AI Literacy in the Language Classroom

Facilitating critical, ethical and responsible use

Zsófia Menyhei and Joanna Szoke

DELTA TEACHER DEVELOPMENT SERIES

DELTA Publishing

DELTA TEACHER DEVELOPMENT SERIES

Series editor: Riccardo Chiappini

AI Literacy in the Language Classroom

Facilitating critical, ethical and responsible use

Zsófia Menyhei and Joanna Szoke

1st edition 1 5 4 3 2 1 | 2029 28 27 26 25
The last figure shown denotes the year of impression.

DELTA Publishing, an imprint of the Ernst Klett Sprachen GmbH, 2025
www.deltapublishing.co.uk
www.klett-sprachen.de/delta

Authors: Zsófia Menyhei, Joanna Szoke, assisted by GenAI as described in Part C, pp 107–112
Series Editor: Riccardo Chiappini
Layout and typesetting: Datagroup Int. SRL
Cover picture: 123RF.com (austler)
Cover: Andreas Drabarek
Printing and binding: Plump Druck & Medien GmbH, Rolandsecker Weg 33, 53619 Rheinbreitbach, Germany
Information and additional products/material for this series can be found here: https://www.deltapublishing.co.uk/series/delta-teacher-development-series-570/?page_id=1#

Printed in Germany
ISBN: 978-3-12-501791-7

From the authors

During my university years as a TEFL and Communication Studies major I found it fascinating to learn about emerging technologies, and how they shape and influence society. When I became an EFL teacher, and later a teacher trainer, I was more interested in the practical aspects. I was always on the hunt for educational technologies that could help me do my job more efficiently and creatively, and help my students achieve their goals. When generative AI technologies entered the scene, I got to explore exciting new 'tools', but it also became clear that GenAI is fundamentally different from the technologies that I had worked with before, and that it would change quite a few things in education and society as a whole. My colleague, Jo and I decided to introduce a focus on AI literacy in our teacher training courses, and participant feedback has reassured us just how important this endeavour really is.

Zsófi

For as long as I can remember, I have been eagerly tinkering with computer software. I have always been interested in how computers can be utilized in everyday life, so it came naturally that during and after finishing my teacher training degrees, I got actively involved in online teaching, teacher training, and digital materials development. When generative AI technologies emerged, I immediately felt that this is something that will potentially revolutionize education if we approach it in the right way. With my colleague, Zsofi, we instantly felt the importance of adjusting our syllabi to include a stronger focus on AI literacy. We believe that embracing GenAI technologies in teaching is the way forward in this new age.

Jo

Contents

Contents

Resources and tools

A selection of the main resources and tools referred to in this book.

A

Introduction

In Part A, we provide an overview of the key concepts of AI. We begin by exploring common teacher reactions to the rise of AI, examining claims about its current capabilities and potential future developments. Next, we discuss how AI is applied in real-world scenarios and its role in language education. Key terms are highlighted in **bold** when mentioned for the first time.

Reactions to AI

Reflection

What was your initial reaction when ChatGPT first appeared? How did your colleagues respond?

Have your views or those of your colleagues changed since then?

Do you or your colleagues currently use ChatGPT? If so, what do you use it for?

Artificial intelligence (AI) has been around for quite some time now. It powers the weather apps we rely on, autocorrects our messages and recommends songs or movies on our favourite streaming platforms. Yet, it wasn't until late 2022, with the release of ChatGPT3, that AI technologies truly started to dominate public discourse. This new form of AI could generate natural, human-like responses and create images based on our requests – and in a matter of seconds! Educators around the world began to realise the potential of **generative AI** (GenAI), and their responses ranged from sheer panic and overwhelming excitement to complete indifference.

These diverse reactions were easy to spot in our own staffrooms. Some colleagues immediately banned GenAI tools like ChatGPT in their classrooms, and hailed the advent of **AI detectors**, which, it was hoped, would help distinguish 'real' student essays from AI generated ones. Others, however, enthusiastically embraced the new tools, creating classroom materials – grammar gapfill activities, short texts for jigsaw reading, images for vocabulary building – and, to everyone else's amazement, even providing feedback on tests and homework! Then there were those who simply refused to acknowledge that anything out of the ordinary had happened, believing AI would have little to no impact on their classrooms or teaching practices. It is easy to relate to all of these teachers. AI offers enormous benefits, especially for the busy teacher, with the potential to reduce our workload by automating various administrative tasks, and the potential to personalise learning experiences by providing our students with tailored materials and real-time feedback. And yes, AI can also seem scary in many ways.

AI is frequently associated with phrases like 'paradigm shift', 'societal implications,' and 'Fourth Industrial Revolution', and many of us weren't formally trained to navigate these changes. It raises serious concerns about **data privacy**, **ownership**, **misinformation**, **bias**, **accessibility**, as well as **misuse** of and **overreliance** on AI assistance. Not to mention the speed at which AI technologies are advancing, making it all the more difficult to stay in the loop, or to hypothesise about what the future holds.

Claims about AI

It has been a few years since we first asked ChatGPT to write a Shakespearean sonnet about cats, yet we find that AI is still a daily topic of conversation among colleagues. Institutions are pressured to provide staff and students with guidance on ethical and responsible use, but this has yet to occur in many schools. At the time of writing, there is still a great deal of uncertainty about an array of social and legal issues, the majority of which are connected to the field of education in one way or another. While many questions remain unanswered, we can already make some key, definitive claims about AI:

AI is here to stay: AI technologies are increasingly being integrated in various sectors of the economy such as finance and healthcare. This is why we strongly believe they should be integrated in all school curricula as well. Banning AI from the classroom isn't going to prevent its emergence in society.

AI has affordances as well as risks: AI offers numerous benefits, but it also presents several challenges. We believe schools should accept their responsibility of preparing students for their future lives and careers, which means that they should provide them with a strong understanding of AI, its advantages as well as its pitfalls.

Teachers need to adapt to the new technological landscape: As AI continues to advance, it compels educators to critically reflect on their teaching practices and the purpose behind them. Teachers may discover that certain tasks can easily be accomplished by AI-powered tools – often requiring minimal effort. In some cases, they might have already encountered instances where students used these tools in their work without providing proper attribution. Sure, they may have AI detectors at their disposal, but the accuracy of such tools is questionable, since they tend to produce false positives (incorrectly flagging content created by humans as AI-generated) and false negatives (incorrectly identifying AI-generated content as created by humans). Therefore, teachers are advised to use them with caution – as one of many solutions to ensure academic integrity (Turnitin, 2023) or avoid relying on them altogether (Code.org et al., 2023). As an alternative solution, teachers might consider other ways to adapt their teaching and assessment to the changing circumstances. (For more on this, please see the ⏩ **Teaching and assessment** section below, as well as Part C).

All of this leads us to the main focus of this book: AI literacy.

As the authors of this book, we are convinced that, since AI is becoming more and more pervasive in our society, all citizens should have the basic knowledge and skills related to AI and its use. We also believe that the best place to start learning about AI is the school, not just in computer science classes, but as part of a curriculum-wide approach including non-technical aspects as well. Therefore, it is also our conviction that teachers don't need a degree in data science to facilitate students' development in this area. Language teachers with no technical background in AI can do so perfectly well – for example, by raising students' awareness about the real-life uses of AI that surround them, and helping them think critically about AI-generated content and the ethical issues associated with it. The aim of this book is to help readers get to grips with AI technologies, and offer them practical ideas to gain new or greater confidence to integrate AI literacy development in their classrooms.

In Part A, we:
- break down what AI is and identify the different types of AI.
- clarify the key concepts connected to how AI works.
- explore some real-world applications of AI.
- discuss how AI technologies can benefit language education.
- explore AI's risks and ethical issues, focusing on their impact on education.
- discuss the concept and areas of AI literacy.
- explain why AI literacy development is necessary in the foreign language classroom.
- provide tips on how to integrate AI literacy in the classroom.

AI: Key concepts

What is AI?

Reflection
Think about the websites and apps you use on your computer. How many of them have AI built in? Can you spot these features? Do you know what they are capable of?

Today, almost everything seems to come with a label indicating an AI feature – from smartphones and home assistants to photo-editing apps and streaming services. Since ChatGPT's release in 2022, AI has become a popular buzzword. But what's really behind it? One key feature of AI literacy is critical thinking, so let us now take a closer look at what AI is. Understanding this will help us see what's required to make it work and how that affects its users, especially in the field of language learning.

AI is 'the science and engineering of making intelligent machines, especially intelligent computer programs' (McCarthy, 2007). John McCarthy, often called 'the father of AI' coined this term in as early as 1956.

The field has come a long way since then, and today we understand it as 'machines that imitate some features of human intelligence, such as perception, learning, reasoning, problem-solving, language interaction and creative work' (COMEST, 2019; UNESCO, 2022, p. 9). For instance, perception is a key aspect of how self-driving cars work. These vehicles use computer vision to perceive their environment, imitating the way humans rely on visual and spatial perception when driving.

There are various types of AI, which are often grouped into paired categories. The first one is **Classical AI** and **Machine learning** (ML). Classical AI, also called 'rule-based AI' or 'good-old-fashioned AI', refers to predefined rules, such as conditional logic (*if-then* statements), to solve problems and make decisions (Miao et al., 2021). An example is a chess programme that follows a set of predefined rules to choose its next move. Machine learning, on the other hand, processes vast amounts of data to identify patterns, builds models from these patterns, and uses the latter to make predictions. Unlike classical AI, ML can 'learn' without explicit programming – it can improve its own performance over time. For example, ML is used to filter spam in emails and gets better at this as it processes more data over time.

Another binary distinction in AI types is between **predictive AI** and generative AI (GenAI) (ISTE, 2023). Predictive AI analyses patterns and trends in the data to make predictions about future events and outcomes. For example, online stores use predictive AI to recommend products based on items the customer has previously viewed or purchased. GenAI, on the other hand, is called 'generative' because it creates new content, like text, images or sound, by applying learned patterns.

AI can also be categorised as either **weak AI** or **strong AI**. Weak AI, or 'narrow AI', is the AI we use today. It excels at performing specific tasks but can't apply its abilities beyond that domain. For example, the AI that drives a car can't diagnose medical conditions (Miao et al., 2021). In contrast, strong AI, which is also called **artificial general intelligence** (AGI), currently exists only in theory. This AI would possess human-like understanding and consciousness, if it ever becomes a reality.

How does AI work?

Dataset, algorithm, prediction

The basic ingredient of all AI-driven technologies is data. The **algorithm** is trained on a dataset, and the larger and more diverse the data, the more accurate the algorithm becomes.

Once we have enough data, we move onto the modelling and decision-making process, which is a black box for most people who aren't AI specialists (thus the term 'black-box decision making' was born, which will be explained later in the chapter).

Finally, we arrive at the so-called **prediction**, which is the 'decision' that the algorithm makes. But what is this magic word that is doing all the work behind the scenes? It is called 'algorithm'. An algorithm is a set of steps or instructions given to a computer programme to perform (Alpaydin, 2021).

One of the simplest computer algorithms is searching for information based on a keyword. For example, in order to find all documents containing the word *report*, a computer follows a series of steps – an algorithm – to do this. AI technologies, however, rely on far more complex algorithms. These systems first analyse input data to detect **patterns**. Pattern recognition can happen with or without human assistance but, once identified, the AI makes predictions and takes action based on those patterns.

Let's look at a concrete, personal example. Jo loves hiking and exploring the outdoors, where she comes across various plants. After watching a video about a dangerous plant called the giant hogweed, which can cause severe burns if touched, she learnt there is also a harmless version called the common hogweed. This sparked her curiosity, so she started searching for this plant in the wild – essentially, she started *collecting data*. While a true scientist might have tested the plants by touching them, she preferred a more hands-off approach. The more plants she

came across, the more aware she became of the distinct features of these two plants, which helped her confidently differentiate between them. If she wanted to, she could use machine learning to help her distinguish the two types of plants.

This means that Jo would need to train the machine on large amounts of **training data** – hundreds of images showing the distinguishing features of the giant and common hogweed. This data should be well 'prepared', which includes making sure that neither of the plants should be under- or over-represented in Jo's dataset. If the machine has **representation problems** or **training bias**, it will make incorrect predictions. For example, if all the model sees is images of the dangerous version of the plant, it won't be able to identify the harmless version because it lacks information about it.

However, if a balanced set of data is provided, the model will be able to make more accurate predictions and draw better conclusions. Once she has fed the balanced training data into the machine, it should be tested with a set of **evaluation data** – typically a smaller set than the training data. Jo would then use additional samples of the two types of hogweed to see if her machine can correctly distinguish between them. If it performs well, she can conclude that she has trained a machine that can accurately *predict* whether a hogweed plant is dangerous or benign, much like our brain makes such judgements. The final step would be to put it to the ultimate test and take it into the wild – the real-world setting.

Supervised and unsupervised learning

Training a machine to sort everything in the world would take forever, so instead we teach it to learn by recognising patterns in data; in other words, to generalise the rule.

In **supervised learning**, the machine is given pre-defined human-labelled datasets, like images of the giant hogweed plant labelled as such. This allows it to identify giant hogweed in new images. Another common use of supervised learning is classifying emails as spam or not spam.

In **unsupervised learning**, the datasets are unlabelled. This means that the machine works with unlabelled data to discover hidden patterns on its own. For example, recommendation systems use unsupervised learning to group users with similar interests and suggest products, movies, or music based on these groups – the so-called **clusters**. While human intervention is needed to validate the results, the machine does the pattern-finding without direct human input (Delua, 2021).

Neural networks and deep learning

Let's take another practical example. Let's say that I'm getting ready for a picnic, but as I look out the window, I see dark clouds in the sky. Based on thousands of previous life experiences, I might make the prediction that there is a storm coming, because my brain has found a causal connection between dark clouds and storm. Every time my prediction is accurate, the connection in my brain gets stronger – it gains **weight**.

This idea is based on one of the earliest machine learning models, which compares human learning to machine learning. According to this model, machines learn with the help of **neural networks**, which are made up of several layers between the input (the data) and the output (the prediction). This model, popular in the 60s, declined in use, but it has seen a resurgence with the advent of **deep learning** (DL) (Morales & Escalante, 2022).

Deep Learning uses neural networks but on a much larger scale, with multiple hidden layers, much more data, and larger processing speed. It can be supervised or unsupervised, but because of its immense size, it typically works in an unsupervised manner. Basically, this upscaling has made it possible to effectively use DL in image recognition, natural language processing, fraud detection, or faster search query completion, to just name a few (IBM Data and AI Team, 2023).

Issues of transparency

Unfortunately, however, this growth in speed and scale also means that the process is becoming so complex and contains so many invisible layers of minor decisions that we can practically consider it impenetrable. That is why it has received the label **black box decision-making**. Just as the workings of black holes in outer space are hidden from view, deep learning decision-making processes are completely opaque to the average person.

This in turn also makes certain AI technologies mysterious and difficult to trust because it becomes unclear how we can be sure that a decision it makes or the answer it gives is actually correct. We will address this issue in more detail later in the book.

In order to make AI more trustworthy, transparent models are needed (Ali et al., 2023), which can explain how the models reach particular decisions. These are also known as **white-box decision-making** models (Loyola-González, 2019) and are already widely used in the fields of medicine and finance (ibid.) to provide more accountability and transparency.

So far we have the following ingredients on our list to create the AI systems that surround us:

1. Our machine or computer programme needs to be able to perform tasks step by step (algorithmically) that would otherwise require some sort of human intelligence [=AI].
2. This machine or programme must have the ability to recognise patterns in datasets to make predictions and then decisions [=ML].
3. It must be able to do this at scale to satisfy billions of queries per second [=DL].

How does GenAI work?

We aren't at the end yet! There are two more important concepts that we need to understand to have a better picture of how we can communicate with a chatbot like ChatGPT so easily and how it is capable of generating, summarising, rewriting, correcting, translating, and analysing all kinds of texts. These are **Large Language Models** (LLMs), and **Natural Language Processing** (NLP).

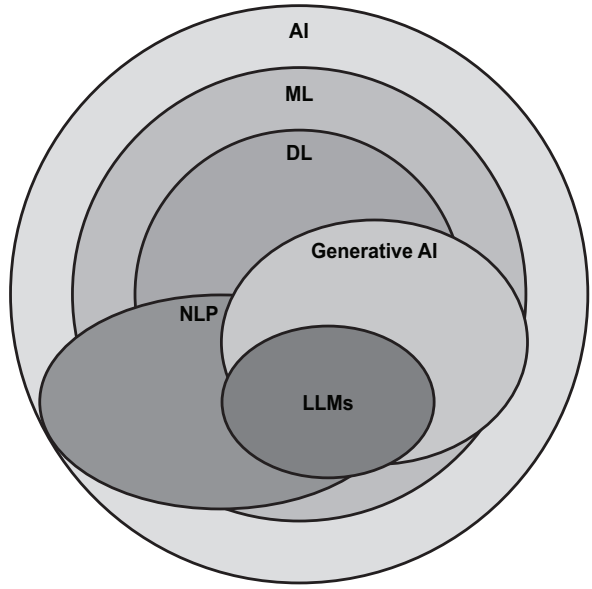

Hierarchical and intersectional representation of AI subfields and technologies

If the reader has already had a conversation with an AI chatbot, such as OpenAI's ChatGPT or Google's Gemini, they might have realised that they can use their own words, including sentence fragments or single-word instructions to communicate with the tool. The instructions given or questions asked to the chatbot are known as **prompts**. These prompts can be phrased just like a text to a friend. In response, the chatbot generates text or code that sounds human. The output can then be adjusted to be longer, shorter, more casual, more formal, filled with slang or enriched with advanced vocabulary, depending on your needs.