



TQI Accreditation
2026, 3 hours



Generative AI Introduction

Questions? Email
bsssenquiries@act.gov.au

Topics in this course	
What is AI and how does it work?	What is generative AI? How does generative AI work? What can generative AI produce?
What are some of the responsibilities and key documents I should know about?	What kind of things can generative AI do in schools? Responsibilities: changed and established Documents to know
AI agnosticism and refusal	Reasons people may refuse AI Environmental Training Data Justice How to deal with AI refusal

HOW DOES THIS ASYNCHRONOUS PL WORK?


This course has a range of reading and activities within a digital workbook. You need to go through these activities and reflect on your practice, either as an individual or as part of a faculty group. Activities are marked out as WATCH, READ, and REFLECT.

The assessment for the course is made up of a multiple-choice quiz to check understanding. This is not intended to be onerous, but to check that the participants have undertaken the course as written.

We hope that you enjoy this course, and that it helps you to reflect on AI in education. We acknowledge that there is no “silver bullet” or “one-size-fits-all” approach, and that different people have different levels of comfort with the technology.

Yours,

The ACT Office of the BSSS

EXTRA CREDIT	
	<p>This course has suggested “extra credit” resources. These are extra readings, websites or videos aimed at enhancing your understanding of content or issues, but they will take you longer than the accredited time for the course. You will see these readings marked out by a box like this one.</p> <p>We suggest that you log interaction with these “extra credit” materials as teacher-identified professional reading. An extra credit log is included so that you can easily keep track of which ones you spent time on, so that when talking with your PL approver your evidence is easy to find.</p>

BSSS Asynchronous AI Courses

AI Introduction

INTRODUCTION

Unless you're happily living under an extremely large rock, you will have heard of AI and you are likely acutely aware of AI's impact in the workplace and in schools.

AI is an abbreviation for "artificial intelligence". We use artificial intelligence all the time in our daily lives. If you open your phone with Face ID, access content on a social media or streaming app feed, use spell check, get frustrated with an ad seeming to follow you around, use a smart home device, or the maps app to check the traffic or route you need to take, you've used AI. There are different types of AI, and not all AIs are equal, but one thing that is common between them is that they are increasingly infiltrating our lives, workspaces, and schools.

The kind of AI that has captured the popular imagination in recent years is called generative AI. This is AI that creates something new: new text, new pictures, new code, or new video. Programs mentioned in the popular press include ChatGPT, Bing AI, DALL-E, Midjourney, Sora, and Gemini.

The big change that happened in early 2023 was the easy availability of a type of generative AI called a Large Language Model. A Large Language Model is a type of algorithm that uses a lot of data – this could be all of Wikipedia, or all of a large code database like Stack Overflow. When asked to generate content, the model uses the data to predict what it should generate; what is most likely to be an acceptable response. AI companies work to finesse these models so that they give a good answer most of the time. Open AI, who made ChatGPT, liken the process to 'training a dog'.

Large Language models are why people are worried about plagiarism. AI like ChatGPT and Bing AI can be used to generate sensible, literate answers to questions, and because these answers are generated, not copied, it is very hard for teachers – and even for automated anti-plagiarism tools -- to "pick up" when students are presenting AI generated content as their own. See the unit on Academic Integrity for more information.

It's okay to mourn



Generative AI is a paradigm shift that has and will affect the way we educate; it can't be ignored.

Some of our hard-won skills are now redundant, and new skills need to be learned.

Students are using generative AI at all stages of learning, including when it is not appropriate to do so.

Please approach this course with an open mind; we are all at different stages in the journey. If you enjoy using AI for education – great! If you're less enthused, you are also welcome. This course is designed to let you make up your own mind.

WHAT IS GENERATIVE AI?

Generative AI is AI that creates an output (e.g., a picture, some words). It is the primary focus of this PL, as it is the most intrusive in the current educational landscape. The uptake has been swifter than many other recent technologies; for example, ChatGPT hit 100 million users after 2 months of being made public. This feat took Facebook and Spotify ~4.5 years. ([Source](#))

The output of generative AI can be:

- Made up from its training data
- Made up from a combination of its training data and internet searches
- A “hallucination” that is confidently expressed but in no way accurate

Programs include ChatGPT, Bing, Midjourney, Dall-E, Sora, Quillbot, Poe, Gemini, Snapchat, Copilot, Claude, Perplexity and many more.

HOW DOES IT WORK?

WATCH: Although this video is a few years old now, it has a very clear and sensible explanation of concepts like neural networks.

AI Art, Explained: <https://youtu.be/SVcsDDABEkM>

In short:

- A computer model is given a lot of data to analyse and learn from.
- The data is broken into small parts that make sense to the computer. These are usually called “tokens”.
- The computer model makes a simply massive number of connections between the tokens, more than a human can make. The term for the imaginary space holding these connections is “latent space”.
- Human or machine learning is used to refine the model and its connections.
- This is how it learns that a token like “dis” could be part of “disagree”, “disappear”, “distinct”, “TARDIS”, “Disney”, “Jadis”, or “look at dis!” and how to tell which one is the most likely to be meant in a particular context.
- This means the model doesn’t ‘understand’ in the way we think of understanding, nor does it ‘copy’ in the way we think of copying
- There are additional layers of programming and protection in place, too.



EXTRA CREDIT

AI and education: guidance for policy-makers:
<https://unesdoc.unesco.org/ark:/48223/pf0000376709>

HOW A COMMERCIAL MODEL OF GENERATIVE AI WORKS WHEN YOU ASK IT SOMETHING:

- The model (e.g., ChatGPT or Gemini) breaks down a large amount of complex information into small bits that are referred to as “tokens”
- You put in a prompt. The program checks that the prompt meets its guidelines, and converts your prompt into something that the model will understand.
- The model uses probability to work out which tokens should go together to make the answer.
- The program checks that the output meets its guidelines, and sometimes converts the outcome into something that you will understand
- The program shows you the output (or doesn't)

Example: Elara and the wizard school

Elara is the witch character made by a young woman who wanted to see her imaginary witch attending wizard school. The white box shows what the user put in, and what ChatGPT returned:

You
Draw Elara in a wizard school with stone walls, dragons, school uniforms and a lake out the front, please!

ChatGPT



Here's the illustration of Elara at a wizard school. The setting includes stone walls with dragon motifs and a view of a lake through large windows. Elara is in the school's uniform, capturing her excitement and curiosity in this magical and scholarly environment.

This dark box shows the prompt that ChatGPT created to send to the model. Notice how much more complex it is!

Prompt


An illustration of Elara, the young witch with curly brown hair and freckles, in a wizard school setting. The school has stone walls and is adorned with motifs of dragons. Elara is dressed in the school's uniform, showcasing a blend of traditional wizard attire and youthful style. The school overlooks a serene lake, visible through large, arched windows. The atmosphere is magical and scholarly, with a sense of adventure and learning in the air. The scene captures Elara's excitement and curiosity as she embarks on her magical education journey.

 Copy

Example: Lila and the Leopard

Lila is Elara's sister. This image shows the guardrails kicking in when Lila prompted a leopard "gobbling up her dinner" – likely because images showing leopards "gobbling" could create a picture with blood and gore.

You
Can you please now do one of Lila's leopard familiar gobbling up her dinner while she talks about spells with her sister. The raccoon is in the background giving it the stink-eye.

ChatGPT
 Error creating image.
 I'm sorry, but I encountered issues generating the image you requested. Unfortunately, I'm unable to provide the illustration at this moment. If there's anything else I can assist you with or if you have another request, please let me know.

You
No worries! Can you do a picture of Lila the witch and her leopard familiar having dinner? The leopard has stolen Lila's dinner and she hasn't noticed.

ChatGPT



Here's the whimsical illustration of Lila the witch having dinner with her leopard familiar, who has sneakily stolen Lila's dinner without her noticing. The scene is set in their cozy kitchen, capturing the playful and mischievous nature of the leopard and the unsuspecting Lila.

Different programs have different guardrails. An education-specific program like NSWEdUChat has guardrails to prevent students offloading their work onto the AI.

READ this short article: <https://education.nsw.gov.au/teaching-and-learning/education-for-a-changing-world/nsweduchat/safety-and-optimisation> to see what kind of guardrails are in place in the NSW education chatbot.

WHAT CAN GENERATIVE AI PRODUCE?

Please note: AI technologies are evolving rapidly, and many products offer more than one of these functions.

Text

Outputs writing in response to a prompt. Can analyse text or images to look for patterns or offer advice. Can be asked to output text in different styles and voices, including the voice of the user, and to incorporate references and quotes. Can be asked to “improve” text for the user and will do so.

Images

Outputs an image in response to a prompt. Can work in a variety of styles, add, or remove elements from an image, including expanding on an existing image. Referred to as a ‘deepfake’ when a person’s image is realistically mapped into a situation that they have never been in.

Analysis and Code

Outputs analysis of data (e.g., graphs, qualitative analysis of words) and/or computer code. This analysis can include the analysis of a person’s study or work habits to aid self-reflection, or the analysis of a piece of work and advice for improvement. This analysis may also be of large groups of data such as student data or workforce data, such as resume-ranking programs.

Design

Outputs visual design such as layout, such as automatically arranging content, choosing colour and font schemes, creating images. A common addition to slideshow or graphic design programs.

Video and Audio

Outputs a video and/or audio clip. This may be entirely AI generated, generated based on an image or series of images, or a ‘deepfake’ where a person’s image or voice is mapped onto a real or digital clip that shows them realistically doing or saying something that has never occurred. Music can also be generated via AI and in recent months has charted at number 1 on Spotify.

Organisation

Technology that is being woven into multimodal tools that allows an AI to offer organisation assistance such as automatically adding follow-up meetings/to-dos to a person’s calendar based on what is said in a meeting, automatically write summaries or minutes of meetings, or undertaking administrative tasks such as planning and booking a travel itinerary.

WHAT KIND OF USES DOES AI HAVE IN SCHOOLS?

Teachers can use generative AI to:

- Generate lesson and curriculum plans
- Generate ideas for innovative activities and assessments
- Analyse curriculum materials as part of constructing an assessment
- Create exemplars, datasets, examples, arguments and other lesson-specific texts
- Visualise ideas for students
- Quickly format or transform resources
- Differentiate class materials
- Generate polished text from a broad/draft description

Students can use generative AI to:

- Get specific, personalised feedback
- Differentiate instructions or explanations, including making documents accessible for students with disability
- Format or transform assessment work (e.g., slideshows for a speech)
- Visualise and ideate (e.g., test concepts in a design class)
- Summarise articles, texts, and videos
- Change the lexical density of text for ease of understanding
- Converse with an AI to finesse understanding
- Brainstorm or discuss ideas for a specific task
- Create a scaffold for a specific task
- Create a study plan or calendar
- Generate study advice and hints

WHAT ARE TEACHERS' RESPONSIBILITIES WHEN USING GENERATIVE AI?

Under the Australian Framework and AISTL standards, most responsibilities are unchanged.


- Monitoring student data and behaviour for trends and acting where needed to address problems
- Explicitly teaching critical thinking and creativity
- Using a wide range of texts, examples, content and contexts, to expose students to a range of ideas and minimise bias
- Explicitly teaching academic integrity and the attribution of sources
- Explicitly teaching what knowledge, ethics and integrity look like in the discipline/s taught
- Compliance with privacy legislation
- Compliance with other Australian laws

New or changed responsibilities:

- Explicitly teaching about the inclusive use of AI technologies and algorithmic biases in the discipline taught
- Teachers teach students about prompting, including the kinds of prompts and information that may breach privacy
- Where generative AI is used in schools (e.g., for marking work or writing reports), teachers and school leaders monitor, and quality assure the decisions
- Teachers do not enter student work or information into AI without consent
- Clearly stating AI policies on assignments, with consistent consequences for breaches. This is likely to be realised at a school level.
- Teachers maintain up to date knowledge about technology, including generative AI. There is an implied responsibility from schools, systems, sectors, and government to provide quality-assured information to teachers.

WHAT CAN/CAN'T THE BSSS HELP WITH?

- Not all uses of AI in school contexts will be permitted by all sectors and/or schools.
- Please note that the Board of Senior Secondary Studies (BSSS) does not provide technology to schools and students, nor does it set policies about the acceptable use of technology in sectors and schools.
- We can help people working in the ACT Senior Secondary System to discuss ideas and proposals, to clarify policy matters, and when AI is involved in an appeal. We cannot restrict or permit specific uses of AI.
- We can help with validation tasks, academic integrity, and penalties.

	<h4>Policy Links</h4> <p>The policies that the BSSS have which may come into play when you are working with AI in the ACT Senior Secondary System include:</p> <ul style="list-style-type: none"> • Academic Integrity Policy • Validation Policy • Special Provisions Policy • Ethical Guidelines for Assessment <p>If you need specific advice, your Certification coordinator, Deputy Principal, or other liaison with the office of the BSSS can organise this.</p> <p>Your organisation may also have specific policy about AI; for example, the ACT Education Directorate has highly specific policy about the practical use of AI, which is not related to or regulated by the BSSS. Always check your organisation's policies before implementing any AI tool.</p>
---	---

IMPORTANT DOCUMENTS TO BE AWARE OF IN 2026

MUST BE READ IN THEIR ENTIRETY AND IMPLEMENTED:

READ: Spend some time now exploring both of these documents. They are important for your understanding of your responsibilities as broadly defined in the Australian context.

The Australian Framework for Generative AI in Schools:

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

NB: [we made a guide for this Framework](#) that may be of interest.

Smartcopying: Using Generative AI Platforms in Schools

<https://smartcopying.edu.au/using-generative-ai-platforms-in-schools/>



GOOD TO BE AWARE OF:

Australian Government: Study Buddy or Influencer:

https://www.aph.gov.au/Parliamentary_Business/Committees/House/Employment_Education_and_Training/~link.aspx?_id=BA9AFC69252E497CB2C885200812B307&_z=z

ACARA: Curriculum Connections AI:

<https://www.australiancurriculum.edu.au/curriculum-information/understand-this-curriculum-connection/artificial-intelligence>

AITSL: Evaluating the evidence for educational technology

Part 1 <https://www.aitsl.edu.au/research/spotlights/evaluating-the-evidence-for-educational-technology-part-1-the-technologies>


Part 2 <https://www.aitsl.edu.au/research/spotlights/evaluating-the-evidence-for-educational-technology-part-2-enabling-learning>

WHY DO SOME PEOPLE NOT WANT TO USE AI?

AI is basically unavoidable in modern society. It remains to be seen if AI is a bubble or a fad, but for now it has moved from being technology that had to be deliberately sought out, to being included in everyday applications by default.

You personally may feel agnostic about or reject the overt or extraneous use of AI; equally, your students may also feel agnostic or outright hostile towards AI. It is good to have a sense of the issues, so that if they come up in class, you are able to discuss them, and so that you know where you stand. This section is included here as a way to acknowledge that there are issues with the unconsidered adoption of AI; a healthy scepticism and considered opinion are important when dealing with new technologies.

Please note that issues around academic integrity and human wellbeing stemming from the use of AI are covered in other courses, as they are broad, complex, and deserving of more time and energy. This introductory module looks at reasons inherent in the creation of AI that may affect a person's choices, particularly those which may be hidden to a regular user.

	<p>EXTRA CREDIT</p> <p>'AI veganism': Some people's issues with AI parallel vegans' concerns about diet https://theconversation.com/ai-veganism-some-peoples-issues-with-ai-parallel-vegans-concerns-about-diet-260277</p>
--	---

ENVIRONMENTAL CONCERNS

Electricity and water use by AI is high, and the information that is available is about the exact usage numbers is sometimes conflicting. Most discussion about environmental concerns centres around a cost-benefit analysis; does the good (or potential for good) from the use of the AI technology outweigh the negative effects on the planet. These discussions do not always consider supply-chain neutrality. These environmental concerns are a great way to work with the "sustainability" cross-curriculum priority, in addition to ICT and critical literacy skills.

READ:

AI's excessive water consumption threatens to drown out its environmental contributions: <https://theconversation.com/ais-excessive-water-consumption-threatens-to-drown-out-its-environmental-contributions-225854>

How much energy does Google's AI use? We did the math: <https://cloud.google.com/blog/products/infrastructure/measuring-the-environmental-impact-of-ai-inference>

Arguments you may encounter in this sphere include a number of logical fallacies, and this kind of argument can be good to use as a proving ground for students' ability to generate and sustain an argument of their own.

Common fallacies you might encounter in this debate:

Fallacy	Description	Example in AI & Environment Debate
Straw Man	Misrepresenting an argument to make it easier to attack.	<i>"People who worry about AI's environmental impact just want to stop technological progress."</i>
Appeal to Nature	Assuming something is good/bad because it's "natural" or "unnatural."	<i>"AI is unnatural, so it must be harmful to the planet."</i>
False Dilemma	Presenting only two options when more exist.	<i>"Either we use AI and destroy the planet, or we reject AI and fall behind."</i>
Slippery Slope	Claiming one step will lead to extreme consequences.	<i>"If we allow AI to grow, soon we'll have no forests left."</i>
Appeal to Ignorance	Arguing something is true because it hasn't been proven false.	<i>"There's no proof AI harms the environment, so it must be safe."</i>
Red Herring	Distracting from the main issue with irrelevant information.	<i>"Why worry about AI's energy use when cars pollute more?"</i>
Bandwagon	Assuming something is good because it's popular.	<i>"Everyone is investing in AI, so it must be environmentally fine."</i>
Hasty Generalisation	Drawing a conclusion from limited evidence.	<i>"One AI company uses renewables, so AI isn't bad for the environment."</i>


TEACHING IDEAS FOR THE DEBATE ABOUT THE ENVIRONMENTAL IMPACT OF AI
Quantitative tasks:

- Examination of existing literature to test the conclusions and discuss how the material is presented – e.g., the scale on graphs, the promotion of some data, the use of image and colour in infographics.
- Use of existing data to explore statistical correlation or other effects.
- A quantitative cost-benefit analysis testing a specific use of AI.

- Teacher-guided data collection to assess the attitudes towards AI in your school community; use data as basis for analysis and drawing conclusions.
- Explicit teaching about how arguments are/are not supported by the use of quantitative data; what conclusions can be drawn from sample sizes, *post hoc* conclusions, etc. In the extra credit reading for this module you can find some teaching materials about the use of graphs and statistics to uphold conclusions.

Verbal tasks:

- Language analysis of the discussion of AI electricity and water data, specifically looking for rhetorical techniques such as emotive language, imperative language, logical fallacies, structure.
- Comparative analysis of Science communication in the scholarly and grey literature: how ideas are developed and dealt with; structure of texts; use of metaphor and other forms of figurative language in both; register and modality of voice. This does not need to be restricted to Science units – note that there are English units that offer opportunity for comparative study of texts.

	EXTRA CREDIT
	<p>How much water does AI consume? The public deserves to know: https://oecd.ai/en/wonk/how-much-water-does-ai-consume</p> <p>Data centres are vital for the future and AI but their environmental footprint can be a problem: https://www.abc.net.au/news/2025-08-27/ai-to-take-up-one-quarter-of-sydney-water-in-a-decade/105700928</p> <p>These extra credit resources are examples of the misleading use of quantitative data and may present an opportunity to explore this concept in more detail.</p> <p>Spurious Correlations https://www.tylervigen.com/spurious-correlations -- (Highly recommended just for fun!) This site uses data mining to make spurious connections between data. It's genuinely funny and worth a visit even if you don't see yourself as mathematically inclined.</p> <p>The course "Calling Bullshit" [sic] is a critical literacy course aimed at US college students. Some of these resources may be of interest to you if you are teaching in this area.</p> <p>The Principle of Proportional Ink https://callingbullshit.org/tools/tools_proportional_ink.html</p> <p>Misleading axes on graphs https://callingbullshit.org/tools/tools_misleading_axes.html</p> <p>Videos of lectures supporting the above materials: https://callingbullshit.org/videos.html</p>

TRAINING DATA

This section will talk about some of the general ethical concerns about the training of AI. The separate course about Academic Integrity, and the separate course about Human and Social Wellbeing, go into more detail about some specific ethical issues around algorithmic biases, scams, mental wellbeing and the use of AI by students to breach academic integrity.

TEXT AND DATA MINING

The issue of text and data mining, and training on copyrighted works, is not covered specifically in current copyright legislation, and many creators have voiced concern about this. There is a related concern, too, that AI will take Arts jobs from humans, from writing to painting to graphic design.

Generative models need huge amounts of data as part of their creation. This data may come from “scraping” websites on the internet – basically, copying a whole lot of data that isn’t protected from being downloaded. This might be your comments on Blogger, or the photos you put up on Flickr and DeviantArt back in the day. Other sources are purchasing commercially available databases, or even illicit databases that contain copies of otherwise protected content (e.g., copyrighted books, paywalled newspapers). Other developers will legally purchase databases of content from the site owners (e.g., [Reddit sold its content](#)) or may already own that content due to the agreements you quickly clicked past when signing up (e.g., Meta’s rights over Instagram and Facebook posts).

This is an extremely difficult issue to resolve, because the way that generative AI works is so different from the way traditional knowledge and copyright has worked. There is no historical legal precedent for the ways AI uses data, and this is actively being pursued by individuals and governments as society tries to respond. Most arguments focus on the copying or manipulation of data as part of the process, and the ability of AI to make products that are in direct competition with the works they have copied.

THE USE OF DATA FOR OTHER THAN ITS INTENDED PURPOSE

Training databases also raise ethical concerns for some people based on privacy. These databases are often made by scraping huge amounts of the internet and contain content that is being used for purposes other than which it was originally intended. One of the more famous incidences of this was the 2024 [revelation by Human Rights Watch](#) that the LAION5 database included photographs of Australian children.

Content note: These links mention the use of AI tools to create child abuse material, or the inclusion of illegal material in large training datasets. No graphic material is discussed or shown and the articles come from professional organisations.

READ: Australia: Children’s Personal Photos Misused to Power AI Tools

<https://www.hrw.org/news/2024/07/03/australia-childrens-personal-photos-misused-power-ai-tools>

The article brings up another problem in the training data – the use of material without consideration for cultural knowledge or cultural safety. We will discuss this further in the “Human and Social Wellbeing” course.

READ: This article from 2024 discusses LAION removing some photographs of children from its dataset, including 362 Australian Children:


<https://www.hrw.org/news/2024/09/03/720-australian-and-brazilian-children-better-protected-ai-misuse>

OUTSOURCING OF WORK TO LOW-PAID WORKERS

Some people may feel uncomfortable with the way AI is trained. Data annotators – people who are paid to categorise data and label it – are a hidden but essential part of AI. Many of the companies contracting this work are based in Africa, particularly in Kenya, Nigeria, and South Africa.

READ: this article from the BBC to get some sense of the issues: Why Big Tech pays poor Kenyans to teach self-driving cars <https://www.bbc.com/news/technology-46055595>

EXISTENTIAL CONCERNS

	EXTRA CREDIT
	<p>US Copyright: https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf (nb: this document has some really fantastic explanations of what AI is and how it works, if you want a deeper dive)</p> <p>Ethics of Artificial Intelligence https://www.unesco.org/en/artificial-intelligence/recommendation-ethics?hub=32618</p> <p>Africa's data workers are being exploited by foreign tech firms – 4 ways to protect them https://theconversation.com/africas-data-workers-are-being-exploited-by-foreign-tech-firms-4-ways-to-protect-them-252957</p> <p>Articles that have some sensitive content in them:</p> <p>Meet Mercy and Anita – the African workers driving the AI revolution, for just over a dollar an hour: https://www.theguardian.com/technology/article/2024/jul/06/mercy-anita-african-workers-ai-artificial-intelligence-exploitation-feeding-machine (content note: this article contains a description of a content moderator unexpectedly learning about the death of a family member through her work moderating)</p> <p>Training AI takes heavy toll on Kenyans working for \$2 an hour 60 Minutes USA: https://www.youtube.com/watch?v=qZS50KXjAX0 (content note: this video briefly discusses matters around moderation of adult content from ~9 minutes)</p> <p>Releasing Re-Laion-5b: Transparent Iteration On Laion-5b With Additional Safety Fixes: https://laion.ai/blog/relaion-5b/ (content note: this article makes substantial but non-detailed reference to the inclusion of abuse material in a large training dataset).</p>

Some users of AI are concerned about AI's possible impact on human life and wellbeing. We will discover many of the risks associated to human wellbeing and society in the module about human and social wellbeing.


You may run into two distinct mindsets at the extremes of belief about the existential risks of AI: "AI Doomers", who foresee that AI will have such a negative effect on humanity as to become apocalyptic, and "AI Accelerationists" who look to AI as the solution to global problems and promoting the longevity of humankind.

LISTEN: Choose between these two NPR podcasts:

- If you feel pretty confident with the ideas around AI and "doomerism", listen to this shorter podcast: <https://www.npr.org/2025/09/24/nx-s1-5501544/ai-doomers-superintelligence-apocalypse> (7 minutes)
- If this is new to you or you want a deeper discussion and understanding, listen to this longer explainer: <https://one.npr.org/?sharedMediaId=nx-s1-5551312:nx-s1-mx-5731549> (21 minutes)

READ: the **Domain Taxonomy of AI Risks** on this website, using the accordions to open the risk profiles: <https://airisk.mit.edu/>

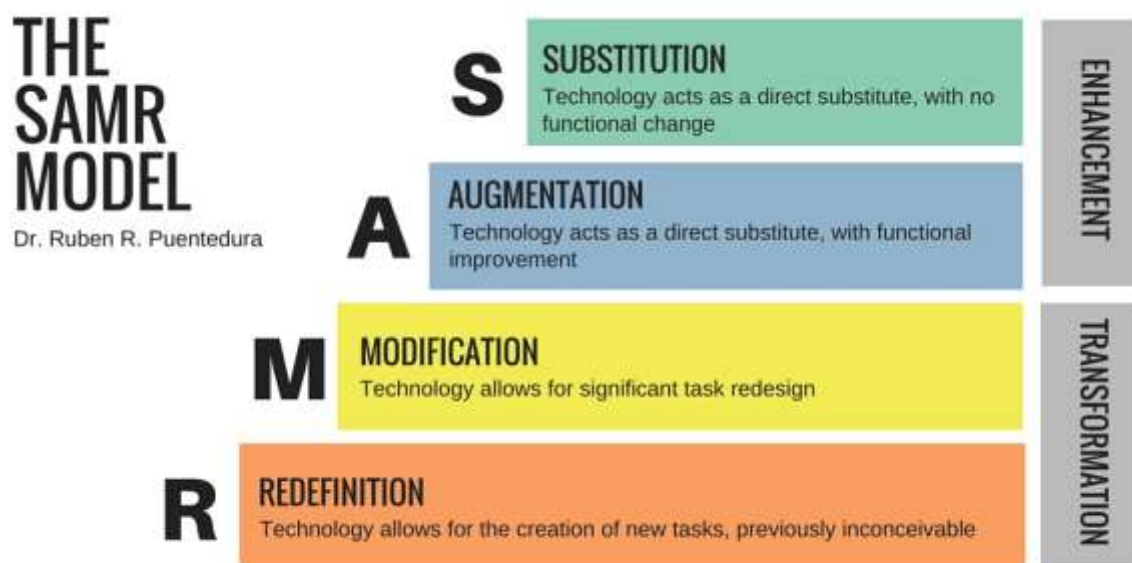
REFLECT: Where do you fall on the Doomer/Accelerationist spectrum? Why? What evidence would change your mind?

EXTRA CREDIT	
	<p>AI could be the ultimate test of Australia's democracy https://www.abc.net.au/news/2025-11-11/ai-could-ultimate-test-australia-democracy-boyer-lectures/105991022</p> <p>Is AI dominance inevitable? A technology ethicist says no, actually: https://theconversation.com/is-ai-dominance-inevitable-a-technology-ethicist-says-no-actually-240088</p> <p>Ascribing consciousness to artificial intelligence: human-AI interaction and its carry-over effects on human-human interaction https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2024.1322781/full</p>

HOW TO DEAL WITH AI RELUCTANCE?

Acknowledge that this is a real and valid response to AI, but one that may be challenging to implement in full due to the embedded nature of AI tools in hidden everyday use. Explain to students why you want them to use or not use particular technologies – “because I said so” is never an acceptable answer. Being able to articulate the reasons for the use of technology will help people to feel okay about using it.


A tool that may be useful in framing how you are using AI technology is Puentedura’s SAMR model, which provides a rubric for introducing technology into educational contexts. It describes how new technologies “fit” into learning.



The SAMR Model.jpg. (2024, January 9). *Wikimedia Commons*. Retrieved November 19, 2025, from https://commons.wikimedia.org/w/index.php?title=File:The_SAMR_Model.jpg&oldid=840044163.

Licensed under [Creative Commons](https://creativecommons.org/licenses/by-sa/4.0/) Attribution-ShareAlike 4.0.

In the SAMR model, the use of technology in education is framed in four ways. When a new technology is introduced, SAMR provides a model for the adoption of that technology in a way that positively affects student learning. It is not necessarily desirable to reach the level of “redefinition” in every task and for every teacher.

	<p>EXTRA CREDIT</p> <p>Evidence for ESSA: https://www.evidenceforessa.org/ This is an American website dedicated to surveying the evidence for educational programs, including EdTech programs. Hopefully Australia will have something like this one day too!</p>
---	--

REFLECT: Answer the following considering the SAMR model and your own work as a teacher. Note that *technology* in this activity is not restricted to AI – you might think of learning management systems, e-books, computers, the internet.

Substitution: Which tasks have you set for students where technology is a direct substitute, without functional change to the task?

This might be making a recording of a class available online or using a digital version of a worksheet with no frills (i.e., the same kind of thing you used to give them on paper, just on a computer).

Augmentation: Which tasks have you set for students where technology is a direct substitute, with functional improvement to the task?

This might be the use of a voice function to put notes on a task for a student; an embedded video in a worksheet; a digital document that enables text to speech (or vice versa).

Modification: Which tasks have you set for students where technology enabled task redesign?

If you have been teaching for a while, think back to the kind of changes you made during the pandemic to modify your existing programs of learning to deal with teaching moving to online platforms. Suddenly, class discussions became breakout rooms, or asynchronous posts on a class forum board. These were probably tasks where technology enabled task redesign (with perhaps some augmentation mixed in).

Redefinition: Which tasks have you set for students where technology created new tasks?

This might be a student putting together a video essay where they splice their own knowledge and presentation style with clips from a film – a type of essay that does not exist without technology. Perhaps this is an adaptive task, where students who are able to answer some questions progress differently. Perhaps it is a video task where a student filmed their own training or dance session and then recorded a reflective commentary.

WATCH: for an explanation of how the model can be used, watch this video: <https://www.youtube.com/watch?v=ZQTx2UQQvbU> (note that the video description also includes links to a version without music if you find the music distracting).

REFLECT:

Where do you feel your practice is on the SAMR model with using AI at the moment? Is it a substitution, or a modification?

How do you think your students are using AI? Are they augmenting their learning, or something else?

NEXT STEPS

To be counted as having completed the work for this course, you need to answer the questions in the form here:

[BSSS Introduction to AI Workshop 1 – Assessment](#)

This part of our suite of AI courses was intended to introduce you to the technology, and to some of the discourse surrounding the technology. We hope you have enjoyed the course.

EXTRA CREDIT CHECKLIST: INTRODUCTION TO AI

If you are planning to use the extra credit readings as part of your professional reading for accreditation, you can check off the readings you've done on the list below. You will need to log this extra reading as teacher-identified professional learning.



Item	Description	Read
AI and education: guidance for policy-makers	Comprehensive guide by UNESCO introducing AI	
Australian Government: Study Buddy or Influencer	Report on the inquiry into AI and education in Australia	
ACARA: Curriculum Connections AI	ACARA's guidance for the introduction of AI across the curriculum in v9	
AITSL: Evaluating the evidence for educational technology	Comprehensive review by AITSL about the evidence-base for EdTech.	
Part 1 Part 2		
'AI veganism': Some people's issues with AI parallel vegans' concerns about diet	Article giving an overview of why people may have concerns about AI.	
How much water does AI consume? The public deserves to know	Article from the OECD questioning AI's water use and offering some explanation of why the tech uses a lot of water.	
Data centres are vital for the future and AI but their environmental footprint can be a problem	ABC article canvassing issues about AI and data centre water use.	
Spurious Correlations	This site uses data mining to make spurious connections between data. An excellent lesson in data interpretation.	
Calling Bullshit The Principle of Proportional Ink Misleading axes on graphs Videos of lectures supporting the above materials:	A critical literacy course aimed at US college students, focusing particularly on data literacy.	
US Copyright report on AI	A report about the status of copyright and AI in the USA. Also contains a good explanation of the technology itself.	

Item	Description	Read
Ethics of Artificial Intelligence	UNESCO's site exploring the ethics of AI	
Africa's data workers are being exploited by foreign tech firms – 4 ways to protect them	A call to action from The Conversation.	
Meet Mercy and Anita – the African workers driving the AI revolution, for just over a dollar an hour:	From The Guardian. This article ranges from AI to content moderation. Content note: this article contains a description of a content moderator unexpectedly learning about the death of a family member through her work moderating, which may distress some readers.	
Training AI takes heavy toll on Kenyans working for \$2 an hour 60 Minutes USA	Video interviewing workers about their work both in training AI and in content moderation. Content note: this video briefly discusses matters around moderation of adult content from ~9 minutes.	
Releasing Re-Laion-5b: Transparent Iteration On Laion-5b With Additional Safety Fixes	From the owners of the LAION-5b database, an explanation about what they did regarding the inclusion of sensitive content in the database. Content note: this article makes substantial but non-detailed reference to the inclusion of abuse material in a large training dataset.	
AI could be the ultimate test of Australia's democracy	2025 Boyer Lecture text, eloquently summarising most of the key issues.	
Is AI dominance inevitable? A technology ethicist says no, actually	If you are concerned about AI taking over, this article from The Conversation might make you feel a little better.	
Ascribing consciousness to artificial intelligence: human-AI interaction and its carry-over effects on human-human interaction	Interesting scholarly article exploring not whether we should ascribe consciousness to AI, but what happens when we do.	
Evidence for ESSA	This is an American website dedicated to surveying the evidence for educational programs, including EdTech programs. Due to the global reach of these programs, some that you use may appear on here!	