

PLENARY SPEECH

From revolution to evolution: What generative AI really means for language learning

Mark Pegrum 

The Graduate School of Education, The University of Western Australia, Perth, Australia
Email: mark.pegrum@uwa.edu.au

(Received 8 April 2025; accepted 12 April 2025)

Abstract

History is littered with unfulfilled promises that emerging technologies – from radios to televisions, and from computers to mobile phones – would completely transform teaching and learning. Now the same promises are being made of generative artificial intelligence (AI). This presentation argues that we should not be focusing on educational revolution, but instead on educational evolution. Education is a complex social, cultural, and political endeavour, serving multiple purposes and multiple stakeholders, and technology is just one of many elements in this large ecosystem.

Focusing on the context of language teaching and learning, this presentation discusses what has changed technologically, and suggests what could and should change educationally. It shows that ChatGPT and a range of other generative AI tools can contribute to language and literacy development in a number of ways, but that we need to be wary of their pedagogical, social, and environmental risks. Educators must develop the AI literacy necessary to take a more nuanced view of generative AI, and we must help our students to do the same.

This paper is based on a keynote presentation delivered at the English Australia Conference in Perth, Australia, on 12 September 2024, with some elaborations for the written version alongside minor updates to reflect more recent developments and publications.

Keywords: AI; AI literacy; artificial intelligence; digital literacies; generative AI; language learning; language teaching

Keynote presentation at the English Australia Conference, Perth, Australia, 11–13 September 2024

Presentation

I would like to begin by acknowledging that teaching and learning have been taking place for thousands of years right here on the traditional lands of the Noongar people, in what we now call Perth, Western Australia. I would also like to recognise that we can learn a lot from past teaching practices that can help to inform our future teaching practices.

1. Introduction: Not another revolution!

My presentation today focuses on generative artificial intelligence (AI), which some people claim is about to revolutionise education. We will start by zooming out to take a technological view, which is relevant to education in general; next we will zoom in to take an educational and assessment view, with examples selected to specifically illustrate language teaching and learning wherever possible; and then we will zoom out once more to take a societal view, which is again relevant to all areas of education. Finally, we will wrap up with some brief comments on AI literacy.

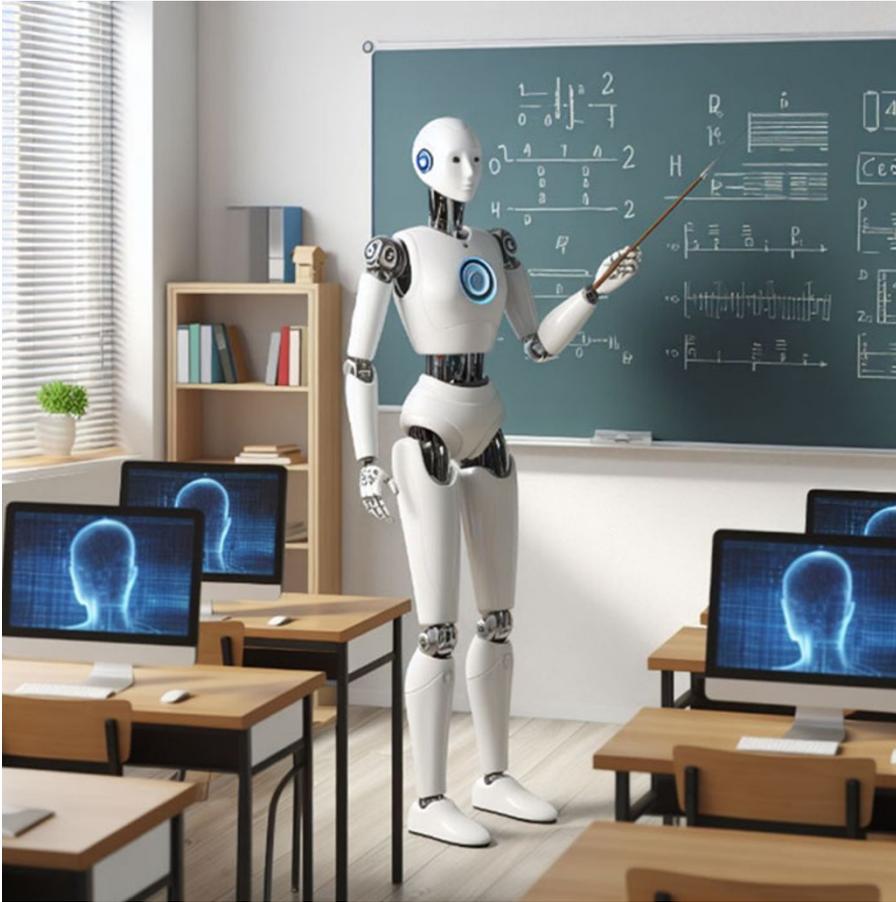


Figure 1. A humanoid robot teaching. DALL-E 3 image created by Mark Pegrum (2024) via Microsoft Copilot.

First, however, I would like to say a few words about revolutions. The term ‘revolution’ can refer to the forcible, often violent, overthrow of an outdated system in favour of a new one. The term can also refer more generically to the complete, radical overhaul of an old system. For more than a century, enthusiasts have been promising us an educational technology revolution. Radios were going to allow widely accessible, up-to-date, home-delivered learning; so, too, were televisions. Computers were going to allow all of the above, along with collaborative networking; so, too, were mobile phones. If we ask ourselves whether these devices have IMPROVED education, the answer is: absolutely. But if we ask ourselves whether these devices have REVOLUTIONISED education, the answer must be: far from it. We still have ministries and departments of education, schools and universities, curricula and textbooks, teachers and classrooms, students and assignments. What these devices have done is not to support an educational REVOLUTION, but rather to support educational EVOLUTION, with the latter term understood to mean a slow, stage-by-stage development.

Undaunted by the failure of past edtech revolutions to materialise, enthusiasts now promise us that the AI revolution is coming for education. The image shown here (see Fig. 1) is a composite of key features of the imagery typically accompanying such predictions: a robot teacher; rows of empty desks, with students either absent or online; and classroom décor which seems to be modelled on schools of a century ago, complete with chalkboards (indeed, chalkboards seem to be everywhere in these images).

While this image is a composite, let us now look at the text of some actual titles, headlines, and slogans from readily accessible online sources – aligned sometimes closely and sometimes more loosely with the fields of education and technology, and sometimes with a more or less commercial bent – as typified by the very generic: ‘Artificial intelligence for education revolution’ (Gurung, 2022). For educators who work in a school, AI will be ‘[r]evolutionizing K–12 education’ (Lim, 2023); for those who work in or attend a university, AI will be ‘revolutionizing higher education’ (ASM Research Publications, 2023); while those who work with international students, as many language teachers do, can look forward to the ‘revolution of international education’ (Focus Education, 2024). Implying that contemporary education is effectively a problem in search of a solution, another source promises, with somewhat unusual phraseology, ‘AI solutions for education industry that can bring revolution’ (Matellio, 2022).

So what exactly might this revolution be expected to entail? As exemplified by the title ‘Learn faster and better with AI’ (Learnboost, n.d.), there is some excitement about acceleration, based on the dubious assumption that rapid learning is automatically better than more gradual learning. However, the greatest enthusiasm is reserved for predictions about the growing ‘personalisation’ of learning; to cite just two recent instances: ‘Education revolution: The rise of the AI-powered personalised classroom’ (Thompson, 2023) and ‘Transforming education: AI-powered personalized learning revolution’ (Zia, 2023). This is highly ironic in the context of accompanying images populated by robots and screens, without a single person in sight. This is a point to which we will return.

Without fail, what these articles and posts promise us, as teachers or students or both, is help in ‘[n]avigating the AI revolution in education’ (Picardo, 2023). This may be just as well, because otherwise – to conclude with a mildly threatening headline – ‘[t]he AI reeducation revolution is coming for you!’ (West, 2023).

Why, then, do I not share the view that generative AI, or any other educational technology, is going to revolutionise education? It is essentially because, as Justin Reich points out in his 2020 book, *Failure to disrupt*, education is a complex ecosystem which serves multiple purposes and multiple stakeholders; these include students, teachers, parents, employers, and governments. Technology is just one element within this large ecosystem. Observing that acceleration may not be an unadulterated good in education, Reich notes: ‘Trying to accelerate learning by ramping up technology is like putting rockets on butterfly wings. More force does not lead linearly to more progress’ (Reich, 2020, p. 9). He later elaborates:

Improvements in education very rarely, perhaps never, come by way of dramatic transformations. They come through deep, long-term commitment to the plodding work of building more robust systems. ... New technologies get introduced into complex learning ecologies, and those complex learning ecologies require multiple changes at multiple levels to take advantage of new technologies. ... changes in educational systems are necessarily incremental, but step change is what continuous, incremental change looks like from a distance. (pp. 243 and 245)

In other words, it is through slow, careful stages that we arrive at step change, or major change, or major forms of evolution in education.

Echoing this point from a practical rather than a theoretical perspective, the British Council’s 2023 report on *Artificial Intelligence and English language teaching*, based on a survey of over 1300 teachers and interviews with almost 20 edtech experts, presents a picture of ‘the future as one of continuity rather than widespread disruption of educational systems’ (Edmett et al., 2023, p. 50).

In my view, therefore, it is time to stop talking about an AI revolution in education, and instead to start talking about how AI can helpfully support the evolution of education in areas where change might be needed. Leaving behind the distractions of revolutionary rhetoric, this presentation will, as indicated, focus on the present and likely future technological aspects of AI, before turning to its educational and assessment implications, and then its broader societal challenges. As we undertake

this journey, we should bear in mind that AI could help to make education better in some ways, but it could also make education worse in other ways. Consequently, as educators, we need to develop the AI literacy to ensure that AI is supporting the evolution of education in what we believe to be a positive direction.

2. A technological view

Generative AI first came to wide public attention with the release of ChatGPT in late 2022. It became the fastest digital service ever to reach one million users, which occurred within five days of its release (though it was surpassed the following year by Threads, which arguably benefited from its link with Instagram; Buchholz, 2023). Let us take a step back for a moment to explore how ChatGPT fits into the bigger picture of artificial intelligence.

You may be aware that the broad field of AI can be roughly divided into two subfields (Pegrum, 2019). **STRONG AI**, or **GENERAL AI**, refers to the ability of machines to apply intelligence to a whole range of tasks, much as humans do. Despite many decades of research, strong AI remains for now largely the province of utopian, or perhaps more often, dystopian science fiction. Think of *2001: A Space Odyssey* (1968), *Westworld* (1973), *Blade Runner* (1982), *The Terminator* (1984), *The Matrix* (1999), *Her* (2013), *Ex Machina* (2014), *I Am Mother* (2019), or *Foundation* (2021 onwards); this has been a long-standing cinematic obsession.

By contrast, **WEAK AI**, or **NARROW AI**, refers to the ability of machines to apply intelligence to very specific tasks, which generally involve very specific forms of pattern matching, and where these machines can frequently exceed human abilities. Recent years have seen enormous advances in this area, thanks to machine learning and, in particular, deep learning based on artificial neural networks. Some of the outcomes are seen in today's image and object recognition, voice recognition, natural language processing, automated translation, and learning analytics. Thus, it can be said that for now the broad field of AI is absolutely dominated by the subfield of weak AI.

Within this subfield of weak AI, we find both conversational and generative AI. **CONVERSATIONAL AI** is trained on huge datasets of human conversations, and its function is to mimic such conversations. Conversational agents, with which humans can interact directly, include organisational or corporate chatbots, and first-generation voice assistants like Amazon's Alexa or Apple's Siri. By contrast, **GENERATIVE AI** is grounded in large language models (LLMs) and trained on vast datasets of texts and other artefacts. Its function is to generate new, or at least remixed, content, in response to human prompts. Examples include GPT.

The major breakthrough made by the San Francisco-based company OpenAI, as released to the public on 30 November 2022, involved combining generative AI in the form of GPT with the ethos of conversational AI in the form of a chat interface, resulting in ChatGPT. It was this combination, and this ability to chat with generative AI, which really captured the public's imagination and which has led to the subsequent proliferation of generative AI tools.

Generative AI does not only involve text generation. One of the best-known of today's image generators is DALL-E, also from OpenAI, which produces images in response to text prompts; other examples include Adobe's Firefly, Midjourney, and Stability AI's Stable Diffusion. To illustrate how fast this technology is advancing, compare the image I produced for my website in 2023 with DALL-E 2, using the prompt 'the future development of web 3.0 and artificial intelligence, in the style of digital art', with the image I produced in 2024 with the next-generation DALL-E 3, using exactly the same prompt (see Fig. 2).

Audio and video are increasingly well-represented among generative AI tools. For instance, Microsoft's VALL-E, given a 3-second recording of a person speaking, can create an audio file simulating the same voice speaking whatever text a user inputs, while VALL-E X can then render it in multiple languages (note that these tools are not currently available to the public). The second half of 2024 has seen the arrival of an array of highly sophisticated video generation software, including



Figure 2. Future of web 3.0 & AI. On left: DALL-E 2 image created by Mark Pegrum (2023) via Microsoft Copilot. On right: DALL-E 3 image created by Mark Pegrum (2024) via Microsoft Copilot.

Luma Labs' Dream Machine, OpenAI's Sora, Runway's Gen-3, and Kuaishou's Kling AI, all of it readily available to everyday end users.

Before returning to ChatGPT itself, I would like to comment briefly on the use of terminology. We often talk about technology in metaphorical terms; all metaphors reveal certain things, but also hide certain things. I am sure you have heard the term *HALLUCINATIONS* used to refer to the incorrect or simply invented information sometimes generated by AI in response to human queries. What is important to understand here is that generative AI is just probabilistic pattern-matching software. It predicts the next word – or the next pixel or the next note – based on the patterns it identifies in its datasets. To put it another way, LLMs are in essence 'incredibly sophisticated autocomplete systems' (Mollick, 2024b). Generative AI has no comprehension of texts or other media, much less of the external world to which they refer. From another angle, then, we might say that *EVERYTHING* produced by generative AI is actually a hallucination (cf. Franks, 2023). It is fortunate for us that most of the time these hallucinations (grounded as they are in human datasets) do in fact correspond to our understandings and our perceptions of the external world.

Then there is the term *GENERATIVE AI* itself. Perhaps a better term, as suggested by Massachusetts Institute of Technology (MIT) Media Lab computer scientist Joy Buolamwini, might be *REGURGITATIVE AI* (Brown, 2024). What is important to understand here is that current AI is incapable of creating, or generating, anything new. All it can do is to remix, at a very granular level, existing content originally created by human beings. As for whether and to what extent a remix can be considered original, that is a question we will leave for another day.

Let us now turn to GPT, the LLM that underpins ChatGPT. Existing since 2018 in evolving iterations, GPT has been used for some time across a range of industries, for example to generate code, to write news articles, and to support the production of marketing materials. When ChatGPT was first released, it was based on GPT-3.5.

To interact with a generative AI chatbot like ChatGPT, we can start by inputting a text-based or voice-based question or command, known as a *PROMPT*. It is vital to use the right prompt to get the right output, so much so that we are now seeing the emergence of a new field, and indeed a new profession, which is being referred to as *PROMPT ENGINEERING*:

Like writing and coding before it, prompt engineering is an emergent form of thinking. It lies somewhere between conversation and query, between programming and prose. It is the one part of this fast-changing, uncertain future that feels distinctly human. (Warzel, 2023)

Interestingly, in 2023, Matt Bell from Anthropic, the company behind the generative AI chatbot Claude, told an interviewer: ‘In fact, our best prompter is a philosopher. Good prompting involves writing extremely clear explanations ... and figuring out reasons for misunderstandings and how to prevent them’ (cited in Purtill, 2023). For many decades, educators have been supporting students in developing *SEARCH LITERACY*, encompassing the ability to compose carefully considered and appropriately detailed search strings in order to access the desired search results. What we might call *PROMPT LITERACY* (Gattupalli et al., 2023; Hwang et al., 2023) is really an extension of this idea. While there may be a small number of specialised – and highly lucrative – job openings for prompt engineers in the future, all users of generative AI will need to develop a certain degree of prompt literacy to support effective everyday interactions with this technology.

While ChatGPT and its underlying LLM are continuing to evolve, there are also numerous competitors, with new iterations appearing almost daily, each one promising improvements over the last. If you are not sure which one to choose, or if you would like to compare results from several chatbots – or even to ask your students to do so – then you could turn to Poe, which pulls together OpenAI’s ChatGPT, Anthropic’s Claude, Google’s Gemini, and many more.

Some (though, as of 2024, not yet all) generative AI chatbots have live internet access, so the line between generative AI and search is beginning to blur. Examples include Google’s Gemini, Microsoft’s Copilot, and Perplexity AI. When you come across such a service, you might well ask yourself: is this a generative AI chatbot; or is this a search engine? The answer is yes; and yes. This fusion of generative AI and search would seem to be the way of the future.

Speaking of the future: current trends include a shift away from publicly available LLMs towards more secure enterprise or on-device tools (the latter sometimes labelled small language models, or SLMs) among users with financial means, especially in the Global North; the integration of generative AI into older, first-generation conversational AI interfaces like Amazon’s Alexa and Apple’s Siri; and the intertwined evolution of semi-autonomous personal agents, likely anchored in on-device SLMs, which can converse in everyday language and navigate and control smartphone (or computer) apps (Low, 2024). In other developments, in February 2024 an intriguing post appeared on X, formerly Twitter, from Logan Kilpatrick, who at the time led developer relations at OpenAI. He wrote: ‘The final form for ChatGPT is not chat :)’ (@OfficialLoganK, 2024). What, then, might it look like?

Given the later announcement of OpenAI’s collaboration with robotics company Figure, one possibility is the emergence of humanoid AI robots, with the founder of Figure, Brett Adcock, describing his aim ‘to bring embodied AI into the world to make a transformative impact on humanity’ (cited in PR Newswire, 2024). A number of other companies are also working in this space. This is potentially quite significant, with some AI researchers advocating a shift away from the traditional emphasis on Cartesian dualism – essentially, the idea that the mind and the body are separable – and towards an emphasis on embodied cognition, arguing that strong AI may only be able to emerge once it has a body through which to experience the world (Paolo et al., 2024). So, might this be the frontier where strong AI will one day arise?

We will have to wait and see. In the meantime, let us turn to the implications of today’s technologies for education and assessment, with a particular focus on the teaching and learning of language and literacy.

3. An educational view

Students and teachers have now been experimenting with generative AI like ChatGPT since late 2022, and extensive reports have emerged of promising educational uses (Pegrum, 2024). Students can, for instance, ask AI chatbots to explain grammar points, contextualise vocabulary, devise revision questions, and take on the role of either a teacher or a learner, depending on the context (see Table 1). As we take a look at some examples, we should bear in mind that when the technology is treated as a

Table 1. Selected student uses of generative AI

Language learners can ask/use generative AI to ...
• explain learning points (e.g. grammar rules, lexis, or themes)
• provide examples of words, phrases, or structures in use
• find and summarise, or simplify, existing documents
• suggest ideas for essays, projects, or other tasks
• produce first drafts of titles, task outlines, or full texts
• improve the grammar, vocabulary, and style of texts
• modify the genre and register of texts
• offer constructive feedback on texts
• co-generate stories in a choose-your-own-adventure style
• engage in conversation (including role-plays) in multiple languages
• create self-study revision questions and games
• take on the role of a teacher or Socratic tutor
• take on the role of a student (with students acting as teachers).

simple replacement for teacher input or student output, it is potentially very problematic; but when it is used in a supporting role, it is potentially very valuable.

In the most controlled environments, students can engage in individualised learning activities and receive individualised feedback. For instance, my Australian colleague Kate Kerr reports on successfully using Microsoft's Reading Coach with young learners of English, who can co-create a story using generative AI, read passages aloud, and obtain feedback on the accuracy of their pronunciation and guidance on words they need to practise.

At higher levels, students can make use of GPT-based tutors such as the Khan Academy's Khanmigo, set up as a Socratic tutor who refuses to give students direct answers to their questions, but rather guides them step-by-step towards finding their own answers. A similar GPT-based tutoring function is seen in Duolingo Max, a premium offering from the language learning app Duolingo, initially available for French and Spanish (Caldwell, 2024). Students converse with an AI tutor who can, for example, explain grammar points in multiple steps, or engage in target language role-plays and then give feedback on them. Some teachers are also experimenting with having their students interact with historical or fictional figures through platforms like Hello History and Character AI (both multilingual).

Fluid conversational scenarios with human oversight may hold even more promise for language practice. An Argentina-based colleague, Doris Molero, reports on supervising an advanced learner of English who engaged in role-plays with the Pi AI chatbot, which could be asked to take on various personas such as, in one case, a car dealer. In another class, Doris used both ChatGPT and Google's NotebookLM with an English learner who was tasked with interviewing one of her favourite singers. When the student initially struggled to come up with in-depth interview questions, inspiration was sought from NotebookLM, which can generate questions pertaining to relevant information sources. The student's homework was to polish up a question set for the next class, where she then role-played a 'live' interview with the singer, who was simulated by ChatGPT.

There are numerous tools which can help students to make a start on their tasks. When studying large amounts of material or revising for exams, students could employ software like Mindgrasp (which supports more than 30 languages) to automatically create summaries and take notes from documents or artefacts ranging from PDF files to videos. When English learners are exploring academic literature, they could employ software like Elicit to identify relevant references and compose

Table 2. Selected teacher uses of generative AI

Language teachers can ask/use generative AI to ...
• design step-by-step lesson plans
• create teaching materials and handouts
• devise assessments and rubrics
• draft model assignments
• generate responses with specifically planned flaws for students to identify
• generate multiple responses to a question (which students may critique)
• provide a first draft of feedback on student work (which students may critique)
• build customised learning chatbots drawing on specific datasets and/or following specific interactional instructions
• analyse student data to improve teaching and/or learning
• draft student reports based on the teacher's notes
• draft meeting summaries based on notes
• draft official documents based on notes.

a first draft of a literature review, which they could then critically edit and supplement with further literature in order to deepen and broaden the analysis. The text below summarises the ‘top 4 papers’ found by the free version of Elicit in response to my question ‘How can language learners benefit from generative AI?’ (updated 16 October 2024):

Generative AI offers significant opportunities for language learners and educators. It can provide personalized practice, feedback, and study plans (Godwin-Jones, 2024), as well as support Universal Design for Learning to benefit all learners, including English learners and students with disabilities (Evmenova et al., 2024). AI-powered tools can enhance language acquisition through customized learning journeys and dynamic content (Creely, 2024). However, challenges exist, including ethical concerns about authenticity and intellectual property (Creely, 2024), and limitations in addressing nuanced social and cultural aspects of language use (Godwin-Jones, 2024). There are also concerns about AI’s potential to standardize language expression and propagate limited cultural narratives (Creely, 2024). Despite these challenges, generative AI is reshaping second language education, with numerous AI-powered language learning tools widely available and used globally (Davoodifard & Eskin, 2024). A balanced approach that leverages AI’s benefits while addressing its drawbacks is recommended for language education in the emerging AI-centric digital landscape (Creely, 2024).

While such tools can help students get started, other tools exist to help them improve on what they themselves have started, for example honing their clarity, formality, or tone. Well-known services Grammarly and Wordtune, though offering some translation features, are focused on composing and revising English-language text, but multilingual alternatives exist.

Varied teacher uses of generative AI have also emerged. The most obvious possibilities include creating lesson plans, teaching materials, and assessments and rubrics (see Table 2). Some teachers have gone further by designing their own dedicated learning chatbots, while others have experimented with using generative AI to provide feedback to students, to analyse student data, and even to draft student reports, but it should be noted that there are permission and privacy concerns with the uploading of sensitive data into publicly available systems, which may retain it and incorporate it into their training datasets.

In a 2023 survey of over 350 teachers of 16 different languages, my colleagues Antonie Alm and Louise Ohashi, based respectively in New Zealand and Japan, found that language teachers were open

Table 3. Language teachers' openness to using ChatGPT to create resource types, from highest to lowest openness (source: adapted from Alm & Ohashi, 2024)

Reading	Generate reading passages or texts on different topics and levels of difficulty
Writing	Generate templates for different types of writing, such as an essay, a letter, or a report
Conversation	Generate scripts for role-playing activities
Quizzes and tests	Generate quizzes and tests on different language skills, such as vocabulary, grammar, reading comprehension, and listening
Games	Generate language learning games such as crossword puzzles, word searches, and hangman games
Grammar	Generate worksheets with grammar rules, examples, and exercises
Vocabulary	Generate flashcards with vocabulary words and their definitions, translations, or synonyms

to using ChatGPT to create the following kinds of resources, ordered here by teachers' levels of openness towards creating them (see Table 3; note that 'Conversation' and 'Quizzes and tests' were equally rated). Certainly, all of these uses are possible.

Let us take a look at some examples of teacher uses of generative AI. Twee is one of the older generative AI tools and is aimed at teachers of English, including English as a foreign, second, or additional language, but could equally be used across a range of other subject areas. Options include automatically generating questions about YouTube videos; generating dialogues or stories at a specified level of the Common European Framework of Reference for Languages (CEFR) and incorporating specified structures and lexis; and building everything from multiple-choice questions to fill-the-gap exercises. A swathe of similar services, each with its own points of emphasis, is now appearing. The multilingual Diffit can automatically build a whole lesson around a given video or text, generating resources and activities at a specified school grade level and in a specified language, including bullet point summaries, vocabulary lists, graphic organisers, and multiple-choice or short answer questions. AI functionality is also being introduced into older platforms, perhaps best exemplified by Padlet (dating back to 2008, when it was known as Wallwisher), which now offers boards created from a series of AI recipes including 'Lesson plan', 'Ideas for class activities', 'Assessment polls', and 'Rubric'.

Some of the most innovative approaches involve the use of multiple platforms in combination to create materials closely tailored to class needs. An Australian colleague, Sandra Houghton, reports how she developed a lesson on community projects for her English learners by using Magic School to generate lesson stages, ChatGPT to produce a reference list of key functional language, Canva's Magic Studio to integrate all of this into a handout, and finally Canva's AI Presentation Maker to create slides signalling lesson stages and reminding students of the key language. In a rather different example, Doris Molero, whom I mentioned earlier, composed an environmentally themed story for her young English learners using ChatGPT, illustrated it with Copilot, compiled the text and images into a book with Book Creator, and then set her students questions using Twee.

In the future, we may also see many language teachers creating their own customised AI chatbots, such as Doris's Pionia Destiny on Character AI, which her students use for homework conversations on new topics covered in class. Meanwhile, Louise Ohashi demonstrates a self-created GPT chatbot to her tertiary level learners of English in Japan before inviting them to experiment with collaboratively creating customised bots on Poe, focusing on their preferred areas of language practice.

Creativity and critique can and should go hand-in-hand. Borrowing an activity suggestion from the *Generative AI idea pack for English language teachers* (CUP [Cambridge University Press], 2024), I entered the prompt 'Create a poster of Australian people' into Copilot to see how my own nationality might be represented; one example is reproduced below (see Fig. 3). This activity naturally leads to analysis and critique of cultural representations, and potentially to the building of intercultural



Figure 3. A poster of Australian people. DALL-E 3 image created by Mark Pegrum (2024) via Microsoft Copilot.

literacy. It can also lend itself to creativity, and potentially to the building of prompt literacy, if students are invited to craft their own detailed prompts to generate images that capture their perspectives and, ideally, reflect a somewhat broader representation of identities and lifestyles than what we see here.

Generative AI has implications not only for learning, but for assessment. One possible approach is to try to catch students out if they overuse generative AI support, as plagiarism detection companies aim to do. Although these companies claim high detection rates for AI-generated texts, many researchers remain unconvinced that AI detectors are reliable (Ardito, 2024; Sadasivan et al., 2024; Weber-Wulff et al., 2023). There is, moreover, an unavoidable lack of transparency in AI detection:

Although text matching software also suffers from false positives and false negatives (Foltýnek et al., 2020), at least it is possible to provide evidence of potential misconduct. In the case of the detection tools for AI-generated text, this is not the case. (Weber-Wulff et al., 2023, p. 27)

More concerning still is research suggesting that the approaches taken by AI detectors, which rely on factors like perplexity and burstiness (essentially, measures of the unpredictability of a text), are inherently biased against those who use language in more constrained ways. This includes not only non-native speakers (Liang et al., 2023; note that it may be possible to correct for such bias with customised detectors, per Jiang et al., 2024) but may also include those with autism spectrum disorder (ASD), as flagged up in a white paper from Anthology (2023), the company behind the Blackboard Learning Management System (LMS). This has enormous ramifications in terms of diversity, equity, and inclusion.

On the other hand, as pointed out in a recent UNESCO guide, generative AI has the potential, albeit still largely untested at scale, to benefit learners with specific needs, including through AI-powered accessibility tools (e.g. for learners with hearing or visual impairments) and as a '[g]enerative amplifier for marginalized users' (e.g. learners from minority linguistic or cultural backgrounds) (Miao & Holmes, 2023, pp. 34–35). There is an apparent contradiction here. A moment ago, we saw that generative AI could serve to EXCLUDE diverse students; here, we read that it could help to INCLUDE diverse students. To me, what this says is that much more research is needed on exactly HOW generative AI can be used to support diverse students in their learning and assessments.

In the era of generative AI, a better approach to both learning and assessment for all students may be to set more tasks at the higher levels of Bloom's Revised Taxonomy, something which has indeed been discussed for many years in the context of developing students' higher-order thinking skills. This might mean students critiquing a text generated by an AI chatbot and seeking to improve on it; or asking the chatbot for feedback on their own work and then offering feedback on the feedback, indicating which elements they would act on, which they would not, and why; or engaging the chatbot in iterative co-creation of artefacts contextualised by the students to address real world needs or challenges. In sum, the most promising educational approaches are likely to be those where students demonstrate their ability to work effectively in human AI-partnerships.

In the *Microsoft new future of work report 2023*, it is suggested that there are many different formats and levels for collaboration between humans and AI, but that human-centred COMPLEMENTARITY should be the goal (Butler et al., 2023, p. 15). This fits neatly with an idea proposed by Cornell law professor Frank Pasquale who, well before the current generative AI explosion, had already suggested that we should be focusing less on AI and more on IA, by which he means INTELLIGENCE AUGMENTATION – that is, machines augmenting human intelligence (Pasquale, 2020). More recently, a similar idea has been proposed by University of Pennsylvania researcher Ethan Mollick under the name of CO-INTELLIGENCE (Mollick, 2024a). Whether we call it complementarity, or augmentation, or co-intelligence, the concept is broadly similar: the future of both study and work is all about humans and AI collaborating.

4. A societal view

We said earlier that, overall, generative AI could contribute to making education better, or it could contribute to making it worse. It brings with it a raft of issues, some of direct pedagogical relevance, and some of broader societal relevance. Let us take a quick look at four key sets of issues.

First, there are issues of accuracy and reliability. On its homepage, ChatGPT offers this casually worded warning: 'ChatGPT can make mistakes. Check important info'. One of the greatest concerns with generative AI is the erroneous responses it can produce; as mentioned earlier, these are known as hallucinations in the narrow sense of the term. They extend from inventing facts to inventing research studies and references (Sharples, 2023). Some hallucinations, moreover, may have real world consequences that end up in law courts. Here are just three news headlines regarding false accusations made by AI in cases spread across Australia, the USA, and Germany:

Hepburn mayor may sue OpenAI for defamation over false ChatGPT claims – *ABC News, Australia* (Mayers, 2023)

OpenAI sued for defamation after ChatGPT generates fake complaint accusing man of embezzlement – *Forbes, USA* (Ray, 2023)

Microsoft's AI falsely accuses reporter of crimes – *MSN, USA* (Landes, 2024)

There is also some danger of hallucinations and misinformation accumulating. If, as is likely, generative AI increasingly comes to be trained on the outputs of other generative AI, this may lead over

time to degraded results – and if this phenomenon becomes widespread enough, it could result in *MODEL COLLAPSE*, effectively, the collapse of generative AI as we know it (Marr, 2024; Shumailov et al., 2024). Of course, aside from unintentionally false information, or *MISINFORMATION*, there is the matter of intentionally false information, or *DISINFORMATION*. One looming dilemma is how best to regulate, and educate about, today's proliferating *DEEPFAKES*, where AI is used to generate fake images, videos, or audio recordings, often of well-known people, that misrepresent their words or actions.

Second, there are issues of standardisation and homogenisation, including of our language. There are copious online lists of overused ChatGPT words such as *DELVE*, *TAPESTRY*, *LANDSCAPE*, and *REALM*, and phrases such as *DIVE INTO*, *BASED ON THE INFORMATION PROVIDED*, *NAVIGATING THE LANDSCAPE*, and *A TESTAMENT TO* (Dan, 2024). These may be useful for teachers as telltale signs of student use of generative AI, while students could be asked to generate such lists themselves as an awareness-raising exercise. Yet it is not just students' writing that is affected; so too is the writing of academics. Following an approach suggested by Swinburne University of Technology lecturer Jeremy Nguyen in a widely read post on X (@JeremyNguyenPhD, 2024), if we conduct a search for occurrences of *DELVE* (the first verb in Dan's list above) in papers in the PubMed academic database, we find a jump from 510 occurrences in 2021 (the year before the release of ChatGPT), to 629 in 2022 (the year of its release on 30 November), to 2,872 in 2023 (the first full year of its availability), and to 9,193 as of mid-November 2024. While correlation does not prove causation, this is certainly suggestive of the influence of ChatGPT on academic writing. Moreover, there is likely to be a feedback loop here. As we, as readers and researchers, start to come across particular terms and formulations more often, we may subconsciously begin using them more often ourselves. In this way, AI could make inroads into homogenising our language usage.

My next example might initially sound like the polar opposite of homogenisation, but I would argue that it is a different, and potentially more acute, manifestation of a similar underlying process of standardisation, not so much of language but of learning itself. I believe that when we are talking about generative AI, educators should avoid the term 'personalisation', which seems to have leaked into our discourse directly from technology companies' marketing slogans (for related recent commentary, see Furze, 2024). In general, AI is anything but personal. It limits our *PERSONAL CONNECTION* with human interlocutors. More importantly, it often limits our *PERSONAL AGENCY* by only presenting us with restricted sets of recommendations and options which are predetermined by hidden and generally unknowable 'black box' algorithms, to borrow Pasquale's (2015) term. Perhaps most importantly of all, it limits our *PERSONAL INTEGRITY* by treating us not as whole people but as disaggregated and reagggregated datasets which can be collected, manipulated, and, not infrequently, monetised. Thus, it seems preferable to avoid the conceptual slippage occasioned by the humanistic connotations of 'personalisation'; at best, I would say, we can call this phenomenon 'individualisation', but with the recognition that it is interwoven in complex ways with standardisation. We should also remember that although there is a place in education for individualised learning, notably in a supporting self-study role, it is potentially problematic in the classroom, where it may undercut decades of pedagogical trends towards collaborative, social constructivist learning that involves people interacting with other people.

Third, there are issues of bias and discrimination. You may recall a controversy from late 2020 when the high-profile computer scientist Timnit Gebru left Google due to a conflict over a then unpublished, co-authored paper about AI (Hao, 2020; Harris, 2023). In this paper, published the following year, the authors warn about the dangers of AI trained on large internet datasets: 'In accepting large amounts of web text as "representative" of "all" of humanity we risk perpetuating dominant viewpoints, increasing power imbalances, and further reifying inequality' (Bender et al., 2021, p. 614). The paper's clarion call has continued to echo through a swathe of subsequent posts, presentations, and publications highlighting the multifarious forms of bias – pertaining to (dis)ability, gender, sexuality, race, language, socioeconomic status, education, and more, not to mention

intersectional bias – which are present in training datasets. Notwithstanding recent efforts by technology companies to diversify the output of generative AI, underlying biases can and do continue to flow through into the texts and images generated, sometimes in obvious ways, and sometimes – perhaps more dangerously – in rather subtle ways (Cao et al., 2023; Cooper & Tang, 2024; Walker Rettberg, 2024).

Last, but certainly not least, there is the issue of environmental impact. All digital tools are energy hungry and environmentally destructive, but this is particularly true of the training and operation of AI, as has been well documented at least since the publication of Gebru's co-authored paper (Bender et al., 2021; Crawford, 2024; Saul et al., 2024; Verma & Tan, 2024). To borrow just two examples from Vicky Saumell's opening plenary at the 2024 International Association of Teachers of English as a Foreign Language (IATEFL) Conference in the UK: creating one AI-generated image can use as much electricity as fully charging your smartphone (AI for Education, 2023); and for somewhere between 5–50 responses, depending on timing and location, the (old) GPT-3 needed to 'drink' half a litre of water (Li et al., 2023). As Vicky herself pointed out, whenever we are refining our prompts over and over again, we should be considering the water consumption (Saumell, 2024).

As John Naughton wrote recently in *The Guardian*: 'AI technology is an environmental disaster in the making; the planet cannot be paved with datacentres' (Naughton, 2024). Gebru and her colleagues use the term ENVIRONMENTAL RACISM to highlight the fact that the negative climate and environmental costs of AI development and deployment are disproportionately borne by those in the Global South, alongside marginalised populations in the Global North, who benefit the least from these kinds of technologies (Bender et al., 2021, pp. 612–613).

In coming years, we will hear increasing mention of movements like GREEN IT and COMPUTING WITHIN LIMITS, focused on more environmentally sustainable approaches to our technologies (Nardi et al., 2018). Maybe it is also time, as Neil Selwyn has suggested, to start thinking about *ed-tech within limits*:

we need to anticipate the ongoing environmental degradation of the planet as radically upending existing assumptions about the continued use of digital technologies in education. (Selwyn, 2021, p. 501, italics in original)

If anything, this has become an even more pressing issue in the era of generative AI. Taken seriously, it may mean switching off our educational technologies at least some of the time.

We are beginning to see educational initiatives addressing the societal challenges of AI. In 2023, for example, the *Australian framework for generative artificial intelligence in schools* identified 'Human and Social Wellbeing' as its second of six core areas, with a focus on wellbeing, diversity, and human rights (Department of Education [Australia], 2023). The next year, the aforementioned *Generative AI idea pack for English language teachers* included an activity called 'How sustainable is AI?', designed to raise students' critical awareness of economic, social, and environmental issues (CUP [Cambridge University Press], 2024). Such educational initiatives represent the first tiny steps towards addressing very large problems.

5. Conclusion: The need for AI literacy

Hopefully it is evident by now that in order to ensure AI supports rather than detracts from the ongoing evolution of education, we as educators need to develop our AI literacy. Ultimately, we also need to help our students develop their AI literacy to prepare them for their future studies, careers, and lives, all of which will be infused with AI. Time does not permit me to go into great detail here, so I will wrap up with a few brief comments.

AI literacy is just one of a web of interrelated literacies, ranging from search (and now prompt) literacy through to ethical and critical literacy (Pegrum et al., 2022). This is abundantly clear in early

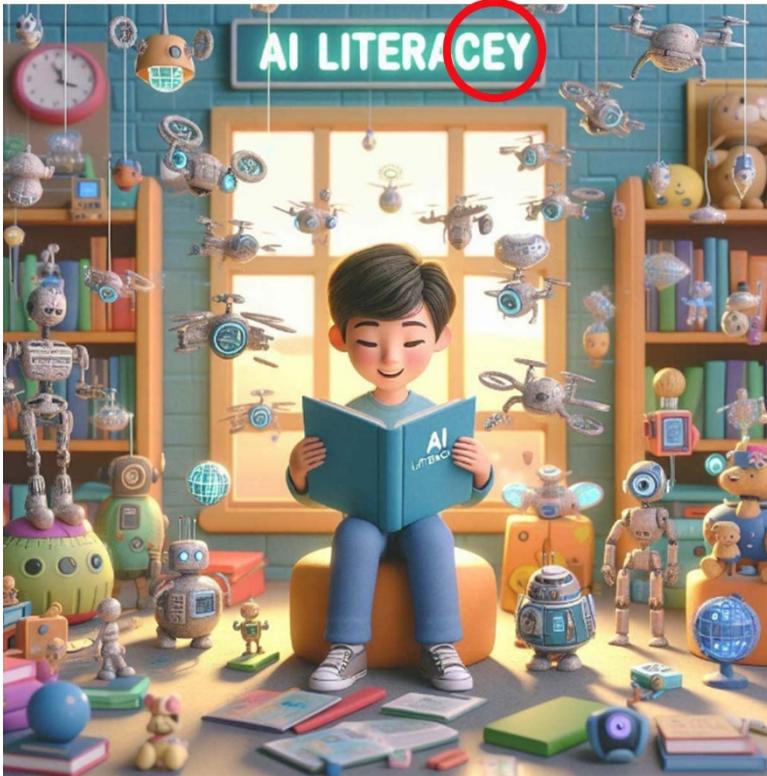


Figure 4. An image of AI literacy. DALL-E 3 image created by Mark Pegrum (2024) via Microsoft Copilot. Circle added by author.

attempts to define AI literacy. Prior to the recent generative AI explosion, the best-known conceptualisations included Long and Magerko's (2020) set of three competencies: critically evaluate AI, communicate and collaborate with AI, and use AI; and Ng et al.'s (2021) set of four aspects: know and understand AI, apply AI, evaluate and create AI, and AI ethics. In the post-ChatGPT era, the latter has been subsumed into a newer pedagogical framework highlighting the affective, behavioural, cognitive, and ethical dimensions of learning about AI (Ng et al., 2024).

My own ongoing work, conducted with my colleagues Ekaterina Tour from Monash University and Shem Macdonald from La Trobe University, focuses on the application of AI literacy to language learning in particular (Tour et al., *in press*). Our model delineates four domains: the operational, contextual, critical, and creative. This is a space to watch over the next two to three years as we will see a rush of presentations and papers attempting to define and operationalise AI literacy for the classroom, including for specific areas like language teaching and learning.

I will close with an image I asked DALL-E 3 to create for me, representing AI literacy. I am sure it will not take you long to spot something out of place here in the form of an extra letter 'E' in 'LITERACY' (see Fig. 4). In brief, there is no danger of this technology replacing teachers or textbooks any time soon, much less revolutionising entire education systems. Nevertheless, in a supporting role, it can be extremely valuable. As educators, we urgently need to develop our AI literacy, both to ensure that we are integrating AI effectively as a support for the ongoing evolution of education, and to ensure that we are well-placed to help our students begin developing their own AI literacy.

References

- AI for Education. (2023). AI's impact on the environment. *AI for education*. <https://www.aiforeducation.io/ai-resources/ai-impact-on-the-environment>
- Alm, A., & Ohashi, L. (2024). A worldwide study on language educators' initial response to ChatGPT. *Technology in Language Teaching and Learning*, 6(1), 1–23. doi: 10.29140/tl.v6n1.1141
- Anthology. (2023). *AI, academic integrity, and authentic assessment: An ethical path forward for education*. White paper. <https://www.anthology.com/paper/ai-academic-integrity-and-authentic-assessment-an-ethical-path-forward-for-education>
- Ardito, C. G. (2024). Generative AI detection in higher education assessments. *New Directions for Teaching and Learning* [early view]. doi: 10.1002/tl.20624
- ASM Research Publications. (2023, December 1). AI revolution in higher education. *ASM Digital Education Hub*. <https://researchpublications.info/ai-revolution-in-higher-education/>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *FAccT'21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). doi: 10.1145/3442188.3445922
- Brown, B. (2024, May 8). *Dr. Joy Buolamwini on Unmasking AI: My Mission to Protect What Is Human in a World of Machines* [interview podcast]. <https://brenebrown.com/podcast/unmasking-ai-my-mission-to-protect-what-is-human-in-a-world-of-machines/>
- Buchholz, K. (2023, July 7). Threads shoots past one million user mark at lightning speed. *Statista*. <https://www.statista.com/chart/29174/time-to-one-million-users/>
- Butler, J., Jaffe, S., Baym, N., Czerwinski, M., Iqbal, S., Nowak, K., Rintel, R., Sellen, A., Vorvoreanu, Hecht, M., B., & Teevan, J. (Eds.). (2023). *Microsoft new future of work report 2023*. Microsoft Research Tech Report MSR-TR-2023-34. <https://aka.ms/nfw2023>
- Caldwell, E. (2024, July 10). Duolingo Max: The future of language learning with AI? *Duolingo Guides*. <https://duolingoguides.com/duolingo-max/>
- Cao, Y., Zhou, L., Lee, S., Cabello, L., Chen, M., & Hershovich, D. (2023). Assessing cross-cultural alignment between ChatGPT and human societies: An empirical study. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)* (pp. 53–67). Association for Computational Linguistics. doi: 10.18653/v1/2023.c3nlp-1.7
- Cooper, G., & Tang, K.-S. (2024). Pixels and pedagogy: Examining science education imagery by generative artificial intelligence. *Journal of Science Education and Technology*, 33, 556–568. doi: 10.1007/s10956-024-10104-0
- Crawford, K. (2024, February 20). Generative AI's environmental costs are soaring – And mostly secret. *Nature*. <https://www.nature.com/articles/d41586-024-00478-x>
- Creely, E. (2024). Exploring the role of generative AI in enhancing language learning: Opportunities and challenges. *International Journal of Changes in Education*, 1(3), 158–167. doi: 10.47852/bonviewIJCE42022495
- CUP [Cambridge University Press]. (2024). *Generative AI idea pack for English language teachers*. Cambridge University Press and Assessment. <https://cambunipress.prod.acquia-sites.com/sites/default/files/media/documents/GenAI-Idea-Pack-for-English-language-teachers.pdf>
- Dan. (2024, August 29). The most overused ChatGPT words. *Plus*. <https://www.plusdocs.com/blog/the-most-overused-chatgpt-words>
- Davoodifard, M., & Eskin, D. (2024). Exploring the role of generative AI in second language education: Insights for instruction, learning, and assessment. *Studies in Applied Linguistics and TESOL*, 24(1), 21–23. doi: 10.52214/salt.v24i1.12863
- Department of Education [Australia]. (2023). *Australian framework for generative artificial intelligence in schools*. Commonwealth of Australia. <https://www.education.gov.au/schooling/resources/australian-framework-generative-artificial-intelligence-ai-schools>
- Edmett, A., Ichaporia, N., Crompton, H., & Crichton, R. (2023). Artificial intelligence and English language teaching: Preparing for the future. British Council. doi: 10.57884/78EA-3C69
- Evmenova, A. S., Borup, J., & Shin, J. K. (2024). Harnessing the power of generative AI to support ALL learners. *TechTrends*, 68, 820–831. doi: 10.1007/s11528-024-00966-x
- Focus Education. (2024). The revolution of international education: A glimpse into how AI is transforming learning. *Focus Education*. <https://www.focusedu.org/the-revolution-of-international-education-a-glimpse-into-how-ai-is-transforming-learning/>
- Foltýnek, T., Dlabolová, D., Anohina-Naumeca, A., Razi, S., Kravjar, J., Kamzola, L., Guerrero-Dib, J., Çelik, Ö., & Weber-Wulff, D. (2020). Testing of support tools for plagiarism detection. *International Journal of Educational Technology in Higher Education*, 17(46). doi: 10.1186/s41239-020-00192-4
- Franks, B. (2023, April 26). Generative AI: It's all a hallucination! *Medium*. <https://medium.com/analytics-matters/generative-ai-its-all-a-hallucination-6b8798445044>
- Furze, L. (2024, October 30). AI and the myth of personalised learning. *Leon Furze*. <https://leonfurze.com/2024/10/30/ai-and-the-myth-of-personalised-learning/>

- Gattupalli, S., Maloy, R. W., & Edwards, S. A. (2023). *Prompt literacy: A pivotal educational skill in the age of AI*. ScholarWorks@UMass Amherst: College of Education Working Papers and Reports Series. https://scholarworks.umass.edu/education_working_papers/6/
- Godwin-Jones, R. (2024). Distributed agency in second language learning and teaching through generative AI. *Language Learning and Technology*, 28(2), 5–31. <https://hdl.handle.net/10125/73570>
- Gurung, D. (2022, December 26). Is artificial intelligence necessary for education revolution? *LinkedIn*. <https://www.linkedin.com/pulse/artificial-intelligence-necessary-education-dhiraj-gurung/>
- Hao, K. (2020, December 4). We read the paper that forced Timnit Gebru out of Google. Here's what it says. *MIT Technology Review*. <https://www.technologyreview.com/2020/12/04/1013294/google-ai-ethics-research-paper-forced-out-timnit-gebru/>
- Harris, J. (2023, May 22). 'There was all sorts of toxic behaviour': Timnit Gebru on her sacking by Google, AI's dangers and big tech's biases. *The Guardian*. <https://www.theguardian.com/lifeandstyle/2023/may/22/there-was-all-sorts-of-toxic-behaviour-timnit-gebru-on-her-sacking-by-google-ais-dangers-and-big-techs-biases>
- Hwang, Y., Lee, J. H., & Shin, D. (2023). What is prompt literacy? An exploratory study of language learners' development of new literacy skill [sic] using generative AI. *arXiv*. doi: 10.48550/arXiv.2311.05373
- @JeremyNguyenPhD. (2024, March 30). Are medical studies being written with ChatGPT? X. <https://x.com/JeremyNguyenPhD/status/1774021645709295840>
- Jiang, Y., Hao, J., Fauss, M., & Li, C. (2024). Detecting ChatGPT-generated essays in a large-scale writing assessment: Is there a bias against non-native English speakers? *Computers and Education*, 217. doi: 10.1016/j.compedu.2024.105070
- Landes, A. (2024, September). Microsoft's AI falsely accuses reporter of crimes. *MSN*. <https://www.msn.com/en-us/money/other/microsofts-ai-falsely-accuses-reporter-of-crimes/ar-AA1pw4XE>
- Learnboost. (n.d.). Learn faster and better with AI: How artificial intelligence improves learning. *Learnboost*. <https://www.learnboost.tech/en/blog/schneller-und-besser-lernen-mit-ki-wie-knstliche-intelligenz-das-lernen-verbessert>
- Li, P., Yang, J., Islam, M. A., & Ren, S. (2023) Making AI less 'thirsty': Uncovering and addressing the secret water footprint of AI models. *arXiv*. V3. <https://arxiv.org/abs/2304.03271>
- Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns*, 4(7). [https://www.cell.com/patterns/fulltext/S2666-3899\(23\)00130-7](https://www.cell.com/patterns/fulltext/S2666-3899(23)00130-7)
- Lim, T. (2023, October 10). Revolutionizing K-12 education: The power of AI homework companions. *Forbes*. <https://www.forbes.com/councils/forbescoachescouncil/2023/10/10/revolutionizing-k-12-education-the-power-of-ai-homework-companions/>
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1–16). doi: 10.1145/3313831.3376727
- Low, D. (2024, September 12). The age of Apple Intelligence and the rise of small language models. *Waracle*. <https://waracle.com/insights/news/the-age-of-apple-intelligence-and-the-rise-of-small-language-models/>
- Marr, B. (2024, August 19). Why AI models are collapsing and what it means for the future of technology. *Forbes*. <https://www.forbes.com/sites/bernardmarr/2024/08/19/why-ai-models-are-collapsing-and-what-it-means-for-the-future-of-technology/>
- Matellio. (2022, October 18). AI solutions for education industry that can bring revolution. *Matellio*. <https://www.matellio.com/blog/ai-solutions-for-education/>
- Mayers, L. (2023, April 6). Hepburn mayor may sue OpenAI for defamation over false ChatGPT claims. *ABC News [Australia]*. <https://www.abc.net.au/news/2023-04-06/hepburn-mayor-flags-legal-action-over-false-chatgpt-claims/102195610>
- Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>
- Mollick, E. (2024a). *Co-intelligence: Living and working with AI*. W. H. Allen.
- Mollick, E. (2024b, October 20). Thinking like an AI. *One Useful Thing*. <https://www.oneusefulting.org/p/thinking-like-an-ai>
- Nardi, B., Tomlinson, B., Patterson, D. J., Chen, J., Pargman, D., Raghavan, B., & Penzenstadler, B. (2018). Computing within limits. *Communications of the ACM*, 61(10), 86–93. doi: 10.1145/3183582
- Naughton, J. (2024, April 13). From boom to burst, the AI bubble is only heading in one direction. *The Guardian*. <https://www.theguardian.com/commentisfree/2024/apr/13/from-boom-to-burst-the-ai-bubble-is-only-heading-in-one-direction>
- Ng, D. T. K., Leung, J. K. L., Chu, K. W. S., & Qiao, M. S. (2021). AI literacy: Definition, teaching, evaluation and ethical issues. *Proceedings of the Association for Information Science and Technology*, 58(1), 504–509. doi: 10.1002/pra.2.487
- Ng, D. T. K., Wu, W., Leung, J. K. L., Chiu, T. K. F., & Chu, S. K. W. (2024). Design and validation of the AI literacy questionnaire: The affective, behavioural, cognitive and ethical approach. *British Journal of Educational Technology*, 55(3), 1082–1104. doi: 10.1111/bjet.13411
- @OfficialLoganK. (2024, February 24). The final form for ChatGPT is not chat :). X. <https://x.com/OfficialLoganK/status/1761081383370326450>
- Paoletti, G., Gonzalez-Billandon, J., & Kégl, B. 2024. A call for embodied AI. *arXiv*, V4. doi: 10.48550/arXiv.2402.03824
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

- Pasquale, F. (2020). *New laws of robotics: Defending human expertise in the age of AI*. Belknap Press.
- Pegrum, M. (2019). *Mobile lenses on learning: Languages and literacies on the move*. Springer.
- Pegrum, M. (2024). Generative AI. *Digital Learning*. <https://markpegrum.com/tools-for-digital-learning/generative-ai/>
- Pegrum, M., Hockly, N., & Dudeney, G. (2022). *Digital literacies* (2nd ed.). Routledge.
- Picardo, J. (2023). Navigating the AI revolution in education: What do schools need to consider? *José Picardo*. <https://www.josepicardo.com/education/navigating-the-ai-revolution-in-education/>
- PR Newswire. (2024, February 29). Figure raises \$675M at \$2.6B valuation and signs collaboration agreement with OpenAI. *PR Newswire*. <https://www.prnewswire.com/news-releases/figure-raises-675m-at-2-6b-valuation-and-signs-collaboration-agreement-with-openai-302074897.html>
- Purtill, J. (2023, April 2). Prompt engineers share tips on using ChatGPT, Midjourney, Bing Chat and other generative AI. *ABC News [Australia]*. <https://www.abc.net.au/news/science/2023-04-02/prompt-engineers-share-their-tips-on-using-chatgpt-generative-ai/102165132>
- Ray, S. (2023, June 8). OpenAI sued for defamation after ChatGPT generates fake complaint accusing man of embezzlement. *Forbes*. <https://www.forbes.com/sites/siladityaray/2023/06/08/openai-sued-for-defamation-after-chatgpt-generates-fake-complaint-accusing-man-of-embezzlement/>
- Reich, J. (2020). *Failure to disrupt: Why technology alone can't transform education*. Harvard University Press.
- Sadasivan, V. S., Kumar, A., Balasubramanian, S., Wang, W., & Feizi, S. 2024. Can AI-generated text be reliably detected? *arXiv*, V3. doi: 10.48550/arXiv.2303.11156
- Saul, J., Nicoletti, L., Rai, S., Bass, D., King, I., & Duggan, J. (2024, June 22). AI is already wreaking havoc on global power systems. *Bloomberg*. <https://www.bloomberg.com/graphics/2024-ai-data-centers-power-grids/>
- Saumell, V. (2024, April 16). The AI factor: Have we figured it out yet? IATEFL Conference 2024 plenary. *IATEFL*. <https://www.iatefl.org/resources/iatefl-conference-2024-plenary-ai-factor-have-we-figured-it-out>
- Selwyn, N. (2021). Ed-tech within limits: Anticipating educational technology in times of environmental crisis. *E-learning and Digital Media*, 18(5), 496–510. doi: 10.1177/204275302111022951
- Sharples, M. (2023). Generative AI and education futures [video]. *University College London*. <https://www.ucl.ac.uk/teaching-learning/case-studies/2023/aug/generative-ai-and-education-futures>
- Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, 631, 755–759. doi: 10.1038/s41586-024-07566-y
- Thompson, T. (2023, November 2). Education revolution: The rise of the AI-powered personalised classroom. *Teraflow*. <https://www.teraflow.ai/education-revolution-the-rise-of-the-ai-powered-personalised-classroom/>
- Tour, E., Pegrum, M., & Macdonald, S. (in press). AI literacy practices. In L. McCallum, & D. Tafazoli (Eds.), *The Palgrave Encyclopedia of Computer-Assisted Language Learning*. Palgrave.
- Verma, P., & Tan, S. (2024, September 18). A bottle of water per email: The hidden environmental costs of using AI chatbots. *The Washington Post*. <https://www.washingtonpost.com/technology/2024/09/18/energy-ai-use-electricity-water-data-centers/>
- Walker Rettberg, J. (2024, January 16). How generative AI endangers cultural narratives. *Issues in Science and Technology*. <https://issues.org/generative-ai-cultural-narratives-rettberg/>
- Warzel, C. (2023, February 8). The most important job skill of this century. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2023/02/openai-text-models-google-search-engine-bard-chatbot-chatgpt-prompt-writing/672991/>
- Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19(26). doi: 10.1007/s40979-023-00146-z
- West, K. N. (2023, December 18). Understanding the AI-driven workforce shift. *Medium*. <https://mskimsworld1.medium.com/understanding-the-ai-driven-workforce-shift-c328772fc0ad>
- Zia, T. (2023, July 19). Transforming education: AI-powered personalized learning revolution. *Techopedia*. <https://www.techopedia.com/transforming-education-ai-powered-personalized-learning-revolution>

Mark Pegrum lectures on digital learning in the Graduate School of Education at The University of Western Australia in Perth, where he is also the Deputy Head of School (International). His current research focuses on emerging technologies, including generative artificial intelligence (AI); digital literacies, including attentional literacy and AI literacy; and the forms digital learning takes in diverse contexts across the Global North and South. His books include: *Mobile learning: Languages, literacies and cultures* (Palgrave Macmillan, 2014); *Mobile lenses on learning: Languages and literacies on the move* (Springer, 2019); and *Digital literacies* (2nd ed., co-authored with Nicky Hockly and Gavin Dudeney, Routledge, 2022). Further details can be found on his website at <https://markpegrum.com/about-mark-pegrum/biodata/>.

Cite this article: Pegrum, M. (2025). From revolution to evolution: What generative AI really means for language learning. *Language Teaching*, 1–17. <https://doi.org/10.1017/S0261444825000151>