remains necessary under the criteria given.

528. OpenAI committed to ensuring that it does, and will continue to, implement the above-listed measures.

529. Consequently, the OPC finds this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

Analysis under Quebec’s Private Sector Act

530. Section 3.2 of Quebec’s Private Sector Act stipulates that any person carrying on an enterprise must establish and implement governance policies and practices regarding personal information that ensure the protection of such information; such policies and practices must, in particular, provide a framework for the keeping and destruction of the information.

531. In this regard, the CAI notes that OpenAI's policies and practices regarding the retention of personal information were not in compliance with section 3.2 of Quebec's Private Sector Act. However, the CAI acknowledges that, during the investigation, OpenAI developed an official policy for the retention and deletion of personal data and adopted a client data retention schedule, as described in paragraph 525.
532. Despite these new measures, section 23 of Quebec's Private Sector Act provides that when the purposes for which personal information was collected or used have been fulfilled, the person carrying on an enterprise must destroy or anonymize it for use for serious and legitimate purposes.
533. In its comments, OpenAI specified that its practice regarding the retention of training data, as described in paragraphs 519 and 520, was directly related to the purpose disclosed when collecting personal information from its users. In order to be consistent with this argument, the CAI considers that the information provided to users when collected by OpenAI should be specific.
534. OpenAI states that, in accordance with subsection 8(4) of Quebec's Private Sector Act, it provides users, **upon request**, with the retention period for the information it collects. However, the CAI considers that providing this information upon request does not bring OpenAI's retention practices into compliance with section 23 of Quebec's Private Sector Act and that users should be specifically informed of all the purposes for which the information is collected at the time of collection.
535. In the absence of properly and specifically informing ChatGPT users of the specific purposes related to the ongoing development of the model and historical reference for scientific integrity purposes, the CAI recommends that OpenAI, once the purposes of the collection have been fulfilled, proceed with the anonymization of the personal information retained²²³ to ensure that its retention practices comply with section 23 of Quebec's Private Sector Act.
536. In this regard, despite the new measures put in place, the CAI concludes that OpenAI's practices regarding the retention of personal information do not comply with section 23 of Quebec's Private Sector Act.

Issue 7: Did OpenAI meet its accountability requirements in respect of the personal information under its control?

537. For the reasons outlined below, we find that OpenAI did not meet its accountability requirements in respect of the personal information under its control.

²²³ Anonymization must comply with section 23 of Quebec's Private Sector Act and the *Regulation respecting the anonymization of personal information*, CQLR c A-2.1, r 0.1.

538. The Acts state that an organization is accountable for personal information under its control. It must designate one or several individuals to oversee the organization's compliance with the Acts and implement policies and practices to that effect.²²⁴

539. OpenAI pointed to various measures the company has implemented to comply with the accountability requirements under the Acts. We recognize that OpenAI has put in place a number of structures, policies and practices to protect the personal information under its control. For example, it has:

- i. Created, during the course of our investigation, a Governance Risk and Compliance team responsible for ensuring that all OpenAI's functions are meeting the contractual, statutory, regulatory and internal executive leadership mandates and requirements. OpenAI explained that this team has worked closely with OpenAI's security and legal teams to elaborate and provide security and privacy training to its employees and contractors, including about escalation measures to quickly respond to incidents.
- ii. Designated an external Privacy Officer and developed privacy communications (i.e., Terms of Use, Privacy Policy, Help Center) and procedures to address privacy complaints and inquiries.
- iii. Implemented risk mitigation measures at various phases of its AI models' development and deployment, as discussed in paragraphs 128, 294 and 378.
- iv. Established security procedures to protect personal information throughout its lifecycle.
- v. Devoted resources to red-teaming²²⁵, announcing in September 2023 that it would be launching a global Red Teaming Network, bringing together experts with specialized knowledge in various disciplinary fields, including privacy.

540. However, while these measures represented steps in the right direction, the various defaults identified in this report demonstrated a lack of accountability on the part of the company.²²⁶ In particular, after having indiscriminately collected the personal information of millions of individuals in Canada and used it to train ChatGPT, OpenAI deployed this service without having first:

- i. established the level of accuracy of personal information disclosed via model outputs – instead, it took a remedial approach to correcting systemic accuracy issues when they were discovered; and

²²⁴ Principle 1 of Schedule 1 of PIPEDA, sections 4(2), 4(3), 5 and 34 of PIPA-BC, sections 5, 6 and 34 of PIPA-AB and section 3.1 of Quebec's Private Sector Act.

²²⁵ "Red teaming" means using people or AI to explore a new system's potential risks in a structured way. See [Advancing red teaming with people and AI](#), OpenAI, November 21, 2024.

²²⁶ Principle 4.1.3 of Schedule 1 of PIPEDA, sections 4(2), 4(3), 5 and 34 of PIPA-BC, sections 5, 6 and 34 of PIPA-AB and sections 3.1 and 3.2 of Quebec's Private Sector Act.

- ii. developed a retention policy for personal information collected for the purpose of developing and deploying its models.

541. Media reported that one of OpenAI’s cofounders expressed that the company had had concerns about ChatGPT’s lack of accuracy and propensity to generate unwanted outputs when it released the tool in November 2022:²²⁷

“Our biggest concern was around factuality, because the model likes to fabricate things. But (...) other large language models are already out there, so we thought that as long as ChatGPT is better than those in terms of factuality and other issues of safety, it should be good to go. Before launch we confirmed that **the models did seem a bit more factual and safe than other models, according to our limited evaluations, so we decided to go ahead with the release.”** (our emphasis added)

542. OpenAI does not dispute that this statement was made but argued that it should not be relied upon for findings on accountability. In our view, however, statements such as the one above are relevant to our assessment in that they demonstrate that OpenAI released ChatGPT without having implemented processes and practices to comply with privacy laws and properly mitigate against risks that were known at the time.

543. This failure to be accountable exposed individuals to risks of harm including breaches of their personal information, inaccuracy of information, discrimination on the basis of accurate and inaccurate information about them, in addition to other easily foreseeable individual and social harms beyond privacy that are outside of the mandate of the Offices.

Findings related to GPT-3.5 and 4

544. Consequently, we find that OpenAI contravened Principle 1 of Schedule 1 of PIPEDA, sections 4(2), 4(3) and 5 of PIPA-BC, sections 5 and 6 of PIPA-AB and sections 3.1 and 3.2 of Quebec’s Private Sector Act.

Recent developments and conclusion under PIPEDA

545. In response to our Preliminary Report, OpenAI submitted that it has implemented a number of measures to enhance its data governance and privacy framework, including (as described in various sections of this report):

- i. development of an internal filtering tool that detects and masks personal information in training data, so the models do not learn from it;

²²⁷ [The inside story of how ChatGPT was built from the people who made it](#), MIT Technology Review, March 3, 2023.

- ii. improvements to individual rights request-handling processes;
- iii. measures to improve the accuracy of model outputs, including through the web search feature;
- iv. development of accuracy evaluations for information about individuals; and
- v. implementation of formal retention policies and schedules.

546. Recognizing the extensive measures (including policies and practices) that OpenAI has implemented since the commencement of our investigation to mitigate privacy risks, and the additional commitments that it has made to the Offices, the OPC finds this aspect of the complaint to be **well-founded and conditionally resolved** under PIPEDA.

Recommendations

547. As initially conveyed to OpenAI in our Preliminary Report, and as further detailed in this report, the Offices found that the company developed and deployed its GPT-3.5 and 4 models in a manner that contravened the Acts.

548. Therefore, with a view to allowing the development and deployment of generative AI in Canada in a sufficiently privacy-protective manner, the Offices made a number of recommendations to OpenAI in the Preliminary Report. These recommendations are listed below.²²⁸

Recommendations
Within 3 months of the issuance of the final Report of Findings in this matter:
<ul style="list-style-type: none"> i. [Limiting Collection / Necessity] Develop and share with the Offices a plan for limiting the personal information used to train its models to that which has been established, through research and testing, to be necessary and proportional for this purpose.

²²⁸ These do not include the additional recommendations that the CAI has issued in accordance with the specific provisions of its law, which are set out in the conclusion of this report.

a. Measures should include:

- the implementation of processes and technical means that minimize the collection of personal information for the purpose of training its AI models, including the cessation of collection of training data from sources containing significant personal information, including but not limited to social media and discussion forums;²²⁹
- the implementation of measures during the pre-training and fine-tuning phases, such as the use of synthetic data or removal of a larger proportion of personal information from training datasets.

ii. **[Limiting collection of sensitive information via user interactions]** Ensure that users are clearly informed of and can reasonably understand the potential consequences of disclosing sensitive information when interacting with ChatGPT.

- a. Measures should include displaying a permanent, conspicuous notice in the ChatGPT interface.

iii. **[Consent]** Develop and share with the Offices a plan for the implementation of measures to ensure that, moving forward, valid consent under the Acts is obtained from individuals in Canada whose personal information it collects, uses and discloses for the purpose of developing and/or deploying ChatGPT. This plan should address the specific requirements under each Act, including the following:

- a. Pursuant to PIPEDA, PIPA-AB, PIPA-BC (with the exception of Quebec's private sector Act): Where the personal information collected, used or disclosed – whether it is from publicly accessible websites, licensed sources or user interactions – is sensitive and/or where the collection, use or disclosure is outside the reasonable expectations of the individuals in question, OpenAI should generally ensure that express consent has been obtained. Where OpenAI wishes to rely on implied consent, it should take measures to ensure that the information is not sensitive and that the collection, use or disclosure falls within Canadian users' reasonable expectations.
- b. Pursuant to PIPEDA, PIPA-AB, PIPA-BC (with the exception of Quebec's Private Sector Act): Where OpenAI collects personal information from licensed third-party sources, for use and/or disclosure, it should implement measures and processes to ensure that those third parties have obtained valid consent from individuals in Canada for their personal information. That consent must permit

²²⁹ See, for instance, the EDPB's [Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models](#), adopted on December 17, 2024 (at paragraphs 105 and 106).

OpenAI to collect the personal information from the third party to be used and/or disclosed for its own purposes. If this consent cannot be obtained in the manner indicated, OpenAI must obtain valid consent directly from the individuals concerned.

- c. Pursuant to Quebec's Private Sector Act: Where OpenAI collects personal information from third parties rather than directly from the person concerned – whether it is from publicly accessible websites through data scraping or from licensed data sources – uses this information or discloses it, it should implement measures and processes to ensure that those third parties have obtained valid consent from the individuals concerned for their personal information to be collected from them, used or disclosed by OpenAI. Otherwise, OpenAI must obtain valid consent directly from the individuals concerned.
- d. Pursuant to all the Acts: With respect to the collection, use and disclosure of personal information in users' interactions with ChatGPT, prior to collection, OpenAI must provide explicit notice to users stating the purposes for, and the means by, which the information is collected, used or disclosed. The notice must also provide details of how users can exercise their rights of access and rectification/correction provided by law and their right to withdraw consent to the disclosure or use of the information collected.

With respect to personal information that has already been collected, it is recommended that:

- e. OpenAI implement measures to untrain its models to no longer use and/or disclose the personal information of individuals in Canada that was collected without consent. If OpenAI takes the position that it is not possible to untrain the models, OpenAI must demonstrate why it is not possible to do so. In any event, OpenAI must ensure that such information is not used for future training of new models (i.e., models that are not founded on previous models trained with existing data). Future training is to be based only on personal information for which valid consent has been obtained.

iv. **[Access and Correction]** Develop and share a plan with the Offices for implementation of measures to ensure that individuals in Canada whose personal information is being or has been used for model training:

- a. are clearly informed of their individual right to access and correct their personal information included in existing and future training datasets and in model outputs, and the right to withdraw their consent to the collection, use and/or disclosure of their personal information used in model training and outputs where valid consent was previously obtained; and

b. can effectively exercise these rights.

v. **[Accuracy]** Develop and share a plan with the Offices for implementation of measures detailed in recommendation vii.

Note: For the purpose of the recommendations above, plans should include a reasonable timeframe and milestones for implementation. The proposed measures and associated deadlines for implementation will be subject to acceptance by the Offices.

Within 6 months of the issuance of the final Report of Findings in this matter:

vi. **[Model transparency]** With respect to its LLMs, provide the general public with plain language, comprehensive and easily accessible information about:

- a. the categories and sources of information used to build the training datasets;
- b. how the models function, including how this can lead to inaccurate outputs; and
- c. the existing limitations on model explainability.

vii. **[Accuracy]** Implement measures to ensure that, when interacting with ChatGPT, users in Canada, including business users, are made aware of the current limitations of the models in terms of the general level of accuracy of personal information included in outputs, so that they can determine whether outputs are sufficiently accurate for the intended purposes. These measures should include, without limitation:

- a. conducting an assessment (research and testing), or alternatively having a qualified third-party conduct such an assessment, to establish the general level of accuracy of the personal information included in ChatGPT's outputs;
- b. being more prominent, explicit and transparent regarding the level of and limitations on the accuracy of model outputs, via measures that could include:
 - systematic disclaimers in model outputs (e.g., date of last model update, advising to check sources, etc.); and
 - a more conspicuous permanent disclaimer about accuracy in the ChatGPT interface, including to communicate the established level of

	<p>accuracy of personal information included in model outputs (see recommendation (vii)(a)).</p> <p>c. providing a mechanism whereby individuals/users can verify the accuracy of the personal information provided in model outputs. These measures should include, without limitation:</p> <ul style="list-style-type: none"> • systematically providing source links for personal information included in model outputs, and where this is not possible, highlighting the specific facts for which there is no source available.
viii.	<p>[Access] Implement measures to ensure that the format of the files/information provided to users using the “Export Data” tool is accessible and user-friendly for the general public.</p> <p>a. This should include giving the users the option to challenge the accuracy of information provided via the Export Data tool or request access more formally, consistent with their rights under the Acts. This option should be clearly explained in the Export Data tool communications and responses.</p>
ix.	<p>[Retention] Develop a formal retention and deletion policy for the personal information it collects about individuals in Canada to ensure that data is retained only as long as necessary for identified purposes. This should include retention periods for the datasets used to train successive iterations of its models, and the unstructured raw data used to build those datasets.</p>
x.	<p>[Accountability] Implement accountability measures to address the above recommendations, including by developing and/or updating governance models, policies and practices, as well as associated employee training, with respect to the changes made to comply with our recommendations.</p>

OpenAI’s response

549. OpenAI expressed its disagreement with our findings. It asserted that it was compliant with the Acts in most respects, through a combination of its existing practices (including newly implemented measures), and associated communications, which it stated would address our recommendations.

550. OpenAI nonetheless engaged extensively with the Offices with respect to our findings and recommendations. It provided details of the measures it had recently implemented, and

further discussed commitments that would resolve the matter. Further to these discussions, OpenAI committed to implementing the following additional privacy enhancing measures:

- i. [Openness and model transparency] Concurrently with the publication of this report, it will publish a Canadian blog post on its website explaining its privacy practices and take measures to promote the post and its contents in the Canadian media. The blog post will inform individuals that, among other things, users' interactions may be reviewed and used to train its models, advise users not to share sensitive information via their interactions with the ChatGPT, address the question of the accuracy of its models (by adding a link to its updated "[Does ChatGPT tell the truth?](#)" article in the blog) and provide information about the categories of content used to train its models.
- ii. [Openness and model transparency] Within three months of the issuance of this report, it will expand its "[How ChatGPT and our foundation models are developed](#)" article to include more plain-language explanation about the sources of information used to train its models;
- iii. [Openness and model transparency] Within three months of the issuance of this report, it will, in the signed-out ChatGPT web experience – before the individual inputs their first user prompt – provide notice that chats may be reviewed and used to train models and advise users not to share sensitive information.
- iv. [Access] Within six months of the issuance of this report, it will (i) provide personal information in a more accessible and user-friendly format in its data exports, and (ii) revise the information it shares with users who are seeking a data export, to inform them about the avenues available to them if they would like to challenge the completeness, accuracy, or nature of the information provided.
- v. [Retention] Within six months of the issuance of this report, with respect to future datasets collected lawfully, which are deprecated and solely used as a historical benchmark for scientific integrity purposes, it will:
 - a. confirm in a report to be provided to the Offices that strong technical and organizational controls are in place to ensure that the datasets retained for related scientific integrity purposes are not used for active model development once they are no longer needed for that purpose;
 - b. to the extent that these retained datasets contain personal information, continue to comply with applicable data subject rights, as required by law; and
 - c. continue its existing process of regularly re-evaluating whether retention of each dataset remains necessary pursuant to its established criteria.
- vi. [Children's privacy] Within six months of the issuance of this report, it will test the addition of a protective measure for the minor family members of public figures (who are not themselves public figures), so that the models refuse requests for the

name or date of birth of minor family members of such individuals, even if such information is publicly accessible through a current online citation.

- vii. [Reporting] It will provide the Offices with quarterly reports to confirm and demonstrate, with detailed submissions and corroborating evidence, compliance with the above commitments until all have been met.

551. Finally, OpenAI informed the Offices that it has deprecated (i.e., retired) its GPT-3.5 and 4 models and confirmed that the new mitigation measures, including the filtering tool, were used throughout the development and deployment of the current models powering ChatGPT.²³⁰

552. Recognizing that OpenAI's commitments vary in many respects from the specific measures recommended in our Preliminary Report, we assessed the adequacy of those commitments in addressing the intention of each of our recommendations.

Conclusion

OPC

553. Given all the above, and in line with a pragmatic and flexible interpretation of PIPEDA and the necessity to balance the privacy rights of individuals with the need for businesses to use personal information for appropriate purposes, the OPC is satisfied that OpenAI's commitments sufficiently address the intent of the recommendations to resolve the contraventions that we identified. Therefore, the OPC finds the matter to be **well-founded and conditionally resolved**.

554. The OPC will continue to work with OpenAI to ensure the final resolution of the matter through its implementation of the agreed upon recommendations.

OIPC-AB and OIPC-BC

555. The OIPC-BC and OIPC-AB similarly took a pragmatic and flexible approach to interpreting the respective legislation, as is consistent with the modern approach, but these statutes are, in certain key areas, more specific and explicit than PIPEDA. In particular, these statutes meet the standard set in PIPEDA related to appropriate purpose and thus are substantially similar but are more specific than PIPEDA. For this reason, the OIPC-BC and the OIPC-AB did not have the latitude to interpret the statutes as broadly as the OPC did.

²³⁰ Furthermore, OpenAI confirmed that it did not use GPT-3.5 and 4 as base models for the training of the current models.

556. As detailed in Issue 2 (Consent), the OIPC-BC and the OIPC-AB find that OpenAI's models are based on scraped data for which OpenAI has not obtained, and cannot obtain, consent under PIPA-BC and PIPA-AB. While the OIPC-AB and the OIPC-BC are encouraged by the new measures aimed at compliance taken by OpenAI since this investigation has been initiated and those which it has further committed to implement, these are not sufficient to meet the foundational requirement for consent in PIPA-BC and PIPA-AB. Despite this finding, the OIPC-BC and the OIPC-AB joined the OPC and the CAI in making the joint recommendations and in monitoring the implementation of the measures to which OpenAI has committed.

CAI

557. The CAI considers the aspects of the complaint related to Issues 1, 5 and 7 (i.e., appropriate purposes, individual rights and accountability) to be **well-founded and conditionally resolved**, and those related to Issues 2 and 6 (i.e., consent and retention) to be **well-founded and unresolved**. Given the specificities of its legislation, the CAI did not issue a finding on Issues 3 and 4 (i.e., openness and accuracy).

558. Furthermore, the CAI has made additional recommendations with respect to consent and retention in accordance with the specific provisions of its law. Specifically, the CAI recommends that OpenAI implement the following measures:

Regarding consent:

- i. When collecting personal information for the purpose of training its models from sources of information accessible on the web, either directly through data scraping or from third parties who have themselves collected this information through data scraping, to implement any reasonable process or measure to ensure that the individuals concerned have been clearly informed at the time of the initial collection of their personal information that, by providing such information, it will be made public and could therefore be collected and used by third parties, in particular for the purpose of training artificial intelligence models, and communicated by them, and to ensure that the disclosure of such personal information does not constitute communication by a third party without consent or a communication of personal information concerning an individual under 14 years of age, without the consent of the person having parental authority or of the tutor.²³¹
- ii. Regarding the information provided to users of the web-based, free version of ChatGPT about the use of their chats for model training, the CAI recommends that OpenAI add a notice or pop-up window to its free online version upon registration

²³¹ The CAI reserves the right to conduct any verification or investigation, and to make any additional recommendations or issue any orders relating to OpenAI's implementation of this recommendation.

and before first use, informing users that their chat may be reviewed and used for model training purposes and therefore not to share sensitive information.²³²

- iii. Regarding section 9.1 of Quebec's Private Sector Act, the CAI recommends that OpenAI modify the privacy settings of its systems so that, by default, user chats are not used for training its models unless users activate the feature.²³³

Regarding retention:

- iv. The CAI recommends that OpenAI specifically inform ChatGPT users about the retention of personal information for historical reference purposes for scientific integrity or, failing that, once the purposes of collection have been fulfilled, destroy or anonymize such information in order to ensure that its retention practices comply with Quebec's Private Sector Act.²³⁴

559. The CAI intends to monitor OpenAI's implementation of the joint, and Quebec's specific, recommendations. The CAI will take this into consideration in its assessment of whether to undertake any additional verification or investigative action and/or issue any further recommendations or orders related to the compliance of OpenAI's practices with Quebec's Private Sector Act.

General considerations and expectations

560. More generally, the Offices expect that OpenAI will employ a privacy-by-design approach to refining its current product line, and to developing future services and products. In addressing the concerns we raised throughout our investigation, OpenAI demonstrated an ability to find innovative solutions to mitigate privacy risks. We expect that the company will continue to innovate in favour of privacy in the future.

561. Finally, while this report aimed at addressing and mitigating the risk to privacy associated with the development and deployment of LLMs, we recognize that this technology raises many other questions and challenges, including societal and ethical ones, which regulators, academics and courts around the world are currently trying to assess and address. We trust that this collective effort will contribute to further shaping and defining a robust framework for the future development of Generative AI, in Canada and around the world.

²³² *Ibid* footnote 231.

²³³ *Ibid* footnote 231.

²³⁴ *Ibid* footnote 231.

Appendix A – Summary of current key mitigation measures implemented by OpenAI at the various stages of its models’ development and deployment

562. The table below provides an overview of the key mitigation measures (described in this report), which OpenAI has implemented to limit the amount and sensitivity of personal information collected, used and disclosed by its models.

Data Collection Stage	
Measures to Avoid Collecting Unwanted or Potentially Harmful Information	<p>When collecting publicly accessible data, OpenAI does not circumvent paywalls or websites’ login requirements, or obtain information from the “dark web” or closed user groups.</p> <p>OpenAI also filters out a range of information from training datasets, such as websites known to primarily offer personal information (e.g., sites for genealogic research, websites that provide personal contact details), and other categories of websites that are known to contain factually inaccurate, unreliable, or potentially harmful information, such as sites containing pirated or other illegal content, erotic content, hate speech, adult content, and spam.</p> <p>This reduces the risk of collecting potentially sensitive information about individuals to train OpenAI’s models.</p>
Means to Object to Web Crawling	<p>Website owners can disallow GPTBot—which is used to crawl and scrape content that may be used in training OpenAI’s models—from accessing their site using robots.txt.</p>

	<p>OpenAI provides instructions that enable webmasters to configure their robots.txt tags to indicate to visiting web crawlers and other web robots which portions of the website they are allowed to visit.</p> <p>Website owners can also contact OpenAI if they do not want their pages used to help teach OpenAI’s models.</p>
<p>Pre-Training Stage</p>	
<p>Unstructured Datasets</p>	<p>Datasets used for pre-training contain no index or relationship network to link given data points to specific individuals, meaning that there is no systematic way to (i) retrieve all information about an individual, (ii) verify if data points pertain to a given individual, or (iii) build profiles about individuals. This mitigates the risk that information may be linked to an individual and used or disclosed in a way that might harm them.</p>
<p>Tokenization</p>	<p>Training data is tokenized, further limiting the risk of potential privacy harms. Tokenizing means converting the information into numerical representations, or visual “patches”, which are representations of videos and images as collections of smaller units of data.</p> <p>By tokenizing training data, OpenAI ensures that all data shown to the model, including personal information, is not used in its original identifying format.</p>
<p>Deduplication</p>	<p>OpenAI reduces the amount of personal information contained in training datasets by</p>

	detecting and removing duplicate copies of identical information.
Masking	<p>OpenAI significantly reduces the processing of personal information by detecting identifying information (like personal phone numbers, email addresses, and home addresses, as well as private individuals' names and social media handles) and masking that information during the training process.</p> <p>OpenAI researched, developed and then implemented an internal tool that can identify additional categories of personal information which may be included in training data and mask this information prior to it being used for training so the models do not learn from it.</p>
Filtering following Individuals' Requests	OpenAI filters individuals' verified personal information from future model training runs following a valid erasure request. The individuals' verified personal information (such as their name) is added to a blacklist and excluded from future training runs.
Post-Training stage (Fine-Tuning)	
Disassociating User Conversations from User Accounts	OpenAI disassociates user conversations from user accounts prior to using such data for model improvement training, thereby limiting the extent to which any information can be directly linked back to an individual.
Filtering of User Conversations	OpenAI has implemented automated filters to remove identifying information from user conversations prior to using such data for training

	<p>or model improvement. OpenAI also developed a tool that identifies a wide range of personal information that may be included in user conversations and redacts this information prior to storing conversation data and using it in the training process, so the models do not learn from it.</p>
<p>Training for Refusals to Limit Personal Data in Outputs</p>	<p>The models are trained to refuse to provide private or sensitive information about people even if the information is publicly accessible on the open Internet and would be provided via search engines.</p>
<p>Training for Refusals to Avoid providing Training Data in Outputs</p>	<p>Models are post-trained to reduce the risk of repeating training data and reproducing training data in the models' output.</p>
<p>Training for Refusals to Avoid (Un-)grounded and Sensitive Inferences</p>	<p>Models are post-trained to avoid making ungrounded or sensitive inferences about individuals based on a video, image or audio ("ungrounded inference" means an attempt to make inferences about an individual that cannot be determined solely from video, audio or image, such as their intelligence, socioeconomic status, or sexual orientation).</p>
<p>Training for Refusals to Avoid Identification</p>	<p>Models are post-trained to not attempt to identify individuals based solely on images, video, or audio data.</p>
<p>Red-Teaming</p>	<p>OpenAI works with internal and external experts (red-teaming) to assess potentially harmful content in outputs, including in relation to privacy.</p>

	<p>OpenAI’s Red Teaming Network consists of a diverse community of trusted external experts—including individual subject matter experts, research institutions, and civil society organizations—who help identify risks across cybersecurity, biological and chemical threats, societal harms, child safety, education, fairness/bias, privacy, mis/disinformation, and many other domains.</p>
--	---

Deployment stage

<p>Web Search Feature</p>	<p>To the extent that personal information may be output by ChatGPT (e.g., non-private or non-sensitive information about public figures), ChatGPT can leverage its web search capabilities in real time to nudge its models to conduct searches to retrieve up-to-date publicly accessible information and sources before responding to a prompt. This addresses the risk of any potentially inaccurate personal information appearing in ChatGPT’s outputs.</p> <p>By providing source references in output, the web search functionality also serves to ensure that users can independently verify the information presented.</p>
---------------------------	--

<p>Giving effect to Individuals’ Rights</p>	<p>OpenAI has implemented a range of technologies (accessible via the Privacy Portal) designed to allow individuals to exercise their rights (such as access, rectification, and deletion) at each stage of data processing, from model training to model output.</p> <p>When OpenAI filters out the individual’s verified personal information from appearing in</p>
---	---

ChatGPT's output (in response to a rights request), it also filters it out from future model training runs.

OpenAI provides users with a variety of user controls, including the option to delete their conversation history from ChatGPT (and other OpenAI services) and opt-out of their conversation data being used for model training.