



AI in government: fundamentals training

This training module will provide you with a foundational knowledge of AI, the principles of safe and responsible use, and how to put these to work as government staff .

This training module was developed for and by the Australian Government to support the implementation of the [Policy for the responsible use of AI in government](#). The module therefore refers to Australian specific resources that may not be applicable to your jurisdiction.

Contents

| | |
|--|----|
| Introduction | 3 |
| AI in government..... | 3 |
| Why does this matter to you?..... | 3 |
| What is AI? | 4 |
| What is AI? | 4 |
| Test your knowledge: Fill in the blank | 4 |
| Where is AI used? | 5 |
| Test your knowledge: Select the correct answer | 6 |
| What makes generative AI different? | 6 |
| How generative AI works..... | 7 |
| Test your knowledge: Multiple choice, quiz 1..... | 8 |
| Using AI with accountability | 9 |
| Safe, responsible use | 9 |
| You remain accountable | 9 |
| Putting AI to work | 10 |
| Before you use AI | 10 |
| Test your knowledge: Multiple choice, quiz 2..... | 11 |
| Deciding if it's right for you | 11 |
| Test your knowledge: Scenario 1..... | 12 |
| Using prompts and information | 13 |
| Using outputs | 13 |
| Test your knowledge: Scenario 2 | 14 |
| Wrapping up..... | 15 |
| Wrapping up..... | 15 |
| Where to find advice | 15 |
| Answers..... | 16 |
| Fill in the blank..... | 16 |
| Select the correct answer | 16 |
| Multiple choice, quiz 1 | 16 |
| Multiple choice, quiz 2 | 16 |

Introduction

AI in government

For decades, government has used Artificial Intelligence (AI) to perform specialised and general tasks, such as:

- investigating patterns of fraud
- providing information through virtual assistants
- digitising and organising physical documents.

Recent advances, including generative AI create new opportunities and challenges. Governments must build and maintain public trust to embrace these opportunities.

[Return to contents](#)

Why does this matter to you?

The community expects government to be an exemplar of the responsible use of AI.

We need to meet this expectation when we choose to use AI in our work.

This training will give you a basic understanding of AI and how to use it responsibly. By completing this course, you will be:

1. **informed** on the applications and risks of AI
2. **capable** of applying relevant advice
3. **confident** about when it's suitable to use generative AI
4. **empowered** to apply your own judgement with AI.

[Return to contents](#)

What is AI?

What is AI?

Artificial intelligence (AI) is a field of technologies used to make predictions or perform tasks that would normally require a person.

AI usually overlaps enough with human intelligence to be useful, even though they learn and work in different ways. These are some common terms you might have heard before.

Machine Learning: Techniques and technology which allow computer networks to find and learn from patterns in data. It has a significant role in most AI fields.

AI Models: A computer program that was trained on patterns in data so it can make predictions or perform tasks.

Computer vision: Techniques to identify patterns and make predictions about real-world objects as seen through images and video.

Natural language processing: Techniques to process, interpret and respond to written or spoken human language.

Generative AI: A type of machine learning that can generate new and convincing text, images or sounds from conversational prompts.

Training data: The starting data which an AI model learns prediction capabilities from. Flawed or unrepresentative training data may embed biases into the model.

[Return to contents](#)

Test your knowledge: Fill in the blank

- An AI model is trained to find *blank* in lots of data or information.
- AI uses its training to make *blank* or perform tasks.
- These technologies perform tasks that would normally require a *blank*.
- Generative AI is well known but it's just one type of *blank*.
- The quality of training data informs if an AI model will contain *blank*.

To check your answers, [head to the final page](#).

[Return to contents](#)

Where is AI used?

AI has existed for decades in specialised software and behind the scenes of popular platforms.

Recent advances, including generative AI, offer new ways to interact with AI and use it across many scenarios.

Here's some examples of where you might encounter AI in everyday life.

Search engines

AI-powered algorithms decide the rank of web pages and personalise results based on your search and browsing history. Social media feeds also use AI in a similar way.

Voice assistants

Intelligent assistants, like Siri or Alexa, process your voice commands, recognise keywords and decide the best action to take in response.

Generative AI chatbots

Generative AI takes conversational instructions or questions to create replies or complete tasks in general-use tools, like ChatGPT or Copilot, and an increasing number of customer service chatbots with a narrow scope or purpose.

Document scanning

Whether it's on your phone or in a digital service, AI is used to scan physical documents, recognise handwritten or printed text and use the results to organise information.

Data analytics

Data analysis tools use AI to identify, sort or classify large volumes of information, in a way that's faster than humans could but easy for us to verify its accuracy.

Maps and directions

Online navigation services will calculate the best route from A to B, using a mountain of information from simple distances to your past journeys and real-time traffic.

Autocomplete

Your smartphone suggests corrections, emojis or even a quick reply based on your past typing habits.

[Return to contents](#)

Test your knowledge: Select the correct answer

Which of these do you think are good examples of AI in everyday work?

- Search engines
- Autocomplete
- Maps and directions
- Data analysis
- Document scanning
- Generative AI
- Voice assistants

To check your answers, [head to the final page](#).

[Return to contents](#)

What makes generative AI different?

Many types of AI have been used in technology, companies and governments for years.

Recently, generative AI has introduced more people than ever with to a hands-on way to use AI. We can **interact with it through natural conversation** and it can **generate convincing text, images or sounds** in response.

This creates new opportunities and risks for government, which is why some of this training focuses on generative AI specifically.

You might encounter generative AI in everyday tools and technology, but it may not always be as obvious as a chatbot.

You might use it in software that lets you:

- type a conversational instruction instead of picking settings and options
- press a button to generate unique text or images.

[Return to contents](#)

How generative AI works

It's easier to make responsible choices once you know the basics of how generative AI works.

How generative AI is trained

Generative AI is 'trained' on large amounts of text, images, sounds or other information. It breaks this data into small pieces to find and learn from common patterns.

Developers can 'fine tune' to prioritise certain patterns, but the quality and quantity of training data has the biggest impact on its predictions and outputs.

Facts, meaning or ideas don't mean the same thing to generative AI as they do to us, so it might produce **the** wrong information.

How it generates output

Usually, the first step to using generative AI is an instruction or question, also called a 'prompt'.

It draws from its training to predict, one by one, the words most likely to follow the prompt. Some generative AI models can do this with images, sounds and video.

This is how generative AI creates an 'output' of convincing text, images or sounds, even if it's factually untrue.

Why generative AI isn't always accurate

Generative AI makes predictions with a touch of randomness – also called 'temperature' – to generate conversational outputs.

This means it will never produce the same output twice and might generate convincing statements that are misleading or factually untrue.

You can offset this by providing clear prompts, specific instructions and relevant supporting information, but you should always check outputs for relevance, biases and accuracy.

[Return to contents](#)

Test your knowledge: Multiple choice, quiz 1

Pick the most appropriate answer.

1. How is generative AI trained?
 - a. It finds and learns from common patterns in large amounts of information.
 - b. It learns answers to facts and ideas from large amounts of information.
2. How does generative AI make an output?
 - c. It considers the meaning of your questions to provide accurate answers.
 - d. It draws on patterns to predict what comes next, based on what came before.
3. What is one reason generative AI is not always accurate?
 - e. You didn't include personal or sensitive information in your prompt.
 - f. Because it adds randomness, it never generates the same output twice.

To check your answers, [head to the final page](#).

[Return to contents](#)

Using AI with accountability

Safe, responsible use

The community expects government to be an exemplar of the responsible use of AI.

People who work in government should have the confidence and support to make responsible choices about using AI.

Many resources about existing technology also apply to AI, like your agency's ICT use policies, values and code of conduct. But some decisions will not have an easy or clear direction.

AI Ethics frameworks such as [Australia's AI Ethics Principles](#), can guide your choices about using AI to:

- benefit individuals, society and the environment and ensure its not inaccurate or misleading
- respect human rights, diversity and the autonomy of individuals
- be inclusive and accessible, and ensure it does not involve or result in unfair discrimination against individuals, communities or groups
- uphold privacy rights, data protection and data security laws.

[Return to contents](#)

You remain accountable

We're individually accountable for how we use AI and its outputs in our work.

Golden Rules

To remain accountable for your use of AI, you must:

- explain, justify and own your advice and decisions
- assume prompts and supporting information could become public
- ensure personal, classified or sensitive information isn't inappropriately used or revealed.

Apply the golden rules whenever you consider using AI in your work, such as when you're deciding:

- if AI is appropriate for a task
- if it's the most effective tool for the task
- what to use in your prompts
- how to use the resulting outputs.

The next section explores these milestones and how to put AI to work by making responsible, effective choices.

[Return to contents](#)

Putting AI to work

Before you use AI

Before you even begin to use AI at work, ask yourself a few easy questions.

What am I trying to achieve?

Keep the requirements and ideal outcome of your task in mind to decide which:

- tools to use
- stakeholders to involve
- processes to follow.

Is it appropriate to use AI?

Some general-use tools, like generative AI, hold unique risks in policy and service delivery.

Alongside AI ethics frameworks, consider how you will manage prompts and outputs for:

- **accuracy** – how will you verify that information is accurate and relevant?
- **bias** – how will you identify and counter biases in prompts and outputs?
- **sensitivity** – how will you prevent, identify or remediate the use of personal, classified, sensitive or copyright-protected information?
- **transparency** – how will you make it clear where AI has influenced your advice or decisions, and allow them to be contested?

What do I need to know before starting?

Not every choice will be easy to decide, so familiarise yourself with:

- your agency's own AI and ICT use policies
- whole-of-government guidance, such as the Australian Government's [Guidance for using generative AI](#)
- your workplace's values, code of conduct and other requirements.

If there's not an easy or clear direction, use AI ethics frameworks such as [Australia's AI Ethics Principles](#) to guide your choices in a safe, responsible way.

[Return to contents](#)

Test your knowledge: Multiple choice, quiz 2

1. We should consider biases when using AI to:
 - a. prevent biases or negative assumptions in our work
 - b. understand the technology behind it
 - c. adapt it to perform well in different social contexts
2. It's important to maintain transparency when using AI because:
 - d. people should understand how it was trained and makes predictions
 - e. we need to know which advice or decisions it informs
 - f. you would otherwise have to explain and justify your advice or decisions
3. We need to consider personal, classified, sensitive or copyright-protected information to:
 - g. evaluate when the benefits of AI outweigh privacy or security risks
 - h. anonymise prompts and supporting information before using them
 - i. meet the obligations, policies and expectations of government staff
4. We should make sure prompts and AI outputs are accurate:
 - j. to test our own understanding of the subject matter
 - k. because it doesn't understand facts, ideas or meaning the same as us
 - l. so people don't have a reason to contest government using it

To check your answers, head to the final page.

[Return to contents](#)

Deciding if it's right for you

AI is only helpful if it's helpful to you.

Some AI tools, including generative AI, might help you complete general tasks faster so you can focus on what you do best. Whether you've used AI already or not, here are some ways to explore if it's the right tool for the job.

Tasks with plenty of background information

AI performs well with lots of information, like generating a summary of a long document or listing action items from a meeting transcript.

Brainstorming, rephrasing and other general tasks

AI performs well on tasks that are general in nature, like brainstorming the structure of an email, setting up a document template or trying different ways to phrase a sentence.

Tasks that take a long time but are easy to verify

AI can be useful on simple, repeatable tasks with outputs that are easy to validate.

Try using it to find names and descriptions scattered throughout documents, assign categories to a group of files or images, or search the internet to find helpful starting points for further research on a subject.

You don't always have to use AI.

Just because you have access to an AI tool doesn't mean you have to use it.

You might have specialist skills or non-AI tools to complete tasks faster, more reliably or to a higher quality. It might not be suitable for the task if you spend more time:

- correcting for inaccurate, biased or sensitive information
- editing or adjusting to meet the desired content style
- verifying and removing unreliable citations or supporting information
- making changes to work that uses AI based on manager feedback.

[Return to contents](#)

Test your knowledge: Scenario 1

Your agency held a survey to inform its next corporate plan.

You've been asked to provide a quick summary of the findings and are considering using a generative AI assistant to help you.

What's the responsible choice?

1. Don't use a generative AI tool to summarise lots of data.
2. Use the AI tool and report key findings to your manager.
3. Confirm the AI tool can be used based on the sensitivity of the information, and before deciding to use the tool.

If you answered 3, you're correct! You can use generative AI and critically analyse its outputs, then report key findings to your manager with the disclaimer that AI was used. You can use AI tools if they are suitable for the task and meet your needs. While an AI tool can help, you remain accountable **for** evaluating if it is appropriate to use and if it is the best tool for the task.

While using an AI tool can be helpful, there are essential steps you must take to ensure you remain accountable. Being accountable means you've taken the right steps to support your use of AI tools.

[Return to contents](#)

Using prompts and information

From training to prompting, what you get out of an AI tool is only as good as what goes in.

Consciously consider the **quality** and **integrity** of your prompts and supporting information, such as the documents or datasets you input into AI tools.

Prompt with quality

Take the time to ensure you provide an AI tool with:

- simple, clear and specific prompts
- essential or relevant supporting information
- robust, complete and up-to-date files or datasets.

Prompt with integrity

Work with colleagues and subject matter experts to:

- counter biases and assumptions in your own prompts and the data you use
- prevent personal, classified, sensitive or copyright-protected information being revealed
- ensure your prompt and supporting information is fair, accurate and relevant.

[Return to contents](#)

Using outputs

We must be able to explain, justify and take ownership of our advice and decisions and consider the risks of using AI.

With that in mind, consider how you will review, identify and adjust AI outputs for:

Accuracy: is it verifiably accurate and relevant?

Bias: is it fair and free of implicit bias or negative assumptions?

Sensitivity: does it reveal private, classified, sensitive or copyright-protected information?

Transparency: is it clear where the output is used and how to contest it?

When an AI tool has broad uses, its risks will be similarly broad and general, making them a bit harder to predict. Some simple choices can help you catch or prevent them outright.

Use AI for early ideation or drafts, not a final pass.

Plan adequate time to review and consult on work that uses AI outputs.

Ensure humans make, can explain and own key and final decisions.

[Return to contents](#)

Test your knowledge: Scenario 2

You need to write a report about leadership in organisations.

This will be informed by publicly available research on the gender identity of people in senior positions. You have access to a generative AI tool which you think could help you write it.

What should you do?

1. Don't use the generative AI tool because it will generate outputs with biases and assumptions.
2. Consider using the tool to generate an early draft, making use of the publicly available research.
3. Use the tool to write the entire report as it normally generates pretty accurate outputs and was trained on lots of data to begin with.

If you answered 2, you're correct! It can be useful for an early draft and publicly available information is trustworthy and safe to use. Take the time to review the draft for biases and inaccuracies before using it. It's great to be cautious, but if you're aware of the risk of inaccuracies and the risk of gender and other biases, you can make responsible choices with generative AI. However, generative AI doesn't understand facts and ideas the way we do. Instead, use clear, specific prompts for an early draft and review it for biases and inaccuracies before you use it.

Generative AI is a useful tool for generating early drafts.

However, if we fail to identify and mitigate biases or inaccuracies in prompts and outputs, it could have an unfair or dangerous impact on people, and marginalised or underrepresented groups.

Before you use outputs in government work, take the time to:

- Consider how AI might inadvertently create biased or unfair outcomes.
- Consider if you have the right skills and knowledge, or access to people and resources, to help you identify and mitigate biases and errors in your prompts and outputs
- Recognise where you may need to consult experts or people with lived experience, such as gender and sexual diversity, gender identity and expression multicultural communities, First Nations peoples, people living with disability and people facing different social or economic circumstances.

[Return to contents](#)

Wrapping up

Wrapping up

Congratulations! You have now completed the AI in government: fundamentals training.

AI technology is always evolving and new developments will present new concepts, opportunities and challenges. However, much of what you've learned today will remain true.

By completing this course, we hope you feel that much more:

1. **informed** on the applications and risks of AI.
2. **capable** of applying relevant advice.
3. **confident** about when it's suitable to use generative AI.
4. **empowered** to apply your own judgement with AI.

[Return to contents](#)

Where to find advice

It's okay if you need a refresh down the track.

There's one easy thing to remember: if you're about to use AI start by **checking advice from your agency and government**.

Follow your agency's advice first

Government agencies and their staff are ultimately responsible for how they adopt AI. Always follow your agency's advice and policies first. Ask for help from the relevant internal teams if you need it.

Check in with whole-of-government advice

Your agency still has to align with the whole-of-government approach to AI, so you can always check in with:

- whole-of-government guidance
- government expectations such as Australia's Public Service Values and Code of Conduct
- AI ethics frameworks such as Australia's AI Ethics Principles.

To quickly find many of these resources in the Australia context, visit digital.gov.au/ai/resources.

[Return to top](#)

Answers

Fill in the blank

1. An AI model is trained to find *patterns* in lots of data or information.
2. AI uses its training to make *predictions* or perform tasks.
3. These technologies perform tasks that would normally require a *person* or *human*.
4. Generative AI is well known, but it's just one type of *artificial intelligence*.
5. The quality of training data informs if an AI model will contain *biases*.

Select the correct answer

All of these options are great examples! AI powers many features of the tools and software we use every day, even though it might not be obvious to us.

Multiple choice, quiz 1

1. The answer is A. Generative AI is trained by finding and learning from common patterns in large amounts of information.
2. The answer is B. Generative AI draws on patterns to predict what comes next, based on what came before.
3. The answer is B. Generative AI is not always accurate because it uses randomness to stay conversational, so never generates the same output twice.

Multiple choice, quiz 2

1. The answer is A. We should consider biases when using AI to prevent biases or negative assumptions at work.
2. The answer is B. It's important to maintain transparency when using AI because we need to know which advice or decisions it informs.
3. The answer is C. We need to consider personal, classified, sensitive or copyright-protected information to meet the obligation, policies and expectations of government staff.
4. We should make sure prompts and AI outputs are accurate because it doesn't understand facts, ideas or meaning the same as us.