



Australian Government

**Department of Infrastructure, Transport,
Regional Development, Communications and the Arts**

Chief Operating Officer Group / Data, Research, Strategy and Net Zero Division / Data, Systems and Information Management Branch

Generative Artificial Intelligence Policy

Released under the FOI Act 1982 by the Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts

Document Change History

Version	Date	Changed By	Description of Change
0.1	03/2024	Data Strategy and Policy	Initial Draft
0.2	03/2024	Data Strategy and Policy	Updates after review
0.3	04/2024	Data Strategy and Policy	Integration of feedback from targeted consultation
0.4	05/2024	Data Strategy and Policy	Integration of feedback from Risk and Integrity Committee
0.5	07/2024	Data Strategy and Policy	Changes to reflect: the CDO's role as the department's accountable officer for AI; and comments from survey circulated to the department's AI Community of Interest and Data Community of Practice
0.6	08/2024	Data Strategy and Policy	Integration of feedback from Enabling Committee

Document Owner

Name	Position
s22(1)(a)(ii)	Director, Data Strategy and Policy Section

Document Approval

Name	Position
Andreas Bleich	Assistant Secretary, Data, Systems and Information Management Branch
Ian Porter	First Assistant Secretary, Data, Research, Strategy and Net Zero Division
Enabling Committee	The Enabling Committee supports the Secretary and the ELT by setting and monitoring business management matters for the department. The Enabling Committee operationalises the strategic direction set by the Secretary and/or Executive Leadership Team.
Executive Leadership Team	The ELT consists of all Deputy Secretaries and the Secretary. The ELT is focussed on significant strategic issues that affect the department.

Contents

Contents	3
1. Purpose.....	4
1.1 . Why do we need a generative AI policy?	4
1.2 What is Generative AI?	5
1.3 Limitations of Generative AI	5
1.4 Ethical Considerations with Generative AI	6
2. Policy	6
2.1 General Principles.....	6
2.2 Requirements	8
3. Roles and Responsibilities	9
4. Governance	11
Appendix A: Related legislation, regulations and policy frameworks	12
Appendix B: Tactical guidance and use cases.....	13
Appendix B: Generative AI terminology	15

Released under the FOI Act 1982 by the Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts

1. Purpose

The purpose of this document is to establish the approach for the deployment and use of generative artificial intelligence (AI) in the Department of Infrastructure, Transport, Regional Development, Communications and the Arts (the department). It provides an overview of generative AI and its challenges, puts forward principles for using it responsibly, and outlines mandatory requirements. This policy applies to all department employees and contractors (collectively referred to as users throughout the rest of this policy).

This policy is created to assure that ethical, privacy, risk, and security implications are appropriately considered when using generative AI for departmental work. The principles, guidance and use cases that follow are informed by the Digital Transformation Agency's (DTA) [Interim guidance on government use of public generative AI tools](#) which draw from the Australian Government's [AI Ethics Principles](#). The DTA guidance may be supplemented with other policies including a risk framework. Users are encouraged to use this policy in conjunction with the DTA resources. The primary principle for use of AI in the department is that AI tools should be used to support human-decision making and not replace it.

Above all, users are encouraged to apply these two **golden rules**.¹

- Users should be able to explain, justify and take ownership of your advice and decisions.
- Know that any information you input into public generative AI tools² can become public and don't input classified, personal or otherwise sensitive information.

1.1 . Why do we need a generative AI policy?

The department recognises that clear guidance specific to the uses and risks that AI technology presents is a priority. This guidance enables users to assure transparency, accountability and fairness when using generative AI to support decision making, policy development, regulation and program delivery. This policy and the department's [Generative AI Hub](#) is aimed at giving users the information they need to stay up to date with rapidly evolving regulations and best practices, and emerging applications.

Generative AI holds significant potential to improve the delivery of public services in Australia. Opportunities exist across the spectrum of activities carried out by the Australian Public Service (APS): from automating administrative or time-consuming manual processes; suggesting content for policy advice; to classifying and organising data and generating computer program code.

Realising these benefits will require the public service to develop, adopt and use generative AI in ways that are trustworthy, and transparent and defensible³. If done well, using AI could see the APS realise substantial efficiency gains, deliver better staff and customer experiences and improve outcomes for the whole community. If done poorly, the APS will fail to capture the benefits of AI and erode trust such that future adoption or use of AI will not be accepted by the community, affecting service delivery innovation by the APS more broadly.

¹ Interim guidance on government use of public generative AI tools, Digital Transformation Agency [Interim guidance on government use of public generative AI tools - November 2023 | aga \(digital.gov.au\)](#)

² Public generative AI tools are widely available third-party AI platforms, tools or software (whether accessed through a browser, phone or application), such as ChatGPT and MidJourney. These have not been security risk assessed by the department and approved for use with classified or sensitive data. This also includes the enterprise version of Microsoft Copilot (previously referred to as Bing Chat Enterprise), where some agencies have enabled it for staff use.

³ How might artificial intelligence affect the trustworthiness of public service delivery?, Long-term Insights Briefing, Department of the Prime Minister and Cabinet [Summary - How might artificial intelligence affect the trustworthiness of public service delivery? | PM&C \(pmc.gov.au\)](#)

1.2 What is Generative AI?

The term 'AI' is often used to describe several technologies used in machine-assisted decision-making applications including machine learning, robotic process automation (RPA), natural language processing, deep learning (neural networks), computer vision and robotics.

AI systems embrace a family of technologies that can bring together scalable computing power, storage and networking, connected devices and interfaces, and data.

Generative AI refers to systems that use data models to generate new examples of content such as text, images, audio, code and other data modalities. Generative AI applications are trained on data and can approximate human generated content from prompts, including prompts that are limited or non-specific. Techniques such as human supervision and alignment can also be applied to improve the outputs, and users can provide feedback or modify their prompt to refine the response.

Generative AI systems use patterns in language and concepts to respond to prompts. These systems do not 'understand' prompts or the responses they generate. This reliance on data patterns means they can replicate the form and style of human-like responses, but factual accuracy is not guaranteed.

Generative AI technology is rapidly being incorporated into common online tools as standalone systems or embedded within other applications. High-level practical applications in use today include:

- **Written content augmentation and creation:** Producing a "draft" output of text in a desired style and length.
- **Question answering and discovery:** Generating answers to input, based on training data and prompt information. Note that many generative AI systems are currently unreliable for this purpose and prone to returning factually inaccurate responses.
- **Tone:** Text manipulation, to soften language or professionalise text.
- **Summarisation:** Offering shortened versions of conversations, articles, emails and webpages.
- **Simplification:** Breaking down titles, creating outlines and extracting key content.
- **Classification of content for specific use cases:** Sorting by sentiment, topic, etc.
- **Image and video generation:** Creating and modifying images and video based on text or image prompts.
- **Software coding:** Code generation, translation, explanation and verification.
- **Creating synthetic data:** algorithm-generated data that mimics the characteristics of real-world data and typically used when real data isn't available or when it's kept private due to data protection laws.

Examples of use cases for generative AI applications are at **Appendix A** and a glossary of generative AI terminology is provided at **Appendix B**.

1.3 Limitations of Generative AI

The current generation of large language models (LLMs) are prone to "hallucinations", generating fictitious information, presented as factual or accurate. This can include citations, publications, biographical information, and other information commonly used in research and academic papers. This is likely to improve as the capability develops, however outputs from generative AI need to be carefully reviewed and fact checked.

Due to their fundamental structure, content from many generative AI models can be difficult to reproduce consistently. This is particularly problematic in research where reproducibility is a cornerstone of establishing credibility. In addition, the outputs of the models are constrained by the prompt's users enter and by the training data, which may be difficult to assess if it is not publicly accessible.

Generative AI models are not databases of knowledge but rather an attempt to synthesize and reproduce the information they have been trained on. Therefore, AI is only as good as the available data. The popular adage, “garbage in, garbage out”, applies to Generative AI. When the consensus data is wrong or outdated, then most likely the generated content will reflect these same errors.

This difficulty understanding context when presented with new information or scenarios outside of training parameters, means that generative AI cannot draw conclusions or make decisions based on complex situations — something that only humans can do at present. Furthermore, generative AI cannot replace human creativity completely as it lacks the ability to come up with novel ideas or recognize abstract concepts such as humour or irony — these things currently require a human touch.

1.4 Ethical Considerations with Generative AI

Ethical use of AI extends beyond compliance with the law as an activity may be lawful but unethical. The APS Values include that ‘the APS demonstrates leadership, is trustworthy, and acts with integrity, in all that it does’ (*Public Service Act 1999*, ss 10(2)).

To embody these values, it is essential that there are ethical considerations throughout all stages of the AI product lifecycle. This includes the process, information, training and data that it uses and the results. For further guidance users are encouraged to read and understand [Australia's AI Ethics Principles](#).

Some specific considerations relating to whether it is ethically appropriate to use a particular AI product include:

- Whether individuals to whom the data relates consent for it to be used in a particular way (including to be used as a secondary purpose).
- Whether the AI product, and training data, infringes intellectual property rights (such as the use of copyrighted information without consent).
- Where the data used to train the AI, product comes from and how the AI product has been trained.
- Whether an AI product might exploit or use information that has been unlawfully published or disseminated.

2. Policy

2.1 General Principles

The following principles contain guidance prepared by the DTA⁴ to help users adhere to Australia's AI Ethics Principles⁵ when using generative AI tools, and is organised into sub-sections as follows:

a) Accountability

Accountability is one of Australia's AI Ethics Principles and one of the five APS Values. To remain open and accountable to the Australian community, you should be able to explain, justify and take ownership of your advice and decisions.

Generative AI tools must not be the final decision-maker on government advice or services.

Users may consider using generative AI tools to brainstorm options, generate code or draft content. However, these tools can produce a convincing but inaccurate response to a query, sometimes referred

⁴ Interim guidance on government use of public generative AI tools, Digital Transformation Agency [Interim guidance on government use of public generative AI tools - November 2023 | aga \(digital.gov.au\)](#)

⁵ Australia's AI Ethics Principles, <https://www.industry.gov.au/publications/australias-artificial-intelligence-ethics-framework/australias-ai-ethics-principles> (Accessed 13/05/2024)

to as 'hallucination'. A human with the appropriate skills, knowledge or experience should review output before using it (especially before entering any coding outputs into government systems).

Finally, to enable you to critically analyse the outputs produced by generative AI tools and to understand the limitations of the technology users must complete a minimum level of training, when available⁶, before using generative AI tools.

b) Transparency and explainability

It should be clear when generative AI tools are being used by government to inform activities. Users could consider including markings in briefings and official communications indicating if generative AI was used to generate any of the information.

Building on the accountability principle, you should critically examine outputs from generative AI tools to assure your advice and decisions reflect consideration of all relevant information and do not incorporate irrelevant or inaccurate information.

Remember, you should be able to explain, justify and take ownership of your advice and decisions.

c) Privacy protection and security

Inputs into public generative AI tools should not include or reveal classified, personal or otherwise sensitive information. All activities need to align with legislation and policies relating to information and data (for example the Privacy Act 1988, and the Protective Security Policy Framework).

Government information must only be entered into public generative AI tools if it has already been made public or would be acceptable to be made public. Employees determining that the information in question is suitable for public release must have the appropriate organisational delegation to do so.

Classified or sensitive information must not be entered into these tools under any circumstances.

You should not enter information that would allow public generative AI tools to extrapolate classified or sensitive information based on the aggregation of content you have entered over time.

Any data entered into public generative AI tools is stored externally to government and we do not know who has access to it. Where available, you should disable any settings or permissions which save chat history.

d) Fairness and human-centred values

Generative AI tools are typically trained on broad sets of data that may contain bias. Bias can arise in data where it is incomplete, unrepresentative or reflects societal prejudices.

Generative AI tools may reproduce biases present in the training data, which could lead to misleading or unfair outputs. Bias in outputs may disproportionately impact some groups, such as First Nations people, people with disability, LGBTIQ+ communities and multicultural communities.

Before using AI-generated outputs, you should consider whether you have a process in place to assure that outcomes are fair and meet community expectations. For example, you could include representatives of relevant communities (with appropriate skills, knowledge or experience) in decision-making on use of data and outputs relating to those communities.

Remember, you should be able to explain, justify and take ownership of your advice and decisions.

⁶ On 17 August the Digital Transformation Agency released the *Policy for the responsible use of AI in Government*. The policy, which comes into effect 1 September 2024, urges departments to adopt a minimum level of training for AI for all staff within 6 months of the policy taking effect. DTA have advised that APS-wide guidance for providing staff with training on AI will be released soon.

e) Human, societal and environmental wellbeing

You should engage with generative AI tools to better understand the potential benefits and risks of the technology, and the role it can play in contributing to the strategic priorities of your agency and the Australian Government.

You should weigh the benefits and risks of a particular use of generative AI, and consider whether it is the right tool for the job. Traditional tools may be cheaper, safer or better suited to the task at hand.

You should only use generative AI tools in a manner consistent with the APS Values, Employment Principles and Code of Conduct for purposes that are consistent with improving the wellbeing of the Australian community. You should also consider intellectual property rights of third parties as well as broader copyright issues when using these tools and seek legal advice where necessary.

Consider whether your use of Indigenous data and generative AI outputs is consistent with the expectations of First Nations peoples, particularly around Indigenous [data sovereignty](#), [data governance](#) and [cultural and intellectual property](#).

2.2 Requirements

Users must comply with all applicable laws, regulations and policy frameworks, and departmental conditions and when using generative AI applications to assure the appropriate and reasonable protection of data. A list of relevant laws, regulations and policy frameworks is at **Appendix A**.

At the request of business areas, the Chief Information Officer (CIO), or their delegate, will assess and endorse AI tools and capabilities for departmental use. Business areas considering generative AI as a technological solution to business problems and processes are encouraged to contact the IT Front Door for advice on potential options for use and investment. Assistant Secretary approval is required for requests to be actioned, after endorsement by the First Assistant Secretary.

The department will maintain and publish a catalogue of all generative AI applications that are endorsed for departmental use, including any conditions on use of these applications. The catalogue will also provide a list of departmental uses cases for applications.


The use of generative AI tools on departmental devices is restricted to work related purposes in accordance with this policy, once an AI tool has been assessed and endorsed for use in the department. Users must complete a minimum level of training, when available⁷, before using generative AI tools.

All works published by the department that are created with the involvement of generative AI must include an accompanying annotation which includes the generative AI technology used and a description of how it was used to create the work.

⁷ Refer to footnote 6

3. Roles and Responsibilities

The diagram and table below detail the roles and responsibilities for generative AI within the department:



Roles	Responsibilities
Strategic	
Chief Data Officer (CDO)	- Accountable officer for AI in the department - Responsible for implementation of this policy to assure the safe and responsible use of generative AI in the department - Engagement in whole-of-government generative AI forums and processes to encourage APS-wide capability uplift over time. - Working with stakeholders to inform implementation of the generative AI policy and development of supporting guidance - Reviewing the generative AI policy to assure it remains up to date and fit for purpose - Reporting, as requested, on departmental use of generative AI, and whole of government initiatives as they relate to generative AI - Approving any exceptions to this Policy if required.

Released under the FOI Act 1982 by the Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts

Roles	Responsibilities
AI Governance Committee	Support and inform the CDO to assure the safe and responsible use of AI in the department.
Chief Information Officer (CIO)	- Promotes and communicates the opportunities and risks of using generative AI tools. - Authorises the use of any new generative AI tools including online services, software or websites requiring sign-in. - Provides strategic oversight of the use of department information assets for generative AI applications.
Chief Operating Officer (COO)	- Sets the overall vision and strategy for AI adoption and governance within the department. - As Chief Security Officer, has accountability for security risk related to the use of AI.
Chief Technical Officer (CTO)	- Manages assessment and endorsement of AI applications/capabilities. - Maintains and publishes a catalogue of all approved AI applications. - Assures the department's technical infrastructure is aligned with the department's overall AI approach.
Chief Information Security Officer (CISO)	- Assures alignment of AI applications with ICT security objectives. - Manages the security certification and accreditation of the departmental IT systems and databases, including AI security assessments and incident management. - Oversees the security of the department's business systems and databases
Chief People Officer (CPO)	- Provides support and guidance regarding the adjustment of job design, job roles and work level standards associated with the introduction of AI. - Provides support and guidance regarding change management and staff consultation requirements. - Central coordination of learning and development in support of AI - Provision of support services to staff and managers in adjusting to change.
Tactical	
SES Band 3 (Deputy Secretary) & 2 (First Assistant Secretary)	Assure acceptable use of AI within their respective Group/Division, in accordance with this Policy.
SES Band 1 (Assistant Secretary)	- Review and endorse business area requests for IT support to identify technology solutions to solve business problems, prior to approval from the CIO. - Assure acceptable use of AI within their respective branch, in accordance with this Policy. - Assure staff have appropriate training and knowledge of AI. - Consider the change impact of AI on job roles and job design, and adjust accordingly, in consultation with staff. - Engage with staff about the benefits, challenges and impact of AI in the workplace

Roles	Responsibilities
Assurance, Integrity, Risk and Governance Branch	The Assurance, Integrity, Risk and Governance Branch is responsible for supporting business areas, the CIO and CDO in their identification, documentation and escalation of relevant risks; and appropriate mitigation activities associated with generative AI products. This may include assistance with templates and risk processes, audits to assure compliance with the relevant legislation and departmental processes and procedures.
Operational	
Supervisors	- Assure acceptable use of AI within their respective team, in accordance with this Policy. - Assure the appropriate use and training of AI tools by their staff and awareness of related policies. - Consider the change impact of AI on job roles and job design, and adjust accordingly, in consultation with staff. - Engage with staff about the benefits, challenges and impact of AI in the workplace
Users (department employees, contractors, temporary workers and others involved in the work of the department)	- Assure acceptable use of AI, in accordance with this Policy and relevant legislation and policies relating to information and data. - Are responsible for all generative AI outputs used in artefacts. - Obtains SES Band 1 endorsement for business use of AI applications and lodges a request for IT to assess AI tools. - Reports through the IT Front Door any instances where this policy cannot be fully applied, for guidance on any potential alternatives. - Assure that appropriate training has been undertaken and to follow conditions of use.

4. Governance

The Chief Data Officer is responsible for overseeing the safe and responsible use of AI in the department. The CDO will assure the use of AI aligns with departmental and whole-of-government policies and strategic goals.

To assure a collaborative approach to the implementation and use of generative AI the CDO will seek advice and guidance through a dedicated committee of key stakeholders representing a broad cross-section of the department. The terms of reference for the committee will be made available to users through the Intranet. The CDO will report to the department's Executive Leadership Team on the progress of integrating AI as part of the department's business model and the extent to which the department approaches AI as a risk or as an opportunity.

The CDO is responsible for assuring this policy remains consistent with relevant APS guidance and reflective of challenges arising from the evolution of generative AI technologies. The CDO will finalise an evaluation of the policy's effectiveness within 12 months of it coming in to effect, and review the policy after three years.

Appendix A: Related legislation, regulations and policy frameworks.

APS conduct

[APS Code of Conduct](#)
[Public Service Act 1999](#)

Artificial Intelligence

[Australia's AI Ethics Principles](#)
[Engaging with Artificial Intelligence \(AI\) guidance](#)
[Interim guidance on government use of public generative AI tools](#)

Cyber and protective security

[2023-2030 Australian Cyber Security Strategy](#)
[Protective Security Policy Framework](#)
[Cyber Security Guidelines](#)
[DITRDCA Cyber Security Policy](#)

Data

[Framework for Governance of Indigenous Data](#)
[Foundational Four](#)
[Archives Act 1983](#)
[DITRDCA Recordkeeping Policy](#)
[Data-matching Act 1990](#)
[Data Availability and Transparency Act 2022](#)
[Notifiable Data Breaches scheme](#)

Privacy

[Privacy Act 1988](#)
[DITRDCA Privacy Policy](#)

Risk

[Commonwealth Risk Management Policy](#)

Intellectual Property

[Copyright Act 1968](#)
[Intellectual property principles for Commonwealth entities](#)
[Protocols for using First Nations Cultural and Intellectual Property in the Arts - Creative Australia](#)

Appendix B: Tactical guidance and use cases⁸

Tactical Guidance – do's and don'ts.

- Check whether your use of generative AI aligns with the department's generative AI policy, ICT obligations and linked policies.
- Do not distribute or click on any links provided or generated by public generative AI tools or bots. These links could lead to phishing sites or malware downloads. Only click on links from trusted sources.
- Treat with care any files generated by public generative AI tools or bots which have the potential to contain malicious code, such as macros in Microsoft Office files. These files must be considered as potentially malicious, and not interacted with or distributed to other people, until vetted and proven to be non-malicious.
- Where developing tools or processes incorporating generative AI (for example, to summarise information for a daily news update), be sure to continue to monitor performance against intended purpose. Generative AI models are updated from time to time and may learn from the data they interact with. This may affect performance and reliability of output.
- Align to relevant principles, guidelines and best practice issued by the Australian Government, including:
 - [Australia's AI Ethics Principles](#) - Department of Industry, Science and Resources
 - [Artificial Intelligence guidelines](#) - Australian Government Architecture
- You should report any instances where you are not able to fully apply this guidance to the department's relevant senior accountable officer (e.g. the Chief Data Officer see **Roles and Responsibilities**).

Use cases

Whilst AI applications require approval in the department, it is acknowledged that users may already be experimenting with AI. Below are high level use cases to guide responsible application of public tools, in conjunction with this policy. *These use cases do not cover built in tools or tools approved for use in the department.*

Generating 'first pass' content

Nick needs to develop a project plan and is wondering if he can use ChatGPT to create a baseline project plan that he can then improve upon.

What should Nick do?

Nick can get the template of a project plan from ChatGPT. Nick must refrain from entering any details of the project, such as the project name, agency, names of systems/software, high level requirements or staff members involved. These details could provide sensitive information about the project to ChatGPT.

⁸ Interim guidance on government use of public generative AI tools, Digital Transformation Agency [Interim guidance on government use of public generative AI tools - November 2023](#) | [aga \(digital.gov.au\)](#)

Generating files or documents

Stephanie is creating an urgent report to present later in the day and realises she is short of time to make a PowerPoint presentation. She is planning to upload the data for the slides to a public AI platform to help create a quick presentation.

What should Stephanie do?

Stephanie must not upload or input any sensitive or classified information into public AI platforms. She could input non-sensitive information, such as information available on a government website or an agency report that has been released publicly. She could also ask the AI for generic slide templates to help her prepare her presentation.

Using generative AI-powered search

Roque is writing technical requirements for a tender that needs to go out urgently. Roque wants to type his requirements into Google Bard to confirm some of the technical specifications around monitor resolutions.

What should Roque do?

Roque can use Google Bard to find information about technical aspects, however Roque should take care not to input any details of his agency's specific requirements or any organisational information. Roque should also take care not to mention the word tender or any market sensitive information that could, combined with his official email address, indicate a tender is being prepared by his agency. The information received should also be thoroughly validated by a human for appropriateness and accuracy before being used.

Exploring datasets

Angus is doing some basic data analysis on a publicly-available dataset. He is wondering whether he can use ChatGPT Plus to help him generate some insights from the data.

What should Angus do?

Given the data Angus is using is publicly available and there is no particular sensitivity to the topic he is researching, Angus can use ChatGPT Plus to assist him with data analysis.

However, Angus should remember that generative AI tools can produce output that is inaccurate or biased. Before sharing more broadly or making any use of the insights, Angus should check that the output fairly and accurately reflects the data.

Appendix B: Generative AI terminology

Algorithm: a series of instructions that allows a computer program to learn and analyse data in a particular way, such as recognising patterns, to then learn from it and accomplish tasks on its own.

Alignment: a process where additional training is applied to steer a large language model towards certain principles.

Artificial intelligence (AI): the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making and translation between languages. Modern AI is usually built using machine learning algorithms.

Bias: a type of error that can occur in a large language model if its output is skewed by the model's training data. For example, a model may associate specific traits or professions with a certain race or gender, leading to inaccurate predictions and offensive responses.

Bot (or Chatbot): A program that communicates with humans through text that simulate human language.

Ethical AI: Principles and practices aimed at ensuring AI technologies are developed and used in a manner that is fair, transparent, accountable, and respects human rights. Understanding ethical AI is crucial to government departments to ensure AI initiatives align with public values and legal standards.

Generative AI: refers to systems that use data models to generate new examples of content such as text, images, audio, code and other data modalities. Generative AI applications are typically trained on large amounts of real-world data and can approximate human generated content from prompts, even prompts that are limited or non-specific.

Hallucination: A well-known phenomenon in large language models, in which the system provides an answer that is factually incorrect, irrelevant or nonsensical, because of limitations in its training data and architecture.

Indigenous Cultural and Intellectual Property (ICIP) - refers to Indigenous peoples' rights to their cultural heritage based on the fundamental right to self-determination. Cultural heritage includes all aspects of cultural practices, traditional knowledge, resources and knowledge systems developed by Indigenous people as part of their Indigenous identity.

Indigenous data sovereignty: Refers to the right of Indigenous people to exercise ownership over Indigenous Data. Ownership of data can be expressed through the creation, collection, access, analysis, interpretation, management, dissemination and reuse of Indigenous Data.

Large language model (LLM): A type of neural network that learns skills — including generating prose, conducting conversations and writing computer code — by analysing vast amounts of text from across the internet. The basic function is to predict the next word in a sequence, but these models have surprised experts by learning new abilities.

Machine learning (ML): describes software components (models) that allow computers to recognise and bring context to patterns in data without the rules having to be explicitly programmed by a human. Machine learning applications can generate predictions, recommendations, or decisions based on statistical reasoning.

Natural language processing: Analyses and derives information from human language sources, including text, image, video and audio data. Natural language processing applications are commonly used for language classification and interpretation. Many natural language processing applications not only process natural language but also generate content that mimics it.

Neural network (NN): Computer programs or algorithms that simulate the behaviour of neurons in the brain. Consists of interconnected nodes, or neurons, that can recognise patterns and learn over time.

Prompt: the textual input supplied to the LLM. Prompts can be brief, for example “Who is the Prime Minister of Australia?”, or a long passage of text describing a problem or task in detail.

Synthetic data: Data that is artificially generated rather than produced by real-world events. Typically created using algorithms, synthetic data can be deployed to validate mathematical models and to train machine learning models.

Training data: Datasets used to help AI models learn, including text, images, code or data. The goal is to achieve a level of knowledge that allows the model to make accurate predictions or decisions.