



information
and privacy
commission
new south wales

IPC Generative Artificial Intelligence Policy

November 2025



Contents

1. Purpose	3
2. Objectives	3
3. Definitions	3
4. Key Risks.....	4
5. Prohibited uses of open-access generative AI tools	5
6. Pre-authorised uses of open-access generative AI tools	5
7. Protocols for the use of open-access generative AI tools	7
8. Records Management.....	7
9. Related Documents.....	8
Appendix A – Risk Assessment Template.....	10

1. Purpose

This policy governs when and how the IPC may (and may not) use generative artificial intelligence (AI) technologies, including in the performance of the Information Commissioner and the Privacy Commissioners' statutory functions. This policy applies to all staff including contractors and consultants engaged to perform work for, or on behalf of, the IPC. It is based on one developed by the NSW Ombudsman: [Generative Artificial Intelligence - Use by NSW Ombudsman officers](#).

2. Objectives

The IPC's Strategic Plan has an objective that IPC has expertise in how data-driven and automated decision-making and other technology is being implemented across all sectors and its impact on privacy and transparency. Through implementation of this Policy, IPC has a framework for reviewing and analysing artificial intelligence technologies being implemented in the NSW public sector.

Consistent with the [NSW Government Artificial Intelligence \(AI\) Ethics Policy](#), the IPC adopts the following overarching objectives for the purpose of developing and implementing this policy:

- Community benefit – AI should deliver the best outcome for the public and key insights into decision-making. Use of AI should always be considered against other analysis and policy tools.
- Accountability – decision-making remains the responsibility of staff.
- Transparency – the public will be informed and have access to an efficient and transparent review mechanism if there are questions about the use of data or AI-informed outcomes.
- Fairness – the use of AI will include safeguards to manage data bias or data quality risks.
- Privacy and security – AI will include the highest levels of assurance. The public's information must be used safely and securely, and in a way that is consistent with privacy, data sharing and information access requirements.

3. Definitions

Term	Definition
Artificial Intelligence (AI)	Refers to computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages. Many such systems are trained using Machine Learning and have the capability to learn from or be trained by data it works with.
Generative AI	Refers to AI algorithms that can create (generate) new content in the form of text, images, audio or multimedia based on natural language prompts provided by humans. Generative AI may be: <ul style="list-style-type: none">• Open-access generative AI tools, which is any generative AI tool hosted on the internet that the public can use freely or by paid subscription. Some current examples of publicly available generative AI tools are ChatGPT and Microsoft Copilot. For the purpose of this policy, the definition of open-access generative AI tools does not include search engines or other similar sites that may use generative AI functionality (such as predictive text and autocomplete suggestions).• Closed-access generative AI tools, which refers to a generative AI tool for internal use only by IPC staff. System and Corporate Services will maintain a list of all approved closed-access generative AI tools.

Term	Definition
Machine Learning	Refers to a process involving a computer program that learns and adapts – either in a structured or unstructured way – by using algorithms and statistical models to analyse and draw inferences from patterns in data.
Prompts	Refer to the questions or instructions that staff enter into the Generative AI tool to generate a response.
Risk Assessment	Risk assessment using the criteria set out at Appendix A.

4. Key Risks

The following table summarises some of the key risks of open-access generative AI, and how those risks are managed by this policy. The risks associated with closed-access generative AI are covered by the Risk Assessment Template:

Risk	Mitigation/Control
Unauthorised disclosure of sensitive, personal or other confidential information by its inclusion in a 'prompt' given to the generative AI tool	<p>This policy prohibits the inclusion of any personal, sensitive or other confidential information in prompts (section 4).</p> <p>This policy prohibits certain uses that have a high risk of inclusion of sensitive or other confidential information being included in prompts (such as decision making about complaints and reviews) (section 5).</p> <p>This policy requires approval for other uses that have a significant risk of inclusion of sensitive, personal or other confidential information being included in prompts (such as editing or report writing) (section 6).</p>
Breach of Copyright, Breach of Moral Rights to attribution and integrity of work, risk of plagiarism	<p>This policy prohibits the use of open-access generative AI to make decisions and perform work involved in writing public and other reports (section 5). Use of open-access generative AI for certain non-confidential content generation or content review is permissible (section 6) subject to restrictions and caveats.</p> <p>This policy requires all information be fact checked, which includes verifying sources and that the content is accurate (section 6).</p>
Incorrect, inaccurate or biased information, including 'hallucination'	<p>This policy requires ownership and accountability for quality content by fact-checking and citing AI generated content (section 5).</p>
Malicious use and deception	<p>This policy is to be read and complied with in conjunction with the IPC Code of Conduct.</p>

5. Prohibited uses of open-access generative AI tools

Staff must not use open-access generative AI to perform, provide input into performing, or otherwise assist in performing, the functions of:

- making any decision in respect of any statutory function, including:
 - decisions about the jurisdiction of the Information Commissioner or the Privacy Commissioner
 - assessing and determining reviews or complaints under the [Government Information \(Information Commissioner\) Act 2009](#), [Government Information \(Public Access\) Act 2009](#), [Health Records and Information Privacy Act 2002](#) or [Privacy and Personal Information Protection Act 1998](#)
 - assessing data breach notifications under the [Privacy and Personal Information Protection Act 1998](#)
 - decisions about what action to take in respect of an enquiry or complaint, including whether to undertake preliminary inquiries or investigation
 - assessing disclosures under the [Public Interest Disclosures Act 2022](#)
 - decisions about whether it is necessary or appropriate to refer a matter to another body
 - decisions about the content of public and other reports, including how the factual findings and recommendations are expressed, and setting out what evidence is relied on.
- searching IPC work email accounts or SharePoint directories – e.g., connecting Outlook and SharePoint accounts to Microsoft Copilot
- forming any legal opinion – excluding the summarising of case law, legislation and standards which is allowed under section 6 below
- other language translations or interpretations of communications and documents that are relied on in exercising statutory or corporate functions¹
- corporate functions, such as data analytics or decision-making in information access requests and recruitment activities, except where the data analysed or used for the decision is deidentified and accessible to the general public – e.g., published statistical data, public NSW government policy and procedure documents.

Although staff are prohibited from using open-access generative AI tools for these functions, it may not preclude staff from using closed-access generative AI tools for these functions, subject to section 7.

6. Pre-authorised uses of open-access generative AI tools

Staff may use open-access generative AI to perform the tasks listed below, without seeking Executive approval, as long as they do not involve the ingestion of and use of Personal Information and are not being used for the matters set out in section 5.

Tasks

Assistance from open-access generative AI tools may only be used for tasks of the following general nature:

¹ This includes website services such as Google Translate.

- Ideation – brainstorming, generating or identifying counter-arguments
- Writing – copy-editing and proofing text, synthesising or summarising text from information that is not sensitive or confidential, generating social media content
- Research – identifying and summarising relevant publicly accessible literature, summarising and explaining complex concepts
- Data analytics – identifying, analysing and projecting trends, and providing conclusions from publicly accessible data
- Information organisation – summarising publicly accessible information, organising and formatting publicly accessible information, preparing chronologies from publicly accessible information
- Presentation and training materials – developing and synthesising training text and handouts, generating unique or specific content such as images, video and audio to match or enhance training materials, designing training and presentations materials.

Approval is required for tasks not listed above. This list is also subject to the restrictions and caveats below.

Prompts

Staff must not include any confidential, personal or sensitive information as prompts (or as part of a prompt) to an open-access generative AI tools. This includes:

- names or any other identifying information about any person (unless the information is public information about a public person relating to information in the public domain e.g. referring to a Minister's public media statement, a Court's published judgments or about an historic figure)
- any health information (within the meaning of the *Health Records and Information Privacy Act 2002*) about any person
- information that has been classified in IPC's electronic document management system (TRIM) as Sensitive (NSW Cabinet, Legal, Law Enforcement, Health Information, Personal, NSW Government).

Ownership and accountability for product and quality

All AI-generated content must be:

- **Critically evaluated.** Staff must assess whether the content includes any unfounded assumptions or biases. Furthermore, any ideas provided by the generative AI tool must be developed and edited such that they are substantially the staff member's own work – i.e., more than 50%. Staff must exercise a high degree of care before relying on AI-generated content that summarises technical or complex material, such as legislation and case law.
- **Fact-checked.** Staff must ensure all information is accurate and up to date. This includes verifying and attributing to sources, checking statistics, and ensuring that any claims made in the content are supported by evidence. Any information relied upon that is obtained from a generative AI tool should be fact checked against a verifiable source.
- **Edited.** Staff must edit all content to ensure that it is logical, well-written, structured in a logical manner and is appropriate for the intended audience. Any text taken from or based on text provided by the generative AI tool must be identified and rewritten so that it is substantially – more than 50% – the staff member's own work to avoid unwitting plagiarism or breach of copyright.
- **Proofed.** Staff must check all content for spelling and grammar, and consistency with IPC style, before it is published or shared.

Transparency

Staff must cite in internal documents, or internal versions of documents that will become external documents, when open-access generative AI content has been used to assist with a task, whether in external or internal documents. An example citing reference is as follows:

* A generative AI system was used to assist in copy editing and proof-reading parts of this document. No sensitive information was provided to the system, and the document has been reviewed and approved by the author for accuracy and appropriateness prior to signing.

A citation does not need to be included in external documents.

7. Protocols for the use of open-access generative AI tools

Check if intended use of open-access generative AI tool aligns with the requirements of this Policy. Some open-access generative AI tools may be available to staff under our shared service arrangements with the Department of Customer Services but not align with the requirements of this Policy.

Staff must check that any intended use of an open-access Generative AI tool is:

- not prohibited under section 5, and
- pre-authorised under section 6.

Approved uses following risk assessment

If an intended use of an open-access generative AI tool is prohibited by section 5 and/or not pre-authorised by section 6 of this Policy, approval may be sought and granted by the Executive for that use, following the submission of a completed risk assessment using the template at Appendix A. In completing the risk assessment please have regard to [DCS's Security Guidelines for AI Platforms](#).

Training

Prior to using open-access generative AI tools, staff must undertake training on effective and safe usage in accordance with this policy from a pre-approved list of training. Completion of training will be recorded by the Executive Assistant and Office Administrator, who is responsible for maintaining a list of staff training.

Login

If open-access generative AI tools are to be used for any work-related purposes in accordance with this Policy, staff must only use an account that has been created with their work email address. However, staff must not use the same password that is used to log into their work account (e.g. computer login). Staff must not use any other open-access generative AI account (for example, an account created using their personal email address) for any work-related use.

8. Records Management

Staff must retain documentation each time generative AI is used, including screenshots of all prompts and outputs generated, and store the documentation in accordance with the IPC Records Management Policy. Documentation related to authorised use of generative AI tools (including the appropriately completed risk assessment) should be sent to the Manager, Systems and Corporate Services, who is responsible for registering authorised uses of generative AI.

Any non-compliance with this policy should be reported to Managers/Directors. Breaches will be evaluated by the Director, Corporate Services and Business Improvement, for further action under the IPC Code of Conduct and Data Breach Policy. Breaches of this policy by staff may result in disciplinary action in accordance with the IPC Code of Conduct.

Roles and Responsibilities

Each of the following roles has specific assigned responsibilities under this policy:

Role	Responsibilities
All staff, contractors and consultants	<ul style="list-style-type: none">• Comply with the requirements of this Policy• Create and capture records of each time generative AI is used and save these in IPC's Electronic Documents Management System (TRIM / Content Manager).• Ensure contractors, consultants and agency staff engaged to perform work for or on behalf of the IPC are made aware of this policy and its requirements and, if required, ensure that appropriate contractual provisions are included in procurement contracts to apply this policy to their work• Check during the procurement stage whether generative AI tools are intended for use during the engagement and ensure any approvals required are obtained
Executive	<ul style="list-style-type: none">• The Executive is responsible for the implementation of this policy in its entirety, including directing all staff in their teams to follow this policy.• Approve applications for intended use of open-access generative AI tools that are either prohibited by section 4 and/or not pre-authorised by section 5 of this Policy.

9. Related Documents

- [DCS's Security Guidelines for AI Platforms](#)
- [Government Information \(Information Commissioner\) Act 2009](#)
- [Government Information \(Public Access\) Act 2009](#)
- [Health Records and Information Privacy Act 2002](#)
- [IPC Code of Conduct](#)
- [IPC Data Breach Policy](#)
- [Privacy and Personal Information Protection Act 1998](#)
- [Public Interest Disclosures Act 2022](#)

Document Information

Identifier / Title:	IPC Generative Artificial Intelligence Policy
Business Unit:	Systems and Corporate Services
Author:	Director Corporate Services and Business Improvement
Approver:	Chief Executive Officer
Date of Approval:	November 2025
Next Review Date:	November 2027
EDRMS File Reference:	D25/015636/DJ
Key Words:	Artificial Intelligence, Generative AI, Open-access generative AI tools, risk management.

Document History

Version	Date	Reason for Amendment
1.0	November 2025	First publication

Appendix A – Risk Assessment Template

Risk Assessment of Generative AI Tool

Completed templates are to be emailed to the Director Corporate Services and Business Improvement for consideration by the Executive Team.

General Assessment

Assessment Item/Question	Response
Assessment Completed by	Name and Position:
Name of Tool	
Name of Tool's Developer/Vendor	
Website/URL Link	
Is the Tool Open or Closed Access?	
Description of use-case for the Tool including benefits to IPC staff, NSW Government Agencies and the public	
Would the Generative AI tool/use case constitute 'operational AI'? (Operational AI systems are those that have a real-world effect. The purpose is to generate an action, either prompting a human to act, or the system acting by itself. Operational AI systems often work in real time [or near real time] using a live environment for their source data. Non-operational AI systems do not use a live environment for their source data. Most frequently, they produce analysis and insight from historical data.)	
Datasets to be used with the Tool and considerations for data governance and de-identification of personal, confidential or sensitive data	
Will the use-case use real-time or near real-time data to make recommendations for staff to act on in real-time or near real-time or take actions itself in real-time or near real-time?	
Does the use-case and the use of data align with relevant legislation?	
Have you completed a privacy impact assessment (either third party or self-assessed)?	
Does the use-case adhere to the requirements in the NSW Cyber Security Policy?	

Assessment Item/Question	Response
Has consultation occurred with stakeholders?	
Is there an easy and cost-effective way for people to appeal a decision that has been informed by your use-case?	
Does the use-case allow for transparent explanation of the factors leading to the AI decision or insight?	
Who is responsible for: use of the AI insights and decisions, policy/outcomes associated with the use-case, monitoring the performance of the use-case and data governance?	
Have you established clear processes to: intervene if a relevant stakeholder finds concerns with insights or decisions and to ensure you do not get overconfident or over reliant on the use-case?	
If you are procuring all or part of a use-case, have you satisfied the requirements for: transparency, privacy and security, fairness and accountability as defined in the NSW AI Assessment Framework?	

Risk Assessment

Assess use-case of Tool against the following risks. Risks to be assessed using the risk matrix detailed in IPC's Enterprise Risk Management Policy and Framework (D25/004858/DJ)

Risk	Likelihood	Consequence	Mitigation/Control
Physical or environmental harms from using the Tool			
Insufficient experienced human oversight of the use-case.			
Using incomplete or inaccurate data			
The potential to cause discrimination or unfair treatment from unintended bias.			
Unauthorised use of health, confidential or sensitive personal information			

Risk	Likelihood	Consequence	Mitigation/Control
Unintended identification or misidentification of an individual			
Incorrect advice or guidance			
The use-case is a single point of failure for the delivery of a service or policy			
Incomplete documentation of use-case design, or implementation, or operation			
No or limited access to model's internal workings or source code ("Black Box")			
Being unable to explain the output of a complex model			
A member of the public being unaware that they are interacting with a use case			
No or limited mechanisms to record use-case decision history			

Assessment Summary: (Include comments on the impact of the risk assessment on implementation of the proposed use-case)

For Executive Team Use

Approval:

Conditions: