



AUSTRALIA

Lesson 3: How do Machines Create Things?

LESSON PLAN

DAYOFIAUSTRALIA.COM

Lesson 3: How do Machines Create Things? | 60 mins

Lesson Summary

In this lesson, students unpack the realm of generative AI, with an emphasis on large language models (LLMs). Through interactive activities, they discern the applications of generative AI, weighing its risks and benefits. Students engage in group discussions on how AI generates text and experiment with a predictive text algorithm game. By delving into case studies and getting hands-on with a large language model, they tackle the complexities and implications of this exciting new AI technology.

Objectives

- Students will explore **Interaction**, the fourth **Big Idea of AI**
- Students will learn how generative AI is different from other forms of AI
- Students will explore the applications of generative AI
- Students will consider the risks and benefits of using generative AI
- Students will practise responsible use of generative AI

Curriculum Alignment

This lesson is linked to the following [Australian Curriculum \(Version 9\)](#) content descriptors:

- **Years 5 and 6**
 - [AC9TDI6K01](#): investigate the main internal components of common digital systems and their function.
 - [AC9TDI6K03](#): explain how digital systems represent all data using numbers.
 - [AC9TDI6P10](#): explain the creation and permanence of their digital footprint and consider privacy when collecting user data.
- **Years 7 and 8**
 - [AC9M8ST02](#): analyse and report on the distribution of data from primary and secondary sources using random and non-random sampling techniques to select and study samples.
 - [AC9TDI8K03](#): explain how digital systems represent all data using numbers.
 - [AC9TDI8P11](#): select and use a range of digital tools efficiently, including unfamiliar features, to create, locate and communicate content, consistently applying common conventions.
 - [AC9TDI8P12](#): select and use a range of digital tools efficiently and responsibly to share content online, and plan and manage individual and collaborative agile projects.
 - [AC9TDI8P14](#): investigate and manage the digital footprint existing systems and student solutions collect and assess if the data is essential to their purpose.
- **Years 9 and 10**
 - [AC9M9ST01](#): analyse reports of surveys in digital media and elsewhere for information on how data was obtained to estimate population means and medians.

- [AC9S9H04](#): examine how the values and needs of society influence the focus of scientific research.

Vocabulary

- **Generative**, n. capable of making or producing new things that didn't exist before

Resources

- [Presentation Slides for lesson](#) (available on website once logged in)
- A whiteboard or chart paper and a marker, for the predictive text game
- A free account with OpenAI, for the teacher only (see **SETUP** note below)
- [ChatGPT](#)

Use of ChatGPT in this lesson

This lesson features an activity where you will facilitate an interactive session and guided discussion using ChatGPT. We have purposefully designed this activity as a teacher-led demonstration and guided discussion for two main reasons: 1) to ensure the safe and responsible use of ChatGPT in line with the [Australian Framework for Generative Artificial Intelligence \(AI\) in Schools](#), and 2) to ensure compliance with [OpenAI's terms of use for ChatGPT](#), which requires people to be **"at least 13 years old to consent to use the Services."**

PLEASE NOTE:

- As of the 1st of April 2024, [ChatGPT no longer requires an account to use](#). You may still be prompted to login, but you can choose to continue without doing so.
- [OpenAI's Terms of Use](#) still apply, meaning students under 13 cannot use the service, and students under 18 require permission from a legal parent or guardian to use the service.
- Data which you provide to ChatGPT via the text prompt interface may be used in training the AI in future. You can opt-out of this behaviour by [following these instructions](#).
- Take care not to include and personally identifying information (PII) about you or your students in your input prompts to ChatGPT, to ensure compliance with the [Privacy Act 1988](#) and other relevant legislation in your state or territory.

As a responsible adult, you can safely guide students through conversations with ChatGPT, taking input from students only as you deem appropriate. We strongly encourage you to pause often during this activity to allow students to ask questions, and give them an opportunity to share their thoughts and feelings around large language models like ChatGPT. Some students may need additional time to reflect on the implications of the technology. If you feel the need, re-emphasise that ChatGPT (and other generative AI models like it) are simply a tool, one that can be used responsibly like any other tool.

Activity Steps

1. **1 min.** Introduce **Human Interaction** as the fourth **Big Idea of AI**.
2. **1 min.** Define the word 'generative' in the context of AI.



Hopefully, everyone should agree that 'generative' means 'able to create or make something' and therefore 'generative AI' is AI that can generate new content.

3. **3 mins.** Examine each example of generative AI and determine what type of content it is generating. In order from left to right:
- DALLE2** is an AI which generates **images**, specifically artwork.
 - ChatGPT** is an AI which generates **text**, specifically human-like conversations.
 - Deepfakes** are AI that generate **video**, specifically replacing human faces with other human faces.

 **NOTE**

Deepfakes can also describe static images rather than motion video, and they can also replace the same human face with an older/younger version of that face, but that distinction is not important here.

4. **4 mins.** In pairs, have students discuss the applications of different forms of generative AI. They should try and come up with 2 examples for each kind of AI: 1 where it is **good**, and 1 where it is **bad**.

 **NOTE**

In the next lesson, we will examine the ethics of AI more closely, and define a simple but formal system of ethics based around harm and benefit. If students have already completed that lesson, you can refer to benefit and harm here instead of good and bad.



Examples of **good** applications:

- Images/art:** creating new artwork in the style of a deceased artist, finishing incomplete works, restoring faded works.
 - Text/conversations:** providing an easy, human-like way to interact with computers for older or less technical people, providing human-like support to customers of a business, creating realistic conversations and dialogue in video games.
 - Video/faces:** ageing an actor up/down so they can continue playing a famous character, fixing poor audio sync or lighting in a film, creating realistic interactions with NPCs (non-player characters) in video games.
5. **2 mins.** Explore the questions 'How does AI generate text? How do you decide what to say next?'



Encourage students to consider what their thought process is when responding to a question, for example. What is the subject of the question? What are they actually being asked? Should the answer be about something that already happened, or something that is going to happen in the future? All of these things come from context information in the question itself.

6. **8 mins.** The students will play a game to act out a predictive text algorithm, by following these steps:

- a. Have students form a circle. If you have a large group of students, you can form 2 or more circles, repeating the following steps for each circle.
- b. Think of a sentence **in secret**. Try not to make a long sentence, as you'll need to remember it later. Don't tell the students what it is yet! Write the **first word** of the sentence on the board.
- c. Choose a random student in the circle to start things off. Beginning with that student, have each student in the circle say **ONE** word out loud. Write each new word on the board for everyone to see, as you go around the circle.



Encourage students to say what they think the **most likely word** is. Try to discourage silly, random answers.

- d. Once you've gone round the circle, read the sentence out loud for the class. If you want, you can continue going around the circle multiple times, to make a longer sentence.
 - e. Tell everyone what your original sentence was. Hopefully, it should be nothing like what the students predicted, but crucially the predicted sentence should still make **some kind of sense** (though not necessarily with perfect grammar).
7. **4 mins.** As a class, unpack how the students just acted out a predictive text algorithm, like phones do. Talk about why it was inaccurate. Of course, they had no idea what your sentence was! Equate this to predictive text on modern phones, which can only do a similarly poor job at creating an entire sentence, because it does not understand how **conversations** are formed. It only knows how **some words** are related and often make sense next to each other. This is due to the phone being trained differently to modern generative AI models like ChatGPT.
 8. **3 mins.** Talk through how students have a **human brain**, which is much better at learning things than their phone. Their own lived experience (conversations they've had, movies they've watched, books they've read, etc.) all create a **very large dataset** that helps their brain learn how humans talk.
 9. **1 min.** Pose the question, what makes generative AI better? How is ChatGPT so much more competent at generating text than our phones?
 10. **2 mins.** The answer lies in the underlying technology beneath OpenAI's ChatGPT, Google's Gemini (formerly Bard), Microsoft's Copilot and other similar products; **Large Language Models**, or **LLMs** for short.
 - a. LLMs are trained on extremely large datasets; much larger than any human could possibly hope to consume in their lifetimes.
 - b. The algorithm behind LLMs acts a little bit like a brain. It's really good at learning what a conversation looks like, beyond a few words!
 11. **2 mins.** Explain what a token is, how it relates to English words. Talk about how tokens are measured as roughly $\frac{3}{4}$ or 0.75 of an English word, using a simple sentence as an example to count tokens.
 12. **1 mins.** Touch on just how big an **LLM** is, comparing the amount of tokens used to train ChatGPT (with GPT-3) compared to the number of words in an average 200 page book (averaging 300 words per page).

 **NOTE**

The measurements here are all very rough. A token is not precisely $\frac{3}{4}$ of an English word; it's simply an average. A 200 page book does not necessarily have 300 words per page; it depends on a range of factors including print margins and font size.

13. **2 mins.** Call out other examples of AI “Chatbots”. ChatGPT isn't the only one! There are other well known examples, like Google's Gemini (formerly Bard) and Microsoft's Copilot. The common factor is that they all use LLMs. This connects in with the case study over the next few slides.



EXTENSION QUESTION: AI companies don't always say exactly what data was used to train their AI. Do you think this information should be open and public?

14. **5 mins. CASE STUDY:** The NSW Department of Education has launched an education-focused version of Microsoft CoPilot in just 16 schools across the state called **NSWEduChat**. This version of Microsoft's LLM is tailored to student use, protecting student data and privacy and providing scoped answers to respond to the NSW syllabus and prevent cheating.
- Ask students what makes **NSWEduChat** different to other “chatbots”? What kind of **dataset** do they think it was trained on?
15. **3 mins.** Touch on some of the key limits of **LLMs** in advance of the next activity.
- LLMs are just generating text.** They are just producing text that looks like a human could have written it.
 - LLMs don't know if something is true.** They don't “know” or understand the things they are saying.
 - LLMs don't always give the same answer.** Even if you give them the same prompt!



EXTENSION QUESTION: What kind of limits might exist if an LLM is trained on text that is English-only? What kind of bias might be involved?

16. **3 mins.** Explain to the students that you are going to guide them through a conversation with ChatGPT. Remind them that the technology is only generating text based on **billions** of lines of text it's seen before. Open a new tab and navigate to chat.openai.com, and login to your OpenAI account that you created earlier.

 **NOTE**

This is a teacher-led activity. Students should not open ChatGPT themselves. As you go through each conversation with ChatGPT, encourage students to reflect on the responses you get from ChatGPT:

- Did ChatGPT understand the prompt correctly?
- Did the response answer the question well?
- Were there any mistakes or inaccuracies in the response?

17. **15 mins.** Start the conversation with ChatGPT. You can choose from one of the example prompts below, or come up with your own.

Example prompts. If you feel you have sufficiently explored the conversation topic, you can start a new conversation with a new prompt. You can replace the bold, highlighted <INPUT> sections with your own input, or input from your students.

- a. Write a short story about <a musician on a magical quest>
 - i. **Follow-up prompt:** What happened next?
- b. Define this word for me: <a complex word e.g. “ostensibly”>
 - i. **Follow-up prompt:** Can you simplify that explanation?
- c. Summarise this text: <a few short paragraphs about a topic of your choice>
 - i. **Follow-up prompt:** Now rephrase the text so a <different target audience, e.g. kindergarten student> can understand it.
- d. How do I <write an essay paragraph / change a car tyre / apply for a job>
 - i. **Follow-up prompt:** Can you give me an example?
- e. I’m interested in <a list of topics or school subjects, e.g. science, maths, reading, video games, sport, music>. What kind of jobs should I apply for?
 - i. **Follow-up prompt:** What subjects should I study to prepare for those jobs?

END OF LESSON PLAN