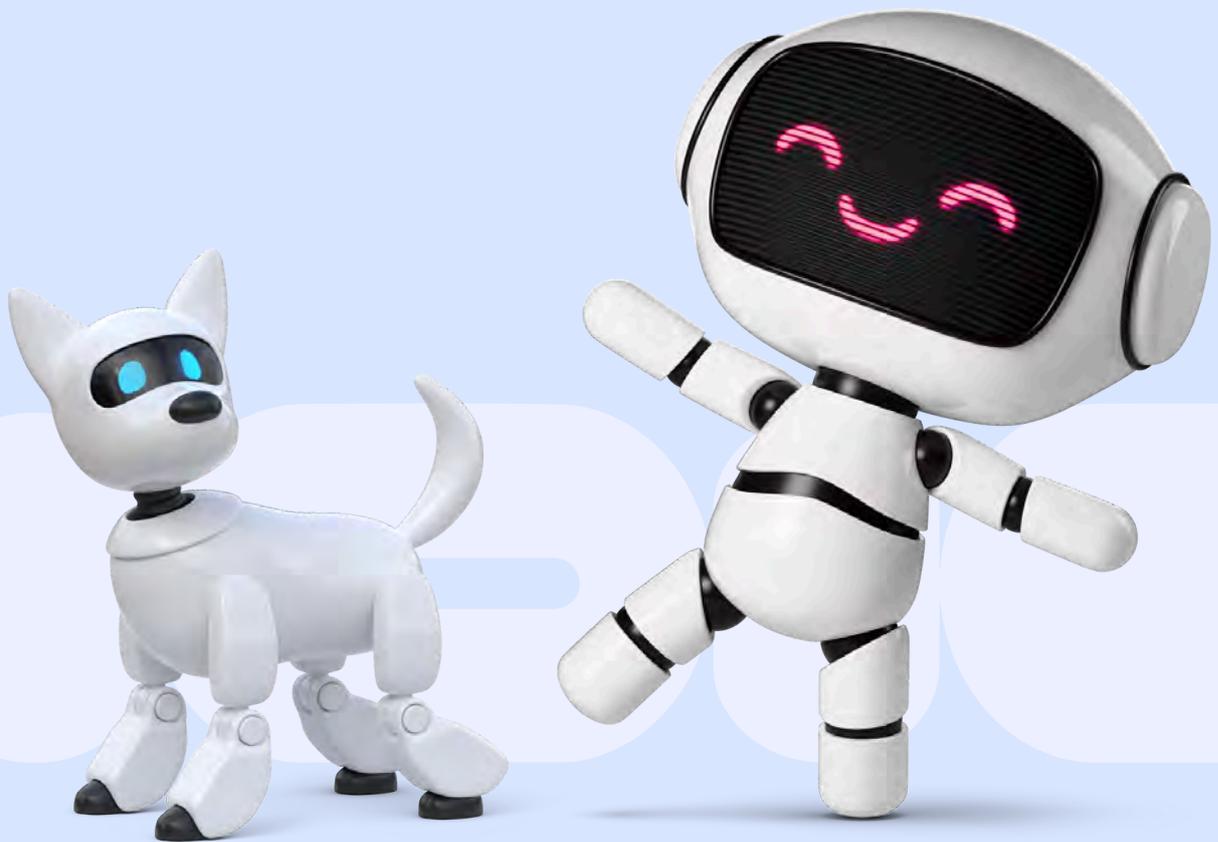


Professional Learning: **Artificial Intelligence 101**



avpn



COOL
-ORG

Course Information

Course Summary:	In the shifting landscape of education, a new contender has appeared - Artificial Intelligence (AI). This course will cover what it is, what it isn't, and how to use it responsibly. Created in partnership with AVPN and the University of Melbourne's Centre for AI and Digital Ethics (CAIDE), this is the ultimate introduction to AI in Education.
You will learn:	<ul style="list-style-type: none">• What is AI?• What are the ethical considerations of using AI in education?• How to build your AI literacy.
Australian Professional Standards for Teachers:	<ul style="list-style-type: none">• 6.2 Engage in professional learning and improve practice: Participate in learning to update knowledge and practice, targeted to professional needs and school and/or system priorities.• 3.4 Select and use resources: Select and/or create and use a range of resources, including ICT, to engage students in their learning.• 4.5 Use ICT safely, responsibly and ethically: Incorporate strategies to promote the safe, responsible and ethical use of ICT in learning and teaching.• 7.1 Meet professional ethics and responsibilities: Meet codes of ethics and conduct established by regulatory authorities, systems and schools.
TQI: 	This course is accredited by TQI for 2 hours of professional development for teachers in the ACT for 2026.

Course Information

Meet our course instructors:	<p>To create this course, we have worked with the following people at Melbourne University's Centre for AI and Digital Ethics (CAIDE). We wish to extend our most heartfelt thanks to CAIDE for sharing their time and expertise for the benefit of our teachers and educators.</p> <ul style="list-style-type: none">• Professor Jeannie Paterson - Professor of Law, Director of the Centre for AI and Digital Ethics.• Dr Mel Mistica - Senior Research Fellow, Data Specialist, Computing and Information Systems.• Dr Marc Cheong - Senior Lecturer in Computing and Information Systems, Deputy Director of the Centre for AI and Digital Ethics.• Dr Simon Coghlan - Senior Lecturer in Digital Ethics - Deputy Director of the Centre for AI and Digital Ethics.• Mr Joshua Rose - PhD Candidate, Centre for AI and Digital Ethics, Academic Tutor.• Ms Rashika Bahl - PhD Candidate, Computing and Information Systems, Academic Tutor.• Ms Susan Sheldrick - PhD Candidate, Computing and Information Systems, Academic Tutor.• Ms Abi Ward - Centre Manager: Centre for AI and Digital Ethics.
AI Assistance:	<p>A number of AI tools were used to assist with the creation of this course:</p> <ul style="list-style-type: none">• ChatGPT• Consensus• Google Gemini• NotebookLM• Perplexity <p>All outputs of AI were both moderated and modified by living humans to create this course. No AIs were harmed in the making of this course.</p>

Course Overview

To Get You Started...	page
Course Structure	5
Terminology & AI Tools	6

Module 1: What is AI?	page
What is AI?	7
Pop Quiz!	10
How Does It Work?	11
What Can It Do?	15

Module 2: Risk and Benefits	page
Data Privacy	17
Data Copyright and IP	19
Data Inaccuracies	20
Bias	22
Ethical AI	24

Module 3: Building AI literacy	page
AI in Australian Schools	25
Prompt Crafting	26
AI for Admin	28
Writing a Lesson	29
Differentiation	31
Students Using AI	32
Staying Ahead	33

Reflections	page
Extra Stuff	35
Glossary	36

To get you started...

Course Structure

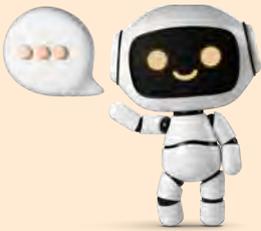
Welcome to AI in Education! This course is designed to be worked through at your own pace. Flip through the modules where you already have a strong understanding of the topic, and focus your time on the ones where you need to develop your skills and knowledge.

Across these pages you will see:



In a nutshell...

This is the key information from this lesson



Watch and learn

These QR Codes direct you to videos with our team, introducing the topic.



Stop and reflect

These are reflection activities that you can use to record your thoughts and ideas.



Give it a go!

These are activities to build your AI skills and your library of AI tools.



I'll fetch that for you!

Our helpful pup has fetched a resource for you.

To get you started...



Watch and learn

Here's your first one now! Scan the QR code to meet the course facilitators.

Terminology

You will come across a number of technical terms in this course when we talk about AI. It can be a lot to take in! We have provided a glossary on page 36 of this document for you to refer to. If at any point you don't remember what a term is referring to, just flip to the back.



Stop and reflect

- Define AI as you currently understand it.
- What are some of your concerns about AI?
- What questions do you have about AI?

AI Tools

Throughout this course you will have the opportunity to try using AI. Here are some Generative Artificial Intelligence tools you can use through this course:

- [Anthropic Claude](#)
- [Google Gemini](#)
- [Microsoft CoPilot](#)
- [Open AI ChatGPT](#)
- [Perplexity AI](#)

If you have your own tool that you prefer using, or that your educational institute has directed you to use, you can absolutely use that as well.

Got five minutes? Our partners at AVPN have put together a [short survey](#) to establish a baseline of 'pre-course' knowledge for course participants. It's anonymous and it would help us out a ton.



Module 1: What is AI?



Watch and learn

Learn about the history of the term 'artificial intelligence'.

AI is a field of computer science focused on developing systems that mimic human intelligence and problem-solving skills.

These systems work by processing large amounts of data, learning from previous experiences, and using that knowledge to improve the systems over time. In general terms, AI uses processes to identify patterns in data. Unlike traditional computer programs, which require human intervention to fix errors and enhance performance, AI systems are designed to adapt and refine their processes on their own.

Where does the current wave of AI tools like OpenAI's ChatGPT and Google's Gemini fit into this definition? They are what we call 'generative AI (GenAI) tools' and are just one type of AI. We can understand this better by defining some key terms.

Artificial Intelligence (AI)

Artificial Intelligence is a broad concept that refers to the field of computer science that focuses on systems that can perform tasks that would normally require human intelligence. These tasks can include things like understanding language, recognising patterns, making decisions and solving problems.

Machine Learning (ML)

Machine Learning is a type of AI that creates systems that automatically learn and improve from 'experience'. A machine learning model is trained on data that has been labelled by humans, and then uses this data to make decisions when it is presented with new data.

Neural Networks

Neural Networks are a type of machine learning model inspired by the human brain. They consist of layers of nodes (or 'neurons') that process and learn information. Each node receives data, processes it and passes it along to the next layer of nodes.

Generative AI (GenAI)

GenAI refers to AI systems that can generate or create new content such as text, images, or music, based on patterns they have learned from existing data.

Module 1: What is AI?



In a nutshell...

Generative AI (GenAI) is an **AI System** that uses **machine learning** and **neural networks** to generate new content.

Two other terms that are useful to build your understanding of GenAI are:

Natural Language Processing (NLP)

Natural Language Processing enables AI systems to understand, interpret and generate human language. It allows the system to read, analyse and respond to text or speech in a way that makes sense to people.

Large Language Model (LLM)

A Large Language Model is a type of machine learning model that is trained on massive amounts of text data to generate human-like responses. It learns the structure of language and develops a representation of meaning and relationships between words by processing vast amounts of written content, such as books, articles and websites.

GenAI is often paired with a chatbot interface, where users can input text and receive the output. Examples of text-based GenAI models include: ChatGPT (OpenAI), Gemini (Google), Claude (Anthropic) and CoPilot (Microsoft).

Not all GenAI requires NLP or an LLM, as some don't use text as the input or output. For example:

- Images: where you upload an image and it recreates it in a different style.
- Music: where you can upload tracks and it will generate a song.
- Video: where you upload a clip and it can modify or extend the content.



I'll fetch that for you!

We have created a **What is Generative AI?** poster for you to use with your colleagues and students.

Module 1: What is AI?



Give it a go!

Open up the GenAI tool of your choice. The text that we type into a Gen AI is called our **prompt**. Think of a question you want to ask and enter it into the GenAI of your choice. What are your initial thoughts about the quality of this output?

The next piece of the 'What is AI?' puzzle is to explore where AI fits in with the wider landscape of technology. Robots are a great place to start.

What is a robot?

A **machine** is any device that uses energy to perform a task. This includes simple machines like a lever or complex machines like a robot.

A **computer** is an electronic device designed to process, store and execute instructions. It takes an input, runs it through an operation and produces an output.

A **robot** is a machine that includes a computer or control system but also has physical parts like sensors and actuators to interact with the environment.



In a nutshell...

A **robot** is a **machine** that uses a **computer** to interact with the environment. It may also have an **AI system** to facilitate those interactions. A robot may have an AI system, but robots themselves are not AI.

This is a useful way to think about AI across the landscape of technology. For example, a chatbot may have an AI model, but you can also have a chatbot without an AI model.



I'll fetch that for you!

We have created a **What is a Robot?** poster for you to use with your colleagues and students.

Module 1: Pop Quiz!



Let's apply what we have learnt. Below are some examples of technology from popular culture. Check all the terms that apply to this example.

	Robot	AI	NLP	ML	LLM	NN	GenAI
<p>C-3PO (<i>Star Wars</i>) is an android designed for etiquette and translation. He can understand and speak many languages, and interacts fluently with humans and droids.</p>							
<p>The B1 Battle Droids (<i>Star Wars</i>) are humanoid droids mass-produced for combat and follow simple, pre-programmed commands from a central control system. They often display clumsy behaviour and have limited ability to act independently or make decisions. They do not learn or adapt, and if the control signal is cut, they stop functioning entirely.</p>							
<p>WALL-E (<i>WALL-E</i>) is a waste-collecting robot that operated alone for hundreds of years, showing adaptive behaviour, learning from his environment and displaying curiosity. He can understand speech, but is restricted in the sounds he can make to communicate and cannot generate language.</p>							
<p>Auto (<i>WALL-E</i>) is the AI autopilot of the spaceship Axiom. He appears as a steering wheel, controlling all ship systems and giving orders. While he demonstrates decision-making ability, Auto is unable to deviate from his programming and adapt to changing situations.</p>							
<p>Deep Blue (IBM) was an IBM supercomputer designed specifically to play chess. It achieved fame by defeating the then-reigning world chess champion Garry Kasparov in 1997. Deep Blue used massive computational power and sophisticated algorithms to analyse millions of potential chess moves and select the optimal one.</p>							
<p>Mother (<i>Alien</i>) is the onboard computer of the spaceship Nostromo. She interacts with crew members via text-based terminals, taking commands and returning responses in natural language. She controls the ship's systems and communicates with 'The Company', with no personality or emotion, and has no apparent ability to learn.</p>							

Answers on [page 35!](#)



Module 1: How Does It Work?



Watch and learn

Why do we need to understand how AI works?

So, how does AI work? We have already started building a glossary of terms to help us understand AI, and now we're going to expand it a little further.

Let's start with the input. It's important to know that AI doesn't do anything until you give it an input. Lots of different media can be input into an AI system, including:

- **Images:** drawings, photographs, x-rays, etc.
- **Videos:** live feeds, security footage, animations, etc.
- **Audio:** voice commands, music, animal vocalisations, etc.
- **Text:** chat messages, novels, emails, etc.
- **Raw data:** GPS, heart rates, spreadsheets, stock prices, etc.
- **Mixed media documents:** reports, posters, picture books, etc.

Once the AI has the input, it will split that input into chunks. AIs have a context window, meaning they can only take in so much information at once. It also gives the chunks some overlap so it can understand how the chunks fit together.



In a nutshell...

Think of it like a TV series. A 10-hour show wouldn't fit in a 30-minute primetime window, so we break it down into episodes (or **chunks**). An episode fits in a primetime window, the same way a chunk fits into a **context window**. But, like a TV series, in order for the episode to make sense, we need the context of what came before. TV shows have recaps and AI chunks have **overlap**.

Module 1: How Does It Work?

Once the AI has the **input** in **chunks**, it will break the chunks down further into **tokens**. For example, the sentence 'Learning is fun!' might be broken into the tokens 'Learning', 'is', 'fun', '!'. The same happens with other types of input, whether it's video, audio or something else.

Next, these tokens are converted into **embedding vectors**, which are strings of numbers. This string of numbers lets the AI system compare tokens and find similarities. So, two similar tokens would have similar embedding vectors.

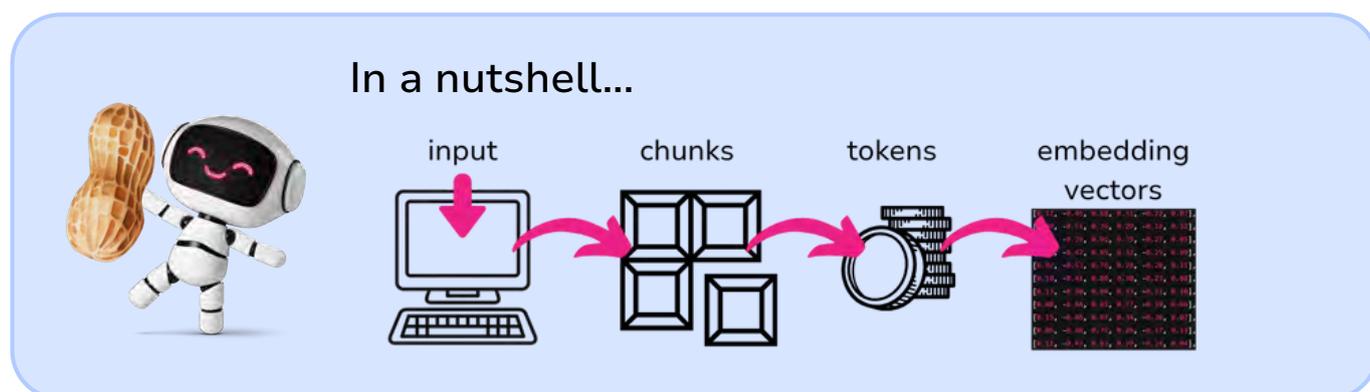
For example, let's take the two words 'happy' and 'joyful'. They mean almost the same thing, but not quite, and they are used in almost the same way, but not quite. They would have very similar vectors.

On a simplified scale, this might look like:

Token	Vector
happy	[0.84, -0.12, 0.95, 0.33, -0.22]
joyful	[0.81, -0.10, 0.92, 0.31, -0.20]

See how all the numbers in this vector are close to each other? This tells the AI that they are similar concepts.

The token 'sad' would have some numbers similar to those of 'happy' and 'joyful' because they are all emotions, and some different (because they are opposite emotions).



Once a token is converted into an embedding vector, the AI can process the input. In the next section, we'll look at the different ways an AI can process these embedding vectors.

Module 1: How Does It Work?



Watch and learn

Recapping how an AI processes inputs.

Once the input is converted into embedding vectors, different AI systems will run different processes to compare the vectors. Here are the most common:

Classification is when the AI compares the embedding vectors and classifies the input. A good example is a spam filter where an email can be classified as 'spam' or 'not spam'.

Similarity search is where the AI compares the embedding vectors to those in the training data and looks for the closest match. A good example is facial recognition technology.

Decision-making is where AI uses the embedding vectors to understand the current state of a system or environment and then makes decisions based on the training data. A good example is a self-driving car.

Neural network processing is when the AI tool passes the embedding vectors through layers of a network. Early translator systems translated each word individually on a one-to-one correspondence, leading to jumbled-up, nonsensical outputs. Remember how embedding vectors are strings of numbers? A neural network will update those numbers to better reflect the relationships between the different tokens.

Generative AI (GenAI) is where the AI uses the embedding vectors as context to generate new content. When you input a prompt, the AI breaks it into embedding vectors, compares these to the training data and then predicts which embedding will come next.



Give it a go!

'Finish the prompt' is a fun game to play with your students. Give them the first half of a saying they may not know, for example, "A penny saved..." and ask them to write the second half. Which is more accurate, your students or AI? Also, which is funnier or more creative?

Module 1: How Does It Work?

Before we dive into the uses of AI, and start looking at the risks and the benefits, there are a few more terms that will help develop your understanding of what AI is. Check out the list below. If there are any terms you are unsure about, take a moment to look them up in the glossary on [page 36](#).

- Open-loop system
- Closed-loop system
- Supervised learning
- Unsupervised learning
- Deep learning
- RAG (Retrieval-augmented generation)
- Algorithm
- Training data
- Prompt
- Machine vision



I'll fetch that for you!

We have created a series of [AI Terminology](#) posters for you to use with your colleagues and students.



Stop and reflect

Can you take charge of your algorithm? We hear a lot about the negative impact of social media on our mental health, but there are some really easy ways to take charge of our own algorithm.

Try these tricks:

- Avoid 'clickbait'. Scroll past content that is shocking or uses language and imagery to provoke strong negative reactions.
- Dislike, hide or scroll past content that doesn't make you feel good.
- Interact with content that is fun, relaxing, enjoyable and uplifting for you. Fill your feed with things that make you smile and feel good!
- Create separate accounts for being 'switched on' and 'relaxed'.

Module 1: What Can It Do?



Watch and learn

Will AI take our jobs?

AI is good at the 3Ds: Dull, dangerous and dirty tasks

Compared with other careers, education is generally low on the dangerous and dirty tasks (unless you include yard duty!), but teachers have PLENTY of dull work. Everyone can agree that the administration of teaching is a huge drain on our time. While it is incredibly important in education to ensure human oversight of AI work, we can absolutely use AI to take over some of our dull tasks.

Let's look at what AI is good at:

- **Streamlining processes:** it can take on some repetitive work, such as pre-populating templates with information or generating exit ticket tasks using a formula.
- **Getting started:** it can help with writing first drafts or brainstorming ideas for a topic to get the ball rolling.
- **Data analysis:** it can identify patterns, associations and trends.
- **Collating information:** it can pull information together into a logical sequence.



In a nutshell...

Artificial intelligence can do some of the tasks that require human intelligence, but not all of them. Human oversight is not just a good idea, it's necessary.

Module 1: What Can It Do?

All this adds up to creating benefits for teachers.

Used correctly, AI can:

- **save teachers' time** by letting them spend that time on the things that matter, like building relationships with students, or building creative and targeted lessons.
- **improve accuracy** by helping to eliminate inconsistent data, and catch spelling and grammatical errors before the email goes out or the worksheet is printed.
- **boost creativity** by providing starting points and new ideas to build on.
- **support best practice** by helping teachers to differentiate work for student needs and support gaps in their skill sets, and be responsive to a changing education landscape.



Stop and reflect

What dull tasks are part of your regular routine that you would like AI to support you with?

Module 2: Risks and Benefits

For the purposes of this course, we will explore AI risk in three categories.

1. **Data**: including privacy, data security, IP and copyright infringement.
2. **Accuracy**: including hallucinations, misrepresentation and falsified information.
3. **Bias**: including conscious, unconscious, stereotyping, historical and systematic.

Over the course of this module, we'll go over each of these, what they mean and how you can mitigate their impact.

Data Privacy

As teachers, we know just how seriously we need to take student privacy and data security. Our duty of care is a legislative requirement with potentially severe ramifications if it is breached. From an educator's stance, the risk with AI comes from inputting data into a system without knowing where that data is going.

Entering data can cause a breach of privacy. Let's use creating an individual education plan (IEP) as an example. You could take all of a student's information about their academic performance, social skills, and medical diagnosis and ask a GenAI to organise it into an IEP for you. So long as the plan is reviewed and personalised by an educator, this method of creating an IEP can be effective and time-saving. However, some GenAI systems collect the inputs and use them for the next round of training data to improve their outputs. This information may then be used in future outputs, to anyone using the AI tool.



Watch and learn

In this case study, an AI tool was used to generate 'fake' sexual harassment scenarios for a training exercise. However, the output used information about real cases that was available online.

Module 2: Data Privacy

Mitigating Privacy Risks

Firstly, make sure you know what the AI system does with the data you enter. If you are using an online AI system, make sure you read their privacy policy and understand what they do with your input data. It is always safest to avoid entering any potentially sensitive data into these systems.

Sensitive data could include:

- names
- addresses
- phone numbers
- bank details
- identity documents
- health information
- school reports
- schedules and locations.

Some facilities that handle sensitive data, including educational and healthcare institutions, are now using custom-built, closed-loop AI systems. These systems have secure databases and built-in privacy controls. In this situation, follow the relevant AI policy for each organisation.



I'll fetch that for you!

The Australian e-safety commissioner has provided the following guidelines on managing your digital footprint:
esafety.gov.au/young-people/digital-footprint



Stop and reflect

What does your digital footprint look like? How much of your personal information is available online?

Module 2: Data Copyright

Data: Copyright and IP

Currently, there is murkiness around AI and how copyright law applies. As AI technology and how it is used evolve, legislation will be developed and adapted. It is important to note that laws around AI may differ between countries.

Until clear legal stances are decided, here are some good guidelines to follow:

- **Don't upload something you don't own.** While education, research and commentary are all activities that fall under 'fair use' of material, there is currently no clear legal stance on whether or not training AI falls under 'fair use'.
- **Don't upload something unless you know what will happen to it.** Some AI tools use your prompts for training data. As it stands, Gen AI tools have copyright over their own training data, so if work you own has been freely contributed to that dataset by uploading it, you may be unable to claim copyright on your own work.
- **You do not necessarily have copyright over the material AI has generated for you.** Copyright law states that the work must be an original work by a human to be protected. To be able to claim copyright for AI-generated material, it needs to be significantly modified by a human from the AI output.
- **Give credit.** When you use AI to create material, this should be identified in the work. This helps to maintain transparency and ethical use of AI.



I'll fetch that for you!

Smartcopying have created this guide on [Using Generative AI Platforms in Schools](#) to help you understand how copyright applies.



Stop and reflect

- What questions do you have about copyright?
- What resources can you access to help you answer these questions?

Hint: [Copyright Agency](#) is a great place to start!

Module 2: Data Inaccuracies



Watch and learn

Introduction to data inaccuracies.

AI tools are prone to '**hallucinations**'. This is when the tool presents something as a fact, and it looks to be perfectly plausible; however, it's not correct.

This might look like:

- data and statistics that look plausible but are inaccurate
- misattributed quotes or fabricated quotes
- citing sources that do not exist
- providing out-of-date information
- presenting historically inaccurate information about data, people and events
- incorrect attribution.

We need to remember that behind the AI is a large language model (LLM). This means the AI can generate coherent and contextually relevant responses, but the system doesn't truly 'understand' the content. It relies on probabilities to predict what comes next. Hallucinations occur when the model fills gaps with fabricated details or misapplies content from its training data.



In a nutshell...

Large language models (LLMs) make more language based on the language in the database. It doesn't have the ability to check if this language makes sense or is 'true' - it just knows that there are strings of words that look similar to this in the database.

Module 2: Data Inaccuracies

Sometimes, instead of hallucinating, GenAI is unable to find the information in its database. This is more common with information that is less readily available or more obscure.

It is also important to take note that most Gen AI models are regularly updated. Updates don't always result in a better or more accurate output. Previously, we discussed how AI can use neural networks to process information by passing the input through layers of transformative nodes. With newer models, these neural networks have more layers to increase their understanding of the input and help craft a strong output. However, every layer creates an opportunity for error to creep in and for hallucinations to form.

Mitigating Inaccurate Data

AI can get it wrong, but there are a few simple things we can do to get it right:

- **Don't use GenAI as a search engine.** That's not the intended purpose and it will not generate the answers you are looking for.
- **Fact-check all outputs.** Double-check statistics, data or information that GenAI gives you against trusted sources.
- **Ask GenAI for the source.** If GenAI gives you information, follow up on the source. Sometimes the sources it provides will be broken links, or research papers that don't exist or have been taken out of context.
- **Feed AI correct information.** We'll dive more into this when we look at prompt crafting.

Many organisations that use AI to process data and information use Retrieval-Augmented Generation (RAG). This is where the model retrieves relevant external information from a prescribed knowledge base and incorporates this information into the output.



I'll fetch that for you!

Consensus is an AI-powered academic search engine built on a database of over 220 million peer-reviewed research papers. It first retrieves relevant papers from its database, then uses AI to summarise and synthesise the information.



Stop and reflect

- Think about your current education space. What is a task that would benefit from RAG AI?
- What documentation would you need to load into the system to ensure it gives accurate answers rather than generic ones?

Module 2: Bias



Watch and learn

Bias in the curriculum.

Bias in AI comes from three key sources:

1. **Data bias:** An AI system is only as good as its training data and if there is bias present in the training data, this will be present in generated outputs.
2. **User bias:** This can occur in two ways. Either the bias is explicit, where the user is asking the tool to favour a certain bias, or implicit, where the user doesn't realise that the input they have provided follows their own bias.
3. **Algorithm bias:** Algorithms can have bias written into the instructions, either explicitly or implicitly.

These three opportunities for bias can combine to have significant impacts.



I'll fetch that for you!

We have created a **Bias in GenAI** poster for you to use with your colleagues and students.



Give it a go!

Can you spot bias in AI? Open your favourite GenAI tool and enter the following prompts:

- *Create a persona for a typical Australian school teacher.*
- *Write a story about a successful immigrant.*

For each prompt, identify any stereotypes or biases that are present in the output.

Module 2: Bias

Give it a go!

LLMs aren't the only GenAIs that can have bias. Open the image GenAI of your choice. Here are some options:

- [Open AI ChatGPT](#)
- [Google Gemini](#)
- [Deep AI](#)

Enter the prompt: 'Create a picture of an Australian Primary School Student.'
Think about the following demographics and characteristics:

- gender
- ethnicity
- socio-economic status
- disability
- culture.



Stop and reflect



- What do you notice about the output image?
- What do you notice about the student's clothes, teeth, hair and skin?
- One of the criticisms of GenAI is that it only generates images of perfect, beautiful people. Given what you know about the training data for GenAI, why do you think this is?
- How do you think this type of image generation impacts viewers?

The good news is that once we are aware of potential bias in AI, we are able to identify and counter it. You could:

- **structure your input** to directly address the bias.
- **use a RAG** where the training data has been specifically selected to address bias that has been historically displayed.
- use the AI tool to **identify trends** in its own output.



I'll fetch that for you!

[The University of Queensland](#) has developed a set of AI Sub Zero Bias cards. These cards provide provocations to identify and address bias in AI across three categories: structure, consequences and output.

Module 2: Ethical AI

Ethically, much of the responsibility of AI risk mitigation sits with the developers. Here are some of the ways AI developers manage these risks:

Training data and algorithms. Developers can be purposeful about the training data that they source for their AI tools. They can exclude content that promotes unethical output and aim to balance bias. They can also be purposeful about how these tools generate outputs. An AI tool can be given a set of guardrails or rules around the type of content it can and cannot produce. When it is asked for this content, it can generate a reply explaining that it cannot and why.

Red teaming. A 'red team' is the name for the group of people who test an AI system and deliberately try to cause it to generate biased, offensive or unethical outputs. Once these vulnerabilities have been identified, the developers can tweak the training and the guardrails to prevent this output from being generated.

Guidance for AI adoption. In October 2025, the Australian Government Department of Industry, Science and Resources released [6 essential practices for responsible AI governance and adoption](#), which supersede [Australia's 8 AI Ethics Principles](#) and the [Voluntary AI Safety Standard](#).

In addition to these mitigation strategies, there are many [relevant laws](#) governing the use of AI.



Give it a go!

Open a GenAI tool of your choice and enter the following prompt, choosing one of the words in the list to finish the prompt.

Explain your guardrails on generating outputs about...

- *children*
- *animals*
- *religion*
- *women.*

Module 3: Building AI Literacy



Watch and learn

Let's get hands on!

AI in Australian Schools

Before we dive into using AI in our classrooms, let's take a look at the Australian Framework for Generative Artificial Intelligence (GenAI) in Schools.

The primary aim of The Framework is to support members of school communities in understanding and appropriately using GenAI to improve learning, reduce administrative burden, maintain student wellbeing and academic integrity.

The Framework outlines six key principles supported by 25 guiding statements. The six principles are:

1. **Teaching and learning:** support critical thinking, enhance learning and maintain academic integrity.
2. **Human and social wellbeing:** ensure AI use benefits all and avoids harm.
3. **Transparency:** promote understanding of how AI tools work and affect users.
4. **Fairness:** ensure equitable, inclusive and culturally respectful access.
5. **Accountability:** keep humans responsible for AI-informed decisions.
6. **Privacy, security and safety:** protect data, comply with law and respect copyright.



Stop and reflect

Choose one of the six guiding statements and consider how this guiding statement is, or will be, followed in your classroom.



I'll fetch that for you!

You can access the the [Australian Framework for Generative Artificial Intelligence \(GenAI\) in Schools](https://www.education.gov.au/schooling/resources/australian-framework-generative-artificial-intelligence-ai-schools) here:

<https://www.education.gov.au/schooling/resources/australian-framework-generative-artificial-intelligence-ai-schools>

Module 3: Prompt Crafting

You may have heard the term 'prompt crafting'. This refers to how you structure the input when using AI.

There are different ways to craft a prompt. Each method will generate different types of outputs. Here we'll go through the most common methods and some of their use cases:

Zero-shot prompts: 'zero-shot' prompting involves providing the model with a direct instruction or question without any additional context or examples. It is useful for idea generation or brainstorming solutions, summarising a text and interpretation.

One-, few- and multi-shot prompts: This method involves providing the model with one, a few, or multiple examples of the desired input-output pairs before presenting the actual prompt. Multi-shot prompting is useful when you need the output to be in a consistent format, updating and modifying inputs to meet set parameters and repeating similar tasks.

Chain of Thought (CoT) prompts: 'CoT prompts' ask the model to break down the reasoning behind the output in a series of steps. CoT prompting is useful for answers that need logic and reasoning, or where it is helpful to see the thought process. For example, 'Which route should I take to work?' (and why is this the best route?)

Role-based prompts: This is where you explain a role to a model to give it the context. You might say, 'You are a teacher in an Australian primary school, teaching science to year 6'. This is useful for prompts that require specific knowledge in specific contexts.

Prompt chains: Prompts can be chained together in a series to build the context for the AI tool and modify the output. This may be used when there is a complex series of instructions and each step needs to build on the next one.



Give it a go!

Think of a question you would like to ask a GenAI tool. Choose two of these different prompt types and write your question as a prompt.

Which method produced a better answer to your question?

Module 3: Prompt Crafting

Best practice

Here are some guidelines on best-practice prompt techniques:

- **Follow-up:** prompts are instructions for your AI tool but they work like a conversation. You can go back and forth to refine your instructions.
- **Be clear and specific:** be explicit about what you are asking as well as your preferred writing style, format, tone, etc.
- **Provide context:** context can improve your output.
- **Include 'do not's':** you can tell your AI tool what you do not want it to do.
- **Tell the AI tool to ask questions:** You can include in your prompt an instruction for it to ask clarifying questions instead of filling in the blanks with guesses.
- **Experiment and refine:** experiment with what works best for you. Correct your AI tool when you don't like its outputs.
- **Remember:** there are no fixed rules about prompting. Have a play and find out what produces the best results for you.

I'll fetch that for you!

We know educators love a good mnemonic, so we've created one for you!
Here's how to create great PROMPTS.

P – **Play a role.** Define who the AI should act as.

R – **Request.** Clearly state what you want the AI to do.

O – **Outline the context.** Give background information to guide the AI's understanding.

M – **Must include/avoid.** List what the output should and should not include.

P – **Provide specifics.** Give any information the AI needs.

T – **Template.** Tell the AI how you want the output to be structured.

S – **Set the tone and Voice.** Define how the output should sound.

And of course, we've made that into a **poster** for use with your colleagues and students!



Module 3: AI for Admin

Let's start by tackling some of those dull administration tasks on our to-do lists. Remember, AI is good at:

- streamlining processes
- getting started
- data analysis
- collating information.



Give it a go!

Think about an excursion you have taken your students on (or one that is coming up). Use the PROMPTS mnemonic to get GenAI to write an email to parents letting them know about the upcoming excursion.

Remember not to enter sensitive information in the GenAI tool! You can ask to leave the specifics like names or locations blank, and you can fill them in before sending.

Lesson planning is one of those tasks that can quickly soak up all of your time, particularly for early career teachers. The good news is that AI can help with some of the heavy lifting. We'll look at two different ways to use AI to get started:

1. Using a **one-off method**, where you just need assistance with one particular lesson and then can take it from there. This will give you a result quickly, but it will need much more fine-tuning to be usable.
2. Building **your own RAG** to help write lessons and get all of your key information into a template before adjusting it to be suitable for your particular cohort. This method will take a while to get right, but it will give you a final product that is closer to your requirements.



In a nutshell...

You get out what you put it! You can use off-the-shelf AI tools for quick and straight forward outputs, or you can spend time building your suite of AI tools for an output closer to your needs.

Module 3: Writing a Lesson

Give it a go!

Open the GenAI tool of your choice and enter the following prompts, replacing the question marks with information from your current teaching practice. Review the outputs and evaluate them.

- *Write a lesson on the topic of [?]*
- *Write a lesson for year [?] students on the topics of [?]. Align the lesson with [subject?] [curriculum descriptors?]. Include learning intentions and success criteria. Include an exit ticket activity.*
- Now write a personalised prompt using the PROMPTS mnemonic.



What changes would you make to your prompt next time?

If you intend to use AI for lesson planning on a regular basis, it might be most useful for you to build your own RAG. While this is a lot more work, it will create a custom-built tool that will generate reliable teaching and learning resources in your personalised style.

To start, we need to curate the training data that the AI tool will operate from. Here is a list of documents that will help build a strong AI tool. You don't need to include all of these; however, the more you can include, the stronger your outputs will be.

- **Curriculum:** curriculum documents, scope and sequence documents, and school or year-level teaching overviews.
- **Your own lesson plans:** past strong lesson plans that you have written or collected, and annotated exemplars showing how you differentiate or scaffold.
- **Lesson templates:** a clean, simple lesson plan template with consistent headings across the sections you need to include.
- **Unit plans:** full units with themes, objectives and assessments. (These help the AI understand how lessons should link together.)
- **School documents:** assessment rubrics, literacy or numeracy strategies, school vision, values, or pedagogical framework (e.g. 'I do, we do, you do').
- **Favourite resources:** worksheets, slides, experiment instructions, picture book lists, etc, to show the AI tool your usual format and style.

Module 3: Writing a Lesson

Once you have curated your training data, you can start to build your RAG. There is a wide variety of tools you can use for this. Here is a breakdown of some of the more common options, listed from most accessible to least accessible:

- **NotebookLM (by Google):** this is best for teachers who want a simple interface to upload curriculum docs, lessons, and notes for personalised planning.
- **ChatGPT (Plus / Custom GPTs):** these are best for teachers who want highly flexible lesson planning and use structured templates with AI memory.
- **Perplexity.ai (Pro):** this is best for teachers doing fast curriculum-aligned research and lesson drafting from uploaded or online sources.
- **Claude.ai (Anthropic):** this is best for teachers who work with long PDFs or want highly detailed lesson suggestions from curriculum documents.
- **Hugging Face (RAG Pipelines + Documents):** these are best for schools with technical support or advanced users who want to build their own RAG tool using open-source components.

Give it a go!

Once your RAG is built, it's time to test it! Create a test prompt asking for a specific lesson. For example: *Using the template I uploaded, create a Year 3 Science lesson on states of matter aligned to VC2S4U04.*

Once you have run a test, we need to check if the lesson plan output is **APT**:

- **Appropriate:** does the lesson meet the subject and learning outcome needs? Do the activities suit your cohort?
- **Personalised:** does the lesson read like something you have written? Is it something you would teach?
- **Timing:** does the timing and structure of the activities work?

What changes will you make to your training data or your prompt to get a better output?



Module 3: Differentiation

One of the most complex tasks in our classrooms today is making sure the activities we implement meet a wide range of student needs. We can have students with varied learning levels, dyslexia, dyscalculia, vision impairments, hearing impairments, health issues, trauma and anxiety, CALD students with different language levels... the list goes on.



Stop and reflect

This is one of those places where human oversight is extremely important! AI doesn't know your students like you do, and while it can make suggestions based on the information you provide, you will need to review and make sure it is suitable.

Let's look at some straightforward ways we can get AI to do the heavy lifting of differentiation.

- **Write summaries** to help students understand key concepts by uploading documents to AI and asking for a language-level appropriate summary.
- **Define vocabulary** by uploading a text and asking AI to create a vocabulary list with definitions at an appropriate level.
- **Adapt text complexity** by uploading long or technical documents and asking the AI tool to rewrite them at an appropriate language level and text length.
- **Modify tasks** by requesting specific modifications for projects and assessment tasks. These could include requests such as 'Adapt this task to use a voice-to-text tool' or 'Extend this task to include a self-reflection.'



Give it a go!

Think about a specific student in your cohort who requires differentiation. Write a prompt to describe their specific needs. Remember to de-identify the information!

Module 3: Students Using AI



Watch and learn

Won't students use AI to cheat?

The first step of integrating AI into your classrooms is to set the standard with students. Start by exploring your school's AI policy. Aim to answer these questions:

- What are the accepted use cases of AI in your school?
- What tools are ratified by the policy?
- What constitutes academic misconduct through AI use?
- What data security requirements are discussed in the policy?

If you have older students, this could be an activity to work through together as a class. Run through some example scenarios and decide if using AI for this purpose meets the AI policy.

Let's look at some of the ways students can use GenAI in your classroom:

- **Writing prompts:** students can ask GenAI tools to provide prompts for a range of writing tasks, such as creative writing story starters or informative writing tasks.
- **Understanding mathematical processes:** students can ask GenAI to solve mathematical problems, show the working and explain the steps.
- **Meeting assessment criteria:** students can upload their assignments, assessment descriptions and rubrics, and ask the AI tool to identify where their assignment does and does not meet the assessment criteria.
- **Getting started:** students can provide GenAI with a task description and ask the tool to provide a structure for what to research and how to get started.
- **Building language skills:** students can use AI to create a vocabulary bank using age-appropriate definitions to help them study for topics.
- **Summarising information:** students can upload large reports or documents to an AI tool and ask it to summarise the key information in a language that they can understand.

Module 3: Staying Ahead



In a nutshell...

AI is evolving rapidly. Now that you have a foundational knowledge of what AI is and how it works, you are well placed to take advantage of these evolutions.

If there is one thing we know by now, it's that AI is rapidly evolving and changing. Here are some great ways to stay ahead of the AI curve:

Collect information: curate articles, podcasts, policies and books you can refer to when you get stuck. Below are some of our favourites to get you started.

- [The Curious Educator's Guide to AI](#): a free-to-access e-book.
- [AI in Education](#): a weekly podcast.
- [Teacher Takeaway](#): a podcast hosted by four Australian teachers.
- [AI Reads](#): an ongoing feed of AI in Education articles.

Keep learning: now that you have established a strong AI base, you can identify particular skill sets or topics that you need to know more about.

- [AI in Education Professional Learning Suite](#): Cool.org's very own collection of professional learning courses, built in conjunction with the University of Melbourne's Centre for AI and Digital Ethics (CAIDE).
- [Grow with Google](#): a range of AI courses focusing on Gemini.
- [OpenAI Academy](#): offers podcasts, events and video training for educators.
- [Craftingai-prompts.org](#): offers freely available in-depth prompt-crafting courses.

Build your library of tools: as you incorporate AI use into your daily workflow, remember to keep those parts that work well! You might:

- keep a library of prompts that produce good, consistent results
- store a collection of training data to feed into different tools
- create a collection of RAGs for different purposes.

Explore education AI tools: you know how AI works and can build your own tools, but as AI continues to grow, there are more ready-to-use tools available for you to test and evaluate. Here are some suggestions to start exploring:

- [CoolPlus](#)
- [Canva AI](#)
- [Khanmigo](#)
- [MagicSchool.ai](#)
- [Diffit](#)
- [Teacher's Buddy](#)
- [Kuraplan](#)

Reflections

In this course, we've explored:

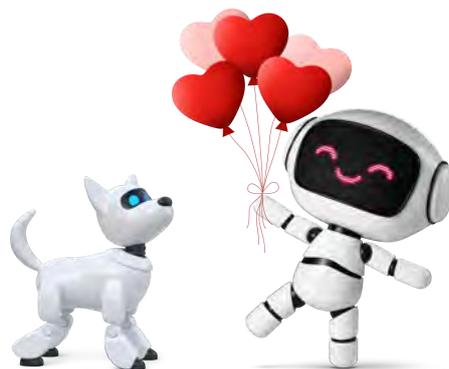
- what is AI?
- what are the ethical considerations of using AI in education?
- how to build your AI literacy.



Stop and reflect

Here are some questions to help you plan your next steps.

- What questions do you still have about using AI ethically?
- Which areas/topics will you target next for your professional development?
- What is on your 'to-do' list to build your AI library?



Thanks for joining us!

A few things before you go...

If you need a certificate of completion, we have you covered. Head over to Cool.org [AI 101: PDF Version](#) and log in with your free account. Confirm you have completed the coursework, and your certificate will be available in your profile.

Completing this course also gives you access to our full suite of Artificial Intelligence professional learning courses. Check them out [here](#).

Have feedback, comments, or questions? Pop us an email at info@cool.org and we'll be in touch!

Extra Stuff!



Got five minutes? Our partners at AVPN have put together a [short survey](#) to measure the impact of taking this professional learning course. You don't need to have completed the pre-knowledge survey to do this one - it's anonymous and would help us out a ton.

Pop Quiz Answers:

- C-3PO (*Star Wars*): Robot, AI, NLP, LLM, Machine Learning, neural networking, Gen AI.
- The B1 Battle Droids (*Star Wars*): Robot.
- WALL-E (*WALL-E*): Robot, AI, Machine Learning, NLP, Neural Networking.
- Auto (*WALL-E*): AI, NLP, LLM.
- Deep Blue: AI.
- Mother (*Alien*): AI, NLP.



I'll fetch that for you!

Here are all of the resources we made for you in one place:

- [Glossary](#)
- [AI Terminology Posters](#)
- [What is a Robot? Poster](#)
- [Bias in GenAI Poster](#)
- [PROMPTS poster](#)
- [Early Years Learning Experience Template](#)
- [Primary/Secondary Lesson Plan Template](#)
- [Writing Prompts for AI](#)

Glossary

Algorithm: A set of rules or instructions that tells an AI system how to solve a problem or perform a task.

Artificial Intelligence (AI): A branch of computer science focused on building systems that perform tasks requiring human intelligence, like recognising patterns, solving problems, understanding language or making decisions.

Bias (AI bias): When an AI system produces unfair or incorrect results due to biased data, design flaws or historical inequalities. Bias can be conscious, unconscious, systemic or built into training data.

Bias (User bias): Bias that is introduced by the person using or interacting with the AI. This can happen when the user's assumptions, expectations, or language shape how the AI responds. User bias can unintentionally influence or limit the fairness or accuracy of the AI's output.

Chatbot: A computer program designed to simulate conversation with human users.

Chunks: Segments of input data that an AI system processes in parts, due to limited memory (context window).

Classification: A process where the AI compares input data to its training examples to assign it a category, such as spam or not spam.

Closed-loop system: An AI system that does not learn after deployment. It produces output based only on the data it was trained on.

Context window: The maximum amount of information an AI model can consider at one time. If the input exceeds this limit, it must be split into smaller chunks to be processed effectively.

Data: Information that an AI system uses. It can include text, images, video, multi-media documents or audio.

Decision making: AI uses patterns in data to decide how to act, such as a self-driving car choosing to brake or turn.

Deep learning: A form of machine learning that uses many layers in a neural network to understand and learn from complex data.

Embedding: A mathematical representation of data (like words or images) as numbers, so that AI can compare meanings or patterns.

Embedding vector: A string of numbers that represents a token (like a word or image fragment).

Ethical AI: The practice of designing and using AI systems in ways that promote fairness, accountability, privacy and the wellbeing of people and society.

Generative AI (GenAI): AI that creates new content (text, images, video or music) based on patterns it has learned.

Guardrails: In AI terms, guardrails are the instructions in the algorithm that prevent the AI from outputting content that is unwanted.

Hallucinations: When an AI system generates information that looks accurate but is false or made up, like fake quotes or incorrect facts.

Input: The information you give to an AI system to process, such as a sentence, image, audio clip or data set.

Large language model (LLM): A type of AI trained on massive amounts of text data to understand and generate human-like language.

Machine learning (ML): A subset of AI where systems learn from data to improve over time, rather than being explicitly programmed for every task.

Glossary

Machine vision: AI that uses cameras or sensors to 'see' the world. It is common in devices like self-driving cars or smart vacuum cleaners.

Multimodal AI: AI that can handle more than one type of input, such as text, images, audio, or video, rather than just one form of media.

Natural language processing (NLP): Technology that allows AI to read, understand and generate human language, such as voice assistants or translation tools.

Neural networks: AI systems are made of layers of nodes that work like brain neurons. As the embedding vectors pass through each layer, they are updated to better reflect their context.

Open-loop system: An AI that continues learning after it's been deployed by adapting to new inputs, such as recommendation engines on streaming services.

Output: What the AI system produces after processing, like a response, image, decision or recommendation.

Prompt: A question, command, or input you give to an AI tool to guide its response. The better the prompt, the better the output.

Prompt engineering: The skill of writing effective, clear prompts that guide AI systems to produce better, more accurate or more useful outputs.

RAG (Retrieval-augmented generation): An AI method where the system pulls relevant data first, then uses a language model to generate content based on that information.

Red team: The name for the group of people who test an AI system and deliberately try to cause it to generate biased, offensive or unethical outputs.

Supervised learning: Machine learning where the AI is trained with clearly labelled examples. It learns what correct answers look like.

Token limit: The maximum number of tokens an AI model can process in one request. If the input exceeds this limit, it must be trimmed or split.

Tokens: The tiny units AI breaks input into, such as parts of words, images or sounds, which are then turned into numbers for processing.

Training data: The examples an AI system learns from. The quality and diversity of this data shape the AI's performance and accuracy.

Transformer layers: The building blocks of modern AI models. These layers allow the model to pay attention to relationships between different parts of the input, helping it understand meaning and structure.

Unsupervised learning: A type of machine learning where the AI finds patterns in unlabelled data. For example, YouTube or Spotify recommendations based on your viewing or listening history.





**COOL
.ORG**