Attachment 1: Guidance for the Responsible Use of Generative Artificial Intelligence

Guidance for the Responsible Use of Generative Artificial Intelligence

Issued On: June 18, 2025
Issued By: Technology Services Division
Subject: Generative Artificial Intelligence
Keywords: Artificial Intelligence, Generative AI

Contents

Executive Summary

Generative Artificial Intelligence (generative AI) refers to a class of AI technologies capable of producing content such as text, images, code, audio, and video based on user-provided inputs. Examples include chatbots, code generators, image and video creators, and tools that summarize or draft documents.

The City of Toronto recognizes the potential of generative AI to enhance service delivery, drive innovation, and improve operational efficiency. The City is committed to a responsible and measured approach that ensures these tools are used safely, ethically, and in compliance with applicable laws, internal policies, and record-keeping obligations. Their use is also contingent upon a clear understanding and effective management of associated risks. The Principles for Responsible, Trustworthy and Privacy-Protective Generative AI Technologies jointly promoted by federal, provincial, and territorial privacy authorities were considered during development of this guidance.

Key risks associated with generative AI include, but are not limited to:

- Privacy breaches resulting from improper handling, unauthorized use, or exposure of personal information;
- Cybersecurity threats due to increased attack surfaces or use of malicious AI-generated content (e.g. computer code);
- Bias embedded in AI models, which can perpetuate or amplify discrimination and inequity;
- Misinformation or inaccurate outputs that may mislead users;
- Reduced transparency and accountability when AI-generated content lack clear traceability or human oversight;
- Copyright infringement and intellectual property rights violations.

This document, organized into **Principles**, **Guidelines**, **Frequently Asked Questions**, and **Use-Case Examples** is intended to inform any individual, staff member, contractor, volunteer, or other representative using generative AI tools for City purposes.

The City's Position on Generative AI

- **Staff should use only City-approved generative AI tools for tasks related to City work**. Using only approved tools ensures alignment with the City's Acceptable Use Policy as well as privacy, security, information governance, and data protection policies. Use of other Generative AI tools or software not approved or supplied by the City is prohibited.

- Generative AI-related guidelines and policies will evolve over time, with input from internal and external stakeholders. This guidance reflects the City's current

approach and will be updated regularly to stay aligned with technological advancements and regulatory changes.

- Technology Services Division actively monitors the use of Generative AI tools across the City. In coordination with Legal Services, the Office of the Chief Information Security Officer (OC) and the City Clerk's Office (CCO), tools deemed high-risk may be blocked.

- This document should be read in conjunction with relevant City policies, guidelines, and legislative obligations, including but not limited to:

    - Acceptable Use of Information Technology Assets Policy (AUP)
    - Cyber Security Policy – Office of the Chief Information Security Officer (CISO)
    - Protection of Privacy Policy
    - Information Collection Policy
    - Privacy Impact Assessment Policy
    - Privacy Breach Protocol
    - Information Management Accountability Policy
    - Human Rights and Anti-Harassment/Discrimination (HRAP)
    - Toronto Public Service By-Law
    - Digital Infrastructure Strategic Framework
    - Municipal Freedom of Information and Protection of Privacy Act (MFIPPA)
    - Personal Health Information Protection Act, 2004 (PHIPA)

Principles

The following principles provide a foundation for the responsible use of generative AI at the City. They are intended to guide the development, implementation, usage, and oversight of generative AI tools.

Human-Centered Design

Generative AI is developed and/or implemented with a human-centered focus that prioritizes the needs, values, rights, and well-being of individuals and communities.

Transparency

The purpose and use of generative AI is communicated and disclosed.

Privacy

Privacy must be preserved in all generative AI usage. City staff must safeguard personal information (PI), personal health information (PHI) and sensitive data from unauthorized collection, use, disclosure, or retention. These types of data and information must not be used with generative AI tools. Generative AI usage must comply with applicable privacy legislation and incorporate data minimization, purpose limitation, and appropriate access controls by design and by default.

Accountability

Clear roles and responsibilities govern the acquisition, deployment, operation, and sustainment of generative AI.

Fairness

Generative AI supports equitable outcomes by mitigating bias and preventing harm to individuals and groups.

Security

Generative AI tools are developed, deployed, and maintained with safeguards that prevent unauthorized access, manipulation, or misuse.

Accuracy

Generative AI outputs are reviewed by humans to ensure they are accurate, based on verified information and are appropriate to the context and purpose for which they are used.

Enablement/Technology Literacy

Users are enabled to use approved generative AI and be supported by mandatory education, training, and collaborations that promote participation and opportunity.

Guidelines

Users must use only approved tools when engaging with generative AI in the City's technology environment. This ensures compliance with corporate policies, privacy legislation, and information security standards. All generative AI tools, including those freely available online, are subject to the same policies, laws, and technology architecture review as any other technology solution.

Before Using Generative AI:

1. **Review and understand the City of Toronto policies, guidelines, and frameworks related to information management, privacy, and security.**

   Users must manage and protect City information in compliance with applicable legislation, by-laws, and policies (see listing above).

2. **Do not use Generative AI to avoid meeting your professional obligations as a City employee. Ensure the Generative AI tool you use is approved by the City for your intended use.**

Staff must use generative AI in a manner that upholds the City's reputation, reflects the professionalism of the Toronto Public Service, and maintains public trust in municipal government.

Generative AI tools are not a substitute for original research, critical thinking, or professional judgment. Rather, they are intended to support, not replace, these processes. Staff are responsible for ensuring that generative AI tools are used ethically, responsibly, and in alignment with City values.
Before using generative AI tools, staff should review the Toronto Public Service Bylaw, the Acceptable Use Policy, and other City Policies, as appropriate.

When Using Generative AI:

3. **Ensure that you are not sharing or submitting any Personal Information, Personal Health Information, or Information that is Exempt from Review (see details below) when using generative AI.**

"Personal Information" has the meaning given to it in MFIPPA and refers to recorded information about an identifiable individual. Examples of Personal Information include:

- home address, personal email address, home telephone number, identification numbers e.g. Social Insurance Number or employee number
- correspondence between the individual and City that directly or indirectly relates to private or confidential matters. Examples of these correspondences could occur in various settings, such as through personal emails, forms, and telephone interviews
- ethnic origin, religion, age, gender, sexual orientation, marital status
- educational, medical, criminal history, employment history, or personal financial transactions
- the individual's name if it appears with other Personal Information relating to the individual or where the disclosure of the name would reveal other Personal Information.

"Personal Health Information" has the meaning given to it in PHIPA and refers to identifying information about an individual in oral or recorded form that is in the custody or control of a Health Information Custodian, if the information any of the following, amongst other items;

- Relates to an individual's physical or mental health, including their family health history
- Pertains to the provision of health care to the individual, including identifying their health care provider
- Is a plan of service under the Home Care and Community Services Act, 1994

- Details home and community care services to be delivered under the Connecting Care Act, 2019
- Concerns payments, eligibility, or coverage for health care
- Involves body part or bodily substance donations, or test results derived from them
- Is the individual's health number
- Identifies the individual's substitute decision-maker

Information that is Exempt from Review refers to records or portions of records that are exempt from disclosure under Part I of MFIPPA. These exemptions are established to protect sensitive information that, if disclosed, could cause harm to individuals, third parties, or public interests.

- Common exemptions include records related to law enforcement, personal privacy, solicitor-client privilege, third-party commercial information, and confidential deliberations.

The Information Collection Policy Personal Information Collection, Use and Disclosure Guideline communicate the City's responsibilities for the collection and use of Personal Information.

For more information on information classifications, review the Information Protection Classification Standard.

4. **Do not use Generative AI to share or submit information that risks re-identification of any data that has been protected.**

   Even if data appears anonymized, it can be re-identified when combined with other inputs that have been submitted to the system.

   Example:
   You input a dataset showing de-identified service requests by postal code and date. Later, you ask the system to summarize complaint patterns about a specific address or person. The AI tool could potentially correlate the postal code and dates with previous inputs, unintentionally revealing identifiable details about a resident, even if you never included a name.

   To avoid this, only use generative AI tools with approved, secure datasets and ensure they are formally authorized for such use.

After Using Generative AI:

5. **Review the content generated by Generative AI to ensure that it is accurate, free of errors, and is not plagiarized.**

Users are responsible for reviewing and verifying all content produced by Generative AI. Outputs should not be assumed to be accurate, complete, or appropriate without careful human oversight.

Generative AI tools may produce content without proper attribution or permission from original sources, and in some cases, may fabricate information or references. As such, human review is essential to ensure the reliability, integrity, and ethical use of AI-generated content.

When evaluating outputs, consider the following best practices:

- Ensure the content is logical, coherent, and meaningfully explainable.
- Verify that any sources cited are accurate and legitimate.
- Assess whether the output could inadvertently reinforce bias or cause harm to equity-deserving groups, particularly based on race, gender, class, disability, sexual orientation, or other protected characteristics. Refer to the City's Human Rights and Anti-Harassment/Discrimination Policy for more guidance.
- Confirm that the content does not replicate or closely resemble copyright-protected materials, and that intellectual property rights are respected.

6. **Don't use content from Generative AI to influence a significant or impactful decision**.

Generative AI tools are not designed to make decisions. These tools only help summarize or synthesize information. They can be prone to inaccuracies, bias, and misunderstanding.

7. **Ensure that proper record managements processes and procedures are observed.**

Any content generated using generative AI for City business purposes (e.g., policy drafts, reports, or recommendations) must be treated as part of an official record and managed in accordance with the City's information management and recordkeeping policies and legislative obligations.

That means users are responsible for:

- Storing records in official recordkeeping repositories authorized for use by your division;
- Capturing version history when AI drafts are refined into final content;
- Not relying on generative AI tools to store records.

When determining whether an output constitutes a record, consider whether it relates to decisions, actions, or policies in the course of your business operations.

If you have any questions regarding records retention requirements, contact Corporate Information Management Services (CIMS), City Clerk's Office, at infomgmt@toronto.ca.

Frequently Asked Questions (FAQs)

What Generative AI tools are approved for use?

The City of Toronto's officially supported and approved generative AI tool is Microsoft Chat, which is integrated within the Microsoft 365 environment. This tool is securely deployed under the City's corporate license and is designed to support productivity, communication, and information management within commonly used applications such as Word, Outlook, Excel, and Teams.

Staff should use Microsoft Chat for generative AI tasks related to City work. Using this tool ensures alignment with the City's security protocols, information governance policies, and data protection standards. It also helps maintain transparency and accountability, as MS Chat operates within authenticated, managed environments and does not access external data sources beyond the Microsoft ecosystem.

The use of any **non-approved generative AI tools**, including public-facing tools, is not permitted.

How accurate is the content generated by Generative AI?

Content produced by Generative AI should not be assumed to be reliable, or a source of truth. Generative AI produces content by leveraging large amounts of data collected from a wide range of sources available on the open internet, text-based sources, and from users who have uploaded data and information to its database. As this data may not always be accurate or reliable, content produced by Generative AI could be false or misleading. This in turn could lead to the spread of misinformation and 'fake news'.

You must double-check the information provided by Generative AI to verify for accuracy, completeness, and logical coherence before using it. You are responsible for any AI-generated data and information that you use to support your work, which must be clearly cited by you in your work product as being AI-generated.

What are intellectual property considerations for content generated by Generative AI?

Generative AI may be trained using data that has been sourced without regard for copyright or licensing permissions and produce content that is identical or substantially similar to copyright protected material. Staff must perform due diligence to ensure that no copyrighted material is published by the City without proper attribution or authorization.

Are there ethical issues with Generative AI content?

Generative AI systems utilize vast amounts of data, some of which may contain errors or bias across race, sex, gender identity, ability, and many other factors. Staff should review any content generated by AI to ensure that instances of bias or discrimination, as well as potentially offensive or harmful material, are changed or removed.

Common Use Cases – Risk Management

Use this section to guide your thinking when applying generative AI. Always consider the risks before you hit "send," "share," or "save".

1. Drafting Emails or Memos

Scenario: You're drafting a quick email to a colleague summarizing a meeting or asking for information and you use generative AI to generate a first draft.

How It Can Go Wrong: The Generative AI system pulls inaccurate details about the meeting; It introduces assumptions or language that sounds overly confident, formal, or impersonal; You don't verify the tone or content, and it creates confusion or reputational risk.

Ask Yourself: Have I checked the facts and context in this draft? Does this sound like me? Is the tone appropriate for the recipient? Would I be comfortable if this message were forwarded outside the organization?

2. Summarizing a Document or Meeting

Scenario: You use generative AI to summarize a report, policy, or meeting notes for quick consumption.

How It Can Go Wrong: Key points are missed or misrepresented; Confidential or sensitive info is included inappropriately; The summary may seem objective but is subtly biased based on how the AI interpreted it.

Ask Yourself: Do I understand the original content enough to assess this summary? Are there nuances missing that a human would catch? Have I ensured this doesn't include or leak sensitive or personal information?

3. Brainstorming or Ideation

Scenario: You ask generative AI to help generate ideas for a public engagement strategy, naming a new service, or drafting policy principles.

How It Can Go Wrong: Outputs can reflect bias, stereotypes, or outdated norms; Generated ideas might conflict with organizational values or community standards; You might over-rely on AI instead of bringing in diverse, human perspectives.

Ask Yourself: Am I treating this as a starting point, not a final product? Have I considered equity, accessibility, and inclusivity in these ideas? Who else needs to weigh in before this moves forward?

4. Research or Information Gathering

Scenario: You ask generative AI to help explain a technical topic or summarize a regulation.

How It Can Go Wrong: AI may hallucinate or invent sources or facts; It may not reflect the most current or authoritative information; You risk embedding misinformation in your work if not verified.

Ask Yourself: Have I double-checked this against trusted sources (e.g., internal docs, government websites)? Is this topic high-stakes enough that a subject-matter expert should be involved? Could this cause harm if the information is wrong?

Contact Information

For general questions and concerns regarding the use of Generative AI:
Reach out to the Technology Services Division (TSD) at ai.support@toronto.ca.

For questions about the security of Generative AI, or to report the event of a cyber incident or issue:
Reach out to the Office of the Chief Information Security Officer at ciso@toronto.ca.

For questions about the use of City records, data, and information within Generative AI:
Reach out to the Corporate Information Strategy & Policy Unit at infomgmt@toronto.ca.

In the event of a suspected or confirmed privacy breach with Generative AI:
Immediately notify the City Clerk's Office at privacy@toronto.ca.

Last date amended 08/28/2025: Contact Information updated to include public-facing City contact emails.