A: Setting the scene

This first section aims to demystify artificial intelligence (AI) by explaining in a very basic sense what AI is, and how it works. We also consider how and why EdTech is prone to hype, and how we and our learners might resist this hype.

1 What is AI?
2 What is generative AI?
3 AI and language learning
4 AI and creativity
5 Technology and the hype cycle

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

1 What is AI?

The field of artificial intelligence (AI) aims to create digital machines that can carry out tasks that typically need human intelligence. How close are we to human-level artificial intelligence?

Early AI in language teaching

With the amount of hype and hysteria that surrounded the arrival of ChatGPT in late 2022, you'd be forgiven for thinking that AI is a completely new technology for teachers and learners. Not so. The earliest and simplest forms of AI in ELT can be traced back to the 1960s, when CALL (Computer Assisted Language Learning) emerged as an area of study. In these early days of AI, computers could be programmed to provide limited responses to prompts. Computers were large and expensive, and tended to be found in universities. The advent of the personal computer, however, meant that by the late 1980s and early 1990s computers began to appear in schools and in people's homes. Language learning software with simple gap-fill and text reconstruction activities became available. As computing power increased, and computers developed multimedia capabilities, other uses for AI in language learning emerged. This was the heyday of the CD-ROM. By the early 1990s, some language learning software began to integrate voice recognition to support pronunciation. Since then, AI has become more powerful, and technology – especially in the form of mobile devices – has become more ubiquitous.

Narrow versus artificial general intelligence

To understand where AI has come from and where it is going, it is useful to distinguish narrow (or weak) and strong AI – the latter is usually referred to as *artificial general intelligence* or *AGI*. Here an analogy may be helpful. Imagine a chair that is mass-produced in a factory. The machine that assembles the individual parts of the chair

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

follows very specific instructions. The machine cannot decide to create another furniture item – let's say a table – unless it is programmed to do so. When the machine breaks down, it cannot fix itself. It is very good at performing a pre-defined task (assembling a chair) quickly and efficiently, but it cannot solve problems and it cannot do new things or adapt to new situations. This is narrow AI. Now imagine a skilled carpenter who makes wooden chairs by hand. She can create unique chair designs. She gets better at her craft over time, learning from her successes and mistakes. She teaches herself to use new and more sophisticated carpentry tools, and she takes pride in her work. The skilled carpenter represents AGI, which is indistinguishable from human intelligence. AGI can plan, problem-solve and learn, and carry out complex multi-faceted tasks. It displays a human-like level of consciousness while doing so. We are not yet in the phase of AGI, but the goal, computer scientists tell us, is to get there.

In our field, the gap-fill computer programs of the 1980s and 1990s are examples of early – and therefore narrow – AI. Voice recognition software, which was notoriously unreliable in the 1990s, has become increasingly accurate. More recently, we have tools like ChatGPT, which are based on generative AI (see **2**), and can generate content in text, image or multimedia formats. ChatGPT is still considered an example of narrow AI by most researchers, but it represents a significant step towards stronger forms of AI, not least in the way it seems to interact with us in a very personable (i.e., pleasant and friendly) manner. It's useful here to imagine AI on a scale, with narrow AI at one end of the scale, and AGI at the other end. Tools like ChatGPT can give us the impression that we are moving quite fast along the scale from narrow to AGI. However, not everyone thinks we can get all the way to AGI, and not everyone is happy at the prospect of this, but there is no doubt that AI is starting to feel more human-like.

AI and consciousness

What 'human-like' actually means is, unsurprisingly, the subject of hot debate. You may have heard concerns that AI might already be 'conscious'. But how can we know whether AI is conscious or not? When I asked ChatGPT this very question, the answer was clearly no,

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

with ChatGPT pointing out that it is a computer program that has no consciousness, thoughts or feelings.

Although current generative AI systems clearly claim *not* to be conscious, the issue underlies many of the debates around AGI. To address the tricky question of what AGI consciousness is, researchers suggest that we need to be scientific in our approach. For example, in one study, a large team of academic researchers first tried to define what human consciousness is from a range of widely accepted neuroscientific theories (Butlin et al., 2023), although it should be noted that there is plenty of debate over exactly what this consists of (Goff, 2023). The researchers then came up with a number of indicators to describe consciousness, based on these theories. The next step was to compare current AI systems against these indicators to see where they match with human consciousness and where they fall short. There were two interesting findings from this study. The first was that no current AI system fulfilled all of the criteria for human consciousness as defined in the study. The second was that there is no reason why future models of AI can't (at least in theory) fulfill all of these criteria. In other words, the researchers concluded that although we may not be at AGI yet, we could get there in the future.

There are plenty of commentators, however, who strongly disagree that generative AI is, or ever will be, intelligent. One theoretical physicist called ChatGPT a 'glorified tape recorder' (Cao, 2023), for example, while others point out that current forms of generative AI are unable to perform basic maths, are prone to hallucinations (i.e., to making up facts) and overall have a very hazy grasp of reality (Marcus, 2023).

Wherever we may be on the scale from narrow AI to artificial general intelligence, the much-publicised arrival of ChatGPT brought generative AI (see **2**) very much to the forefront of public consciousness. Freely available at the time, ChatGPT reached 100 million users within two months of its launch. In comparison, TikTok took nine months to reach a similar number of users, and Instagram took two and a half years. The unprecedently fast and widespread uptake of ChatGPT had the effect of focusing minds on where AI may be leading us. Generative AI tools like ChatGPT did not come out of nowhere. They were based on years of work in the field of natural language processing. ChatGPT in particular

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

though, seemed to awaken the ELT community to the potential advantages – and challenges – of AI in language teaching. Teachers and learners quickly realised that this was going to be a game-changer for our field. But as we will see in this book, AI encompasses much more than a tool like ChatGPT, and in many ways, the game has already changed.

Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., Deane, G., Fleming, S. M., Frith, C., Ji, X., Kanai, R., Klein, C., Lindsay, G., Michel, M., Mudrik, L., Peters, M. A. K., Schwitzgebel, E., Simon, J. and VanRullen, R. (2023). Consciousness in Artificial Intelligence: Insights from the Science of Consciousness. Available at: https://arxiv.org/abs/2308.08708. Accessed 24 December 2023.

Cao, S. (2023). A. I. Today Is a 'Glorified Tape Recorder,' Says Theoretical Physicist Michio Kaku. *Observer Newspaper*, 15 August 2023. Available at: https://observer.com/2023/08/michio-kaku-ai-chabot/. Accessed 24 December 2023.

Goff, P. (2023). Understanding Consciousness Goes Beyond Exploring Brain Chemistry. *Scientific American*. Available at: https://www.scientificamerican.com/article/understanding-consciousness-goes-beyond-exploring-brain-chemistry/. Accessed 22 January 2024.

Marcus. G. (2023). Reports of the birth of AGI are greatly exaggerated. Blog post. Available at: https://garymarcus.substack.com/p/reports-of-the-birth-of-agi-are-greatly. Accessed 24 December 2023.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

2 What is generative AI?

The appearance of tools like ChatGPT brought the potential benefits and challenges of generative AI into sharp focus for educators. Understandably, teachers have responded to it with both joy and fear. Understanding some of the principles behind generative AI can help demystify it.

Knowledge-based AI

To understand generative AI, it's helpful to first look at earlier types of AI – known as *knowledge-based AI*. The spell-check program is an early and useful example of this. How does a spell-check program typically work? First, a dictionary or lexicon of correctly spelled words (the knowledge base) is defined. Then, a programmer creates explicit rules that tell the spell-check program to compare each word in a text against the words in its knowledge base. If a word is not found in the lexicon, it is considered a potential spelling error. The program then suggests possible corrections based on words in its knowledge base. The corrections are generated by the program applying pre-written rules and algorithms.

These days, most spell-check programs will take context into account when suggesting corrections. For example, imagine that the program encounters the word 'their' in the sentence, 'Their going to the park'. The word 'their' is correctly spelled, but it's not grammatically correct. The program may suggest 'they're' as the correct spelling because it considers the sentence as a whole. Some spell-check programs may offer a choice of corrections for this sentence, for example, by offering not just 'they're', but 'there'. This shows us that there is a rule (or algorithm) that tells the program to suggest words that are phonologically similar when it spots a possible error. By providing a choice, the program is also allowing for human judgement. Other terms you may come across for knowledge-based AI are *predictive AI* or *rule-based AI*. Knowledge-based

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

AI has been used extensively in language learning, including in apps (see **6**), intelligent tutoring systems and chatbots (see **8** and **9**), automated translation and testing (see **15**).

Data-driven AI

As well as knowledge-based AI, we have so-called *data-driven AI*. This approach is based on something called *machine learning*. Simply put, machine learning involves the development of algorithms and models that can learn from data and improve their performance on a specific task *without being explicitly programmed*. More on this below.

Data-driven AI has been around for several decades, too, and it is used in areas like language testing, speech recognition and machine translation. Indeed, knowledge-based and data-driven approaches are often combined. For example, spell-check programs are knowledge-driven, but they are also now based on large datasets of language. This means that they can more accurately identify the context a word is used in, and therefore offer more accurate corrections. If you, like me, have used spell-check programs for years, you'll have noticed how much more helpful they have become over time.

Generative AI

Within the field of data-driven AI, we have *generative AI,* which is trained on massive quantities of user-generated online data. Generative AI can become more knowledgeable and accurate over time, as it receives more data, is retrained and refined, and as internal parameters are adjusted to better process the input data. There are different types of generative AI tools. Some generate text and are based on huge amounts of online texts, including resources like Wikipedia and online books. Text-generation AI tools included ChatGPT, Gemini and Bing at the time of writing, although these are very likely to be joined or replaced by other tools in the future. Other generative AI tools can generate images, videos or sound, and are based on immense datasets of these media; these media and texts are harvested from the internet – often without permission (see **22**). Simply put, generative AI uses the algorithms it develops to teach itself; for example, it can draw conclusions about things that may not have been in the original data, and it can generate new data in the form of images, videos or text.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

Large language models

Text-based generative AI tools are based on so-called *Large Language Models* (LLMs). This is a complex area, but let's try to get a basic understanding of LLMs.

A Large Language Model is a type of artificial intelligence system that is designed to analyse and generate human-like text. It's essentially an advanced computer program that's very good at analysing and producing words in context. LLMs are built upon deep learning techniques and algorithms based on neural networks. The term *neural network* comes from biology, and refers to the way the human brain is made up of billions of interconnected neurons. Neural networks in computing attempt to simulate these biological networks. They are complex mathematical models based on algorithms with billions of parameters that can identify patterns, correlations and relationships in data, and make predictions or decisions based on these data. I find it helpful to think of an artificial neural network as a digital brain made up of tiny decision-making units. These units work together to solve problems or recognise patterns. Information goes in, and the neural network processes it and gives an answer. Artificial neural networks underpin many computer tasks, like recognising pictures and understanding language.

Large Language Models are called 'large' because, as we saw above, they are trained on datasets containing vast amounts of text taken from the internet. The training of LLMs can either be supervised or unsupervised. In supervised training, humans tell the model when the content it produces is right or wrong, or how it can be improved. The model uses this human feedback to adjust its algorithms and improve its performance, making it better at understanding and generating text. This feedback loop helps train the AI to be more accurate, useful and aligned with human preferences and intentions. This approach is called *Reinforcement Learning from Human Feedback* (RLHF), and it requires a significant amount of human labour (something we explore in **14**). Exactly *how* a LLM uses a neural network to teach itself to adapt the parameters of its algorithms is unclear though. This lack of clarity is, understandably, a concern for many computer scientists – and has led some commentators to claim that we are getting close to

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

artificial general intelligence (AGI – see **1**). One thing is clear though – the quality and quantity of training data that generative AI platforms use is important, as is the quality of human feedback they receive in supervised learning. Poor or biased input, or poor or biased feedback, is likely to lead to poor or biased outputs (see **18**).

What this means for language teachers

In short, LLMs learn language patterns, semantics and context from large amounts of internet data, and they use this knowledge to generate text that is coherent, contextually relevant – and very human-like. We're not sure exactly how they do some of the mathematical/algorithmic bits in this process. LLMs don't make many grammar and spelling mistakes and they can sound very knowledgeable and convincing, even if they are not always factually correct (see **1**). Because of their advanced language capabilities, text-generation AI tools based on LLMs are finding their way into language learning in the form of chatbots, translation, content generation and more – as we explore in the rest of this book.

Comparing knowledge-based AI with generative AI can help us understand how the latter is an important developmental step in the field of AI. Knowledge-based AI has been around for decades, and most of us have probably experienced it in our personal lives (for example, in spell-check programs) or with our learners (for example, by encouraging them to use language learning apps). Generative AI however, which can generate realistic, seemingly new content, is a more recent development. As we will see in this book, there is room for both knowledge-based *and* generative approaches in language learning.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

3 AI and language learning

Now that we have a basic understanding of both knowledge-based and data-driven AI (including generative AI), we can consider its potential to support language learning.

Learning a language

When I moved to Spain in my early 20s, I knew not one word of Spanish. In my first year of living in the country, I had a very old car that kept breaking down. I had to learn Spanish fast, not least to describe my latest car problems to the local mechanic. *Gears, handbrake, clutch*, *windscreen wiper* … these low-frequency vocabulary items were vital for me to learn in my first few months. None of these words in Spanish sounded anything like their English equivalent, and I found them very difficult to remember. So, I took a piece of paper, wrote the words in English on one side and the Spanish translation on the other side, and kept the piece of paper in my pocket for several weeks. Every time I put my hand in the pocket, I'd remember the piece of paper, try to remember the Spanish words, and then check whether I was correct by looking at my paper. Within a couple of weeks, I'd managed to memorise all the words, and I remember them to this day.

Although we don't know everything about how languages are learned, we do know several things from decades of second language acquisition research. For example, learning a language entails storing a large number of words and phrases (often called *lexical chunks*) in memory. In my case, it was car vocabulary in Spanish. I used paper, but had it existed at the time, an AI vocabulary app could have helped me with this. Research shows that knowledge-based (see **2**) AI language learning apps can help learners commit new words and phrases to memory, especially if these are grouped together in thematically connected lexical sets. It also helps if words are shown in context (so in sentences rather than as individual unconnected words) and if the app uses *spaced repetition*. Spaced repetition occurs when the learner comes across a target word or phrase many times over increasingly longer time

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

intervals. I unwittingly used spaced repetition by looking regularly at the car vocabulary items on my piece of paper, until, finally, I didn't need to look at it anymore.

To learn a language, we need more than words. We also need grammatical structures, and we need to know how to pronounce things (as I soon discovered when trying to describe issues with the clutch to my mechanic in beginner level Spanish!). Most language learning apps include grammar-based activities, and you can also listen to how words and sentences are pronounced.

In the examples above, I've described how you can set out to deliberately learn the lexis, structure and pronunciation of the language – for example, by noting and reviewing key vocabulary on a piece of paper, or, if you have a mobile device, by using an app. We can also acquire language more informally by being exposed to it. You might notice a word that you've never heard before in an English language movie, and then decide to use it. If you're into digital gaming in English, for example, where you will typically work in teams to complete a mission online, you're likely to pick up terms related to gameplay; you will also most likely pick up some of the social language used to communicate with your team members. This is often referred to as *incidental* language learning, and it is no less valuable than formal language learning.

Motivating learners

A learner's progress in a language learning app is typically tracked by data-driven AI (see **2**), and additional lexis or activities are suggested by the app depending on that progress (this is known as *adaptive learning* – see **8**). Some language learning apps include elements of gamification, where learners can win points and move up through levels depending on their performance in games and quizzes. These gamification elements, including immediate feedback on progress and an attractive interface, are designed to maximise engagement, encouraging learners to spend more time on an activity while in a heightened state of attention. Engagement has the potential to support learning by encouraging learners to spend more time on an activity while in a state of high concentration.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

Practising the language

Second language acquisition research also tells us that apart from learning the nuts and bolts of a language (the lexis, grammar and pronunciation), we need to actually *use* it. This means applying our knowledge of the language to reading, writing, listening and speaking, in the conditions that research suggests are optimal. What are these optimal conditions? Research suggests that learners' motivation and engagement need to be high, that they are provided with help (or 'scaffolding') at the point of need (that is, when they need to use the language), and that feedback on what they say or write is personalised and constructive (see Thornbury, 2016). Clearly, teachers can provide all of these things for learners. The question is – can AI?

Generative AI tools (see **2**), it is argued, can potentially provide language learners with many of these optimal conditions, by acting as a personal tutor and language partner. Simpler knowledge-based AI chatbots can typically hold conversations based on narrow, predefined scripts because they respond to a series of pre-programmed prompts. If you ask these chatbots questions that they have not been programmed to respond to, the conversation soon breaks down. Chatbots based on generative AI, however, can engage in more naturalistic dialogue. One of the first and (at the time of writing) best-known of these generative AI chatbots is ChatGPT. A chatbot that is based on generative AI can 'remember' what has already been said during a conversation, and it can correct the learner if asked to do so. It can also provide another of the key elements needed for language learning that we identified above – scaffolding, or help at the point of need. Because generative AI can draw on previous elements in a conversation, the learner can ask it to explain things, to provide more examples or to simplify its language, as needed. While providing feedback or more examples, the chatbot can, if asked, help the learner notice key words or structures by highlighting their importance. Noticing is also an important part of language learning, and it happens when the learner pays conscious attention to specific words or structures, and then tries to use the words or structures correctly in their output. Duolingo, a well-known language learning app that was launched in 2012, was one of the first to integrate a generative AI chatbot (in this case, based on Open AI's GPT-4 technology). We explore the use of chatbots based on generative AI for language practice in more detail in **9**.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

As we've seen in this chapter, AI can provide support for learners in a range of areas that research has shown are important for language learning. However, there is one area in which AI inevitably falls short – that of providing a human connection. Language is, after all, about communication, and communicating with a machine is simply not the same as communicating with another human, with all the nuance, empathy and connection that this entails. Humanoid robots powered by generative AI underpinned by large language models (see **2**) are, inevitably, on their way. To what extent they may replace human conversation partners remains to be seen.

Thornbury, S. (2016). Educational Technology: Assessing its Fitness for Purpose. In McCarthy, M. (Ed.). *The Cambridge Guide to Blended Learning for Language Teaching*, pp. 25–35. Cambridge: Cambridge University Press.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

4 AI and creativity

We consider what the creative potential of generative AI means for human creativity. More specifically, we consider how this creativity can be used by language teachers and learners.

The first generative AI to arouse widespread public interest included tools that could generate *text* (for example, ChatGPT and Gemini) and tools that could generate *images* (for example, DALL-E and Stable Diffusion). Generative AI's ability to produce new content in text, image and multimedia formats, based on large amounts of data collected from the internet without permission, has led to major concerns around copyright and attribution (see **22**). It has also led to concerns over whether AI has the potential to replace human creativity.

Human versus AI creativity

Let's start with the thorny question of whether generative AI's creativity is as good as human creativity. If we measure the worth of creative content by its ability to win prizes, then the answer may well be 'yes'. In one well-known case, a photographer won a 2023 Sony World Photography award with an AI-generated photographic image. He stated that he had deliberately submitted the image to spark debate around the use of AI in creating images. AI-generated content is not always of high quality, though. Low quality books generated by AI are readily available for purchase on platforms like Amazon's Kindle, where there are no quality controls on self-published content. Unscrupulous academic journals have long been known to accept nonsensical articles generated by AI for publication, motivated by profit (Aldhous, 2009). In response to AI-generated content, competition organisers, journal publishers – and educational institutions – tend to have guidelines and principles around the acceptable use of AI. There are also laws, such as the European Union's AI Act (see **24**), around transparency and disclosure in the use of AI.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

There is, perhaps unsurprisingly, no clear agreement on exactly what human creativity *is*, although there are psychological tests available that try to measure creative original thinking. However, generative AI can beat humans at these tests. For example, in one controlled experiment, AI beat 99 percent of humans in the widely used Torrance Tests of Creative Thinking (Shimek, 2023). In another experiment, AI beat 91 percent of humans in the Alternative Uses Test for Creativity (Haase and Hanel, 2023). This may reflect the flawed nature of these tests more than it reflects a lack of human creativity, though. Another study found that AI can be beneficial for creativity, by helping humans come up with better creative ideas than they have on their own – although interestingly, very creative people seem to need less AI support – for now at least (Doshi and Hauser, 2023).

Creative ways of using AI with learners

For teachers and learners, the creative aspects of generative AI can be harnessed in the language classroom in several interesting ways. Below are a few ideas for using image-generation and text-generation tools.

- Teachers can generate images for learners to use as discussion prompts. Teachers can also add AI-generated images to worksheets.
- Learners can generate their own images on a theme. They can then compare the images with each other, discussing how their wording of the prompt affected what appears in the image. Or learners can try to guess the key words used in their partner's prompt, from what they see in the image.
- Learners can brainstorm ideas on a topic, and then use a text generation tool to help them come up with more ideas.
- Learners can learn about the history of art by generating images in the style of a particular painter, or in the style of an art movement such as impressionism or cubism. These images can then form the basis of research projects, in which pairs of learners find out more about an artist or movement. Learners can later present their findings to the class with both original and AI-generated images (and ask whether their classmates can tell the difference!).
- Learners can also learn about literary styles and genres by generating short texts on a topic in the style of a particular writer or in the style of different forms of poetry (e.g., a sonnet, a haiku). Learners can then compare and discuss these texts.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

- Learners can create a short, AI-generated poem on a topic. They can then change some of the language in the poem (for example, by replacing nouns, adjectives and verbs with other words), to make a new and personal version of the poem. This can be particularly effective at low levels, where learners often don't have the linguistic resources to create poetry from scratch.
- Learners can change the style of a single text. For example, they can generate a formal email from a text message or vice versa. Again, analysing and discussing the differences between the language used in these texts can be a helpful language learning activity for learners.

Exploring bias in images

Here is an activity that uses image-generator tools in a communicative language learning activity for learners. This activity gives learners the opportunity to discuss AI bias (and possibly their own biases) around jobs and gender.

- Use several different generative AI image tools to generate eight to ten images of a scientist using the prompt: *Make an image of a scientist in white coat standing in a laboratory, holding a test tube*.
- In class, show your learners the images and ask them to discuss the differences and similarities. Ask them to focus on gender – how many of the images are of men? (When I tried this activity in late 2023, *all* of the images generated were of men.)
- Tell your learners about a well-known experiment that has been carried out multiple times over decades with school children (Miller et al., 2018). The children are given a piece of blank paper and asked to draw a picture of a 'scientist'. Most children draw a man in a white coat. Fewer children draw a female scientist. Ask your learners why they think this happens. Point out that this experiment shows that we can internalise gender bias early in life. This drawing activity can help uncover unconscious bias with children.
- Discuss bias in generative AI with your learners (see **18** for more on this), and point out how the images you generated of a scientist reflect gender bias. If, by the time you try this activity, the images you generate show equal numbers of male and female scientists, you can discuss whether generative AI is making progress in addressing bias.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

- Put learners in pairs and ask them to create a list of jobs, including jobs that are typically associated with women (e.g., *nurse, secretary*), with men (e.g., *engineer, electrician*), and with both genders (e.g., *teacher, journalist*).
- Ask learners to use different AI image tools to generate multiple images of some of these jobs, and to discuss to what extent gender bias is reflected in the images produced.
- Hold a class discussion about the importance of recognising bias in AI. To extend this activity, you could put learners in pairs or small groups and ask them to research and share other examples of bias in AI. Several examples are provided in **18**, and learners will find plenty more examples online.

Apart from recognising bias, teachers and learners should be clear and transparent about using AI to generate images, text and ideas. AI-generated content should always be clearly labelled with the tool used, for example, 'Image generated by Stable Diffusion' or 'Additional ideas provided by ChatGPT'.

Aldhous, P. (2009). CRAP paper accepted by journal. *New Scientist*, 11 June 2009. Available at: https://www.newscientist.com/article/dn17288-crap-paper-accepted-by-journal. Accessed 24 December 2023.

Doshi, A. R. and Hauser, O. (2023). Generative Artificial Intelligence Enhances Creativity but Reduces the Diversity of Novel Content. Available at: https://ssrn.com/abstract=4535536 or http://dx.doi.org/10.2139/ssrn.4535536. Accessed 24 December 2023.

Haase, J. and Hanel, P. H. P. (2023). Artificial muses: Generative Artificial Intelligence Chatbots Have Risen to Human-Level Creativity. Available at: https://arxiv.org/abs/2303.12003. Accessed 24 December 2023.

Miller, D. I., Nolla, K. M., Eagly, A. H. and Uttal, D. H. (2018). The Development of Children's Gender-Science Stereotypes: A Meta-analysis of 5 Decades of U.S. Draw-A-Scientist Studies. *Child Development*, *89*(6), 1943–1955.

Shimek, C. (2023). AI Outperforms Humans in Creativity Test. *Neuroscience News*, 6 July 2023. Available at: https://neurosciencenews.com/ai-creativity-23585/. Accessed 24 December 2023.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

5 Technology and the hype cycle

Learning technologies have always attracted a lot of exaggerated claims and false promises. AI is no exception. Here we examine some of the reasons for this.

AI hype and hyperbole

When released in November 2022, ChatGPT received widespread public attention. Newspaper headlines predicted the death of education. School districts and entire countries banned – and then unbanned – its use. Some technology experts warned of generative AI posing an existential threat while others called it 'robotic, incoherent, unreliable, and untrustworthy' (Marcus, 2023). ChatGPT was not the first generative AI application to be reported on in mainstream media. A year earlier, Google's LAMDA (Language Model for Dialogue Applications) made headlines when one of its engineers claimed that it was an entity with thoughts and feelings of its own – in short, that it was a conscious being.

Digital learning technologies are often prone to hype and hyperbole, and they can give rise to polarised debates. There are many examples of this hype in the field of language learning, some of which you may have seen yourself. When Interactive Whiteboards (IWBs) first appeared in classrooms in the early 2000s, there was much talk of how they would improve learning outcomes. There was no evidence that this was true then, and there is no evidence that it's true now (see Hockly, 2013 for an overview of this). The new kid on the block, generative AI, has experienced its fair share of hype and hyperbole, from speculation that it is going to revolutionise learning, to fears that it will kill creativity (see **4**), replace teachers (see **17**) and even exterminate humankind.

Three ways to understand AI hype

Why are learning technologies so prone to exaggerated claims and false promises? A key reason is economic. Educational technology (EdTech) is

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

very big business, generating billions of dollars a year globally. Creating hype around an EdTech product by promising that it will improve learning – or better yet, 'revolutionise' learning through 'innovative' approaches – can help get the product into schools, generating profits for the EdTech company, even when there is no evidence to back up the company's claims. This is essentially hype for profit.

A second reason is the underlying belief that a certain technology (for example, a device, or a piece of software) is essential to language learning. This is often referred to as *technology solutionism*, and it ignores the fact that learning is a complex and personal process. Generative AI (see **2**) is a prime candidate for technology solutionism. Indeed, the release of ChatGPT was followed by a flood of chatbot apps that claimed to be 'based on AI' that supports language learning better than ever before; however, whether an app was underpinned by knowledge-based AI or by generative AI was rarely made clear in any publicity. Simply putting the word 'AI' into an advert allowed the app to tap into the prevailing hype, despite the fact that AI has been used in language learning software for decades (see **2**) and no miraculous language learning has taken place.

A third reason is due to an essentially 'mechanistic' view of language and of learning. If one believes that language is made up of small components that can be learned by progressively assembling these components over time, then certain learning technologies and approaches (such as adaptive learning – see **8**) fit well with this belief. However, although learning the essential elements of a language (grammar, vocabulary, pronunciation, etc.) is necessary, there is a lot more needed, as we saw in **3**. AI's much praised reliance on data suggests that all aspects of learning can be quantified, and outcomes predicted. This is not the case. A 2020 study carried out at Princeton University in the USA (Salganik, 2020) challenged teams of AI researchers and data scientists to predict outcomes for children (such as long-term future exam results and perseverance with their schoolwork) based on data. Thirteen thousand data points on over 4,000 families over a period of 15 years were provided, but none of the teams were able to develop effective statistical models that could explain the actual outcomes. Some things – such as decision-making based on personal

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

experience, reflection, empathy, emotion and imagination, as well as the host of social and environmental factors that affect these nuanced processes – cannot be quantified.

Resisting AI hype with learners

Overall, then, like many new educational technologies, generative AI has been through a cycle of hype and hyperbole. It is clear that generative AI is a key development in the field of computing, and it is likely to be deployed increasingly widely in the field of education. However, as one researcher puts it, 'we need to be mindful that education remains vulnerable to what can be termed AI theatre' (Selwyn, 2022, p. 621). Resisting the hype can be challenging, though. One way to address this issue is to look at the metaphors we use to talk about AI (and educational technologies, or EdTech, in general). As one researcher found, we tend to see these tools in five different ways (Mason, 2018). These are:

1 manual labour – we see EdTech as a tool
2 construction – we see EdTech as supporting scaffolding and knowledge construction
3 mechanism – we see EdTech as a machine
4 biological life – we see EdTech as an ecosystem or natural evolution
5 journey – we see EdTech as a progressive journey leading us towards improved learning.

It's interesting to reflect on how AI seems to attract many (if not all) of these metaphors. Metaphors influence how we see and use technology. They shape what we expect from it, and what we perceive as good and bad about it. We often use metaphors without realising how much they affect our ideas about something, and recognising these metaphors helps us think more critically about them.

Here's a short classroom activity you could carry out with higher proficiency learners to explore this:

- Write '*AI is like …*' on the board, and ask your learners to each write at least five different endings to this sentence.
- Put the learners into small groups to compare their sentences. Can they group the sentences in any way (for example, do some refer to machines, or to biology/nature, or to a journey, etc.)?

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press

- Get feedback from the group, and ask them what their similes (a form of metaphor) suggest about their underlying beliefs and feelings around AI.
- Share the five metaphor types listed above. Did their similes reflect any of these?
- Finally, discuss to what extent your learners think current AI reflects each of the five metaphors, and what this means for how we might understand and use AI.

An activity like this can help develop our own and our learners' critical digital literacies, which are essential to identify and resist hype. We examine the area of digital literacies in more detail in **25**.

Hockly, N. (2013). Interactive Whiteboards. *English Language Teaching Journal, 67*, 3, 354–358. Oxford: Oxford University Press.

Marcus, G. (2023). The Rise and Fall of ChatGPT? Blog post, 23 August 2023. https://garymarcus.substack.com/p/the-rise-and-fall-of-chatgpt. Accessed 25 December 2023.

Mason, L. (2018). A critical metaphor analysis of educational technology research in the social studies. *Contemporary Issues in Technology and Teacher Education*, *18*, 3, 538–555.

Salganik, M. J. (2020). Measuring the predictability of life outcomes with a scientific mass collaboration. *PNAS*, 117, 15.: 8398–8403. Available at: https://www.pnas.org/doi/epdf/10.1073/pnas.1915006117. Accessed 22 January 2024.

Selwyn, N. (2022). The future of AI and education: Some cautionary notes. *European Journal of Education*, *57*(4): 620–631.

https://doi.org/10.1017/9781009804509.001 Published online by Cambridge University Press