INVITED PAPER

The future of language learning in the era of ChatGPT

Nick Saville[a,1]

aCambridge University Press and Assessment, ALTE Secretary-General

(published: 31/10/2023)

1. Introduction

Originating in the 1950s, Artificial Intelligence (AI) refers to the capability of machines to emulate human-like intelligence by executing tasks humans can do, such as understanding language, and problem-solving. However, in recent years AI has become a buzzword in the media and it needs to be demystified and explained more clearly if we are to take advantage of it in educational contexts. Practitioners must now navigate potential pitfalls and risks, including ethical considerations, and come to terms with the implications.

2. The era of ChatGPT

"The Age of AI has begun", Bill Gates claimed in his blog on 21 March 2023, and like Gates, many commentators now think that AI will be as revolutionary as the internet in the 1980s and the invention of smart phones that ushered in the 4th Industrial Revolution about 15 years ago (the era of Data and Devices). However, the speed of change and potential for disruption, has taken many by surprise. In less than two months following its release in November 2022, Open AI's ChatGPT gained over 100 million users, and for the first time the global public became familiar with the jargon that accompanies this kind of *generative AI* (GenAI): large language models (LLMs), machine learning (ML), natural language processing (NLP), black box algorithms, amongst others.

As the influence of these latest AI models widens and developments continue at a breakneck pace, established ways of thinking and current practices will be disrupted in many sectors of society, including language education. The challenge for educationalists is to come to terms with the key concepts and jargon and to

---

[1] corresponding author - email: nick.saville@cambridge.org

understand better the risks and benefits that educational uses of AI might bring in their own teaching, learning and assessment contexts.

In the first months after its release there were mixed reactions to ChatGPT from politicians and policy makers, and in some cases there were attempts to ban its use in educational and other societal contexts. This was not an unprecedented early response to disruptive technologies, and there have been several other examples in the past where the new tech was banned from the classroom when it first came out. Perhaps the most memorable was the pocket calculator which was first released in 1972 by Hewlett-Packard. Some educationalists believed that this handy device would have a negative effect on children's learning of arithmetic, and initially calculators were not allowed into schools in some jurisdictions. However, once the devices became popular and students had access to them at home, it became clear that they were being used, despite the policy e.g., to do homework. Moreover, it was realised quite quickly that the *potential benefits* of using calculators outweighed the risks. Similarly, there was concern over Wikipedia in the early 2000s based on fears that students might use it to avoid conducting their own research in libraries. In both cases, teaching and learning practices adapted to accommodate the new tech. So, a more realistic option is not to ban ChatGPT and similar types of AI, but to seek to understand them better and to explore ways to use them with caution.

In considering the future of language learning in light of the recent AI developments, a number of questions can be posed related to the role that humans will play in combination with machines:

- In what ways will practitioners – teachers, assessment specialists etc. – need to adapt their roles and current practices to benefit from the opportunities that AI-based systems can offer?
- What practical steps can be taken to ensure that AI-based systems are used to make improvements over current practices?
- How can policy makers and practitioners working together address the risks associated with the latest AI models?

In using artificial intelligence (AI) for educational purposes, teaching and learning must remain the central concern for policy-makers and practitioners. As a first step in answering these questions, it worth remembering that existing challenges in delivering educational goals supported by educational technologies (EdTech) are the starting point. Hitherto the take up take been slow and difficult to integrate into existing practices on a large scale. This is partly due the wide range of, sometimes competing, objectives that formal education has to deliver.

In language education, the priorities for language educators, both at policy level and in schools, are still related to the following: *delivering the curriculum; setting and raising standards; improving teaching and learning outcomes;* and *innovating and managing change.* Added to this list most recently are concerns for *diversity* and *sustainability*, in light of global crises related to migration and climate change. The real world impacts of language education in schools, relate to important societal goals, such as the need for higher levels of plurilingual communication skills and intercultural competence, as well as transferable competences that enhance employability (cf. 21st century skills).

Taken together, this means that the emerging technologies need to enhance the communicative approach to language teaching and ensure that the social and educational values that are associated with it remain the priority. The well-being of learners must continue to be emphasised alongside the cognitive and meta-cognitive aspects of their development.

EdTech has been around for many years, and traditional models of AI already play an important role in language assessment, for example, in carrying out the automated assessment of writing tasks and compositions. However, it has proved difficult to scale up digital technology and to integrate it widely into school contexts. This in turn has limited the use of AI for educational purposes as it depends on the digital tech, and especially the collection and use of digital data. But the last three years have provided a watershed moment brought about by the following two events: a) the pandemic (2020-22) that forced us all to adopt new practices using the available digital tech (e.g. video conferencing); b) the wider familiarity with AI across all sectors of society since the arrival of ChatGTP in 2022.

With this in mind, the questions posed above can be answered by considering the following three points:
- Hybridity as a design principle;
- Understanding AI in order to retain focus on educational values;
- Ethical considerations.

I believe that by addressing these together, we can ensure that the *Technology Tail* does not wag *the Educational Dog,* and that we can open up exciting new possibilities that have so far been impractical on a wide scale: e.g., personalised learning experiences; interactive learning environments that enhance accessibility and improve learner engagement; streamlined assessment processes that offer real-time feedback for diagnostic and formative purposes; etc.

3. Hybridity

Hybridity (or hybridisation) became an established concept during the pandemic (2020-22) and a range of blended solutions were deployed as a necessary response to the emergency conditions: e.g., physical plus online classes; synchronous and asynchronous sessions; etc. This was not entirely new. *Cloud-based learning* had been around for nearly 10 years, supplementing in-class learning, and the *flipped classroom* concept was already established. So the existing tech that was available for *blended solution* became widely used in the mainstream in response to the pandemic conditions. For example, we all became familiar with the use of Zoom and similar applications for many 'authentic' communicative purposes, including for teaching and learning. As a result, innovative uses of video conferencing rapidly emerged for specific educational purposes, such as in the assessment of spoken language. With schools closed, traditional examination providers also had to make use of the available tech to deliver their tests online, outside of the school context. In many cases they also took advantage of AI-systems to support this, e.g. to provide remote proctoring or for the automation of the marking of speaking and writing.

By building on the concept of *hybridity as a design principle,* (rather than an emergency measure), we can exploit the interplay between humans and machines more effectively, taking advantages of machines to do things faster and more efficiently, but retaining the 'human touch' to ensure that educational values remain the top priority - and with the goal of bringing about long-needed improvements in the areas that have been highlighted above.

But before this vision can be put in practice at scale, the concerns recently raised in using ChatGPT have to be addressed. This takes us to the second point: what do we need to learn about AI in order to use it effectively, and what measures do we need to take to prepare for its deployment in our own educational contexts?

4. Understanding AI: what is ChatGPT?

Traditional EdTech and AI models have paved the way for a new era of context-aware, dynamic language education. Educators now have the opportunity to harness the potential of EdTech *plus EdAI* to develop innovative pedagogical practices to support plurilingual learners in an increasingly interconnected, multilingual world. To do so, it is important to provide clearer and more accessible information about the technology itself in language that is designed for non-specialists.

The technical complexity of AI is challenging but an understanding of some basic concepts and related terminology is certainly achievable, and is necessary to inform choices and decisions we need to make about AI-systems. It is also important to grasp the differences between traditional AI and the latest innovations represented by ChatGPT and other generative AI systems. The following four points are particularly important, and their implications need to be clearly understood:

• Collection and use of Data, including concerns for privacy;
• Bias;
• Malpractice (cheating etc.);
• Transparency and Explainability (cf. the "black box" problem).

A good starting point is to find out about two related AI concepts that are central to the development of AI systems that are already being used in educational contexts: *Machine Learning* (ML) and *Natural Language Processing* (NLP).

ML is a subset of AI that uses statistical techniques to enable machines to improve what they do through experience – hence the concept of the machine doing the learning. ML systems are trained using large amounts of data, e.g., written or spoken language. Some significant challenges and risks of AI-based systems from the social perspective are related to the collection and uses of that data.

NLP is concerned with interactions between computers and human language, and how computers are programmed to process and analyse large amounts of natural (i.e., human) language data of all kinds. NLP is central to many applications of ML relevant to language learning and assessment systems, including speech recognition, text-to-speech synthesis, and machine translation.

In building an ML model, *training data* is the foundation for the model's learning and decision-making process. The data provides *input* that "feeds" the model, and the output is the prediction based on the data. The goal is to achieve *a mapping from inputs to outputs*, so that the AI can *accurately predict* the output for new, unseen inputs. The quality and quantity of the training data is therefore crucial; the larger the data set, the more patterns and relationships can be learned from the data. It is important that the training data is diverse, representative and unbiased to avoid biases in the data and in the predictions (see below re: ethical considerations).

The choice of algorithm is also an important consideration. An algorithm is a finite sequence of instructions to perform a computation, and a variety of them are used in AI for specific types of ML. How the algorithms work and how they can be explained to those who use the AI systems are important considerations, especially as some algorithms are inherently more difficult to explain than others. Those that cannot be easily understood and interpreted are known as 'Black Box' algorithms. They are based on artificial neural networks are inspired by the structure and function of the human brain (see also *deep neural networks* - DNNs).

Until recently, ML was largely limited to *predictive models*, used to observe, and classify patterns in content. For example, a traditional machine learning problem was to use texts of written language (such as essays) as input and then to predict the level of proficiency (e.g. the CEFR level). In other words, to address the problem of accurately scoring and classifying the texts according to the features of the writing.

ChatGPT is an example of a *generative model of AI* where GPT stands for *generative pretrained transformer. Generative models* can generate human-like text through seemingly natural interactions with users and can perform many different kinds that can be fine-tuned by users to meet their specific needs.

Generative AI systems are based on Large Language Models (LLMs), and unlike traditional AI systems, LLMs learn from vast amounts of text data in an unsupervised way and can create new, human-like content to fulfil the requests of the user that is prompting it. This generative capability makes these AI systems versatile tools in language teaching and learning. A critical dimension is the way the system is prompted and will be aspect of AI that teachers will have to come to terms with.

5. Ethical concerns - Building Trust in AI systems

As discussed, banning new technologies is not a sensible or practical approach and is out of touch with the realities of technological evolution. There is no inherent reason to prevent the use of ChatGPT in schools or universities, but there are factors that might make it less appropriate or less effective in certain contexts. Cambridge University's pro-vice-chancellor for education, Prof Bhaskar Vira, believes that ChatGPT is simply "a new tool" that is now available (Varsity, Feb 2023). The response is not whether it should be used, but how.

There are many emerging suggestions for AI uses in language education. Here are a few on using it as an 'assistant' to the human teacher in providing language practice with feedback:

- Chatbots can be used as virtual tutors or language partners to help students practice speaking and writing in a foreign language.
- AI can be used to develop adaptive tests that can adjust the difficulty level of questions based on the learner's performance.
- AI speech recognition technology can be used to assess pronunciation and provide real-time feedback to learners.
- AI can be used to analyse student writing, providing feedback on grammar, vocabulary, and coherence.
- AI can be used to analyse students' writing and speech to determine their emotional state and provide feedback on how to improve their communicative skills.
- AI machine translation can be used to help learners understand written and spoken content in a foreign language.

However, there are also pitfalls and risks that need to be acknowledged and understood in order to build trust and protect the best interests of learners.

Known risks are associated with the *collection and use of personal data*, including concerns for the security and privacy of learners. In this respect, bias is a long-standing concern in machine-learning as the AI models depend on the data they are trained on; if the training data is biased, the model's output will also be biased. The danger is that the AI perpetuates underlying inequalities and prejudices through the outputs of the system. One of the main concerns in generative AI is related to the *origin* of the data and the ownership of it in the large language models. And despite the size of the data used to build the models, they are still prone to lapses and factual errors in the output (known as *hallucinations*).

Other concerns related to deployment of AI in assessment include the loss of the 'human touch', an over dependence on technology that offers no improvements in outcomes, and the Black Box problem referred to above. Moreover, ChatGPT has created opportunities for *cheating and malpractice* (plagiarism) in computer-delivered examinations which potentially undermine the integrity of existing assessment systems.

Perspectives should now be sought on these issues from a wide range of stakeholders to ensure that AI systems are designed and used in keeping with the overarching educational goals and values (notes above regarding language education). They should complement the roles that humans play in a hybrid way – rather than

replace them – and a responsible approach should seek to bridge both the technical and social aspects.

Systems that allow for the integration of human intelligence, knowledge, and decision-making with AI system are sometimes referred to as a 'human-in-the-loop' approach. By bringing together AI specialists with domain experts, practitioners and decision-making bodies (policy makers, regulators etc.) it should be possible ensure that AI systems are designed with societal goals in mind (society-in-the-loop) and that the following aspects are attended to so that systems are:

- *Interpretable* by developing understandings of how the technology works.
- *Explainable* by providing accessible explanations so users can understand why/how an outcome was reached.
- *Transparent* by providing information about the data and algorithms used.
- *Justifiable* by providing valid arguments in support of outcomes.
- *Contestable* by providing information enabling stakeholders to challenge an outcome.
- *Sustainable* by developing awareness of the environmental impact of AI and how it can be managed.

Emerging regulations and legislation can also contribute to this debate. Already under the EU's *General Data Protection Regulation* (GDPR), 'data subjects' can request *human review* when automated decision-making is used, and in 2023, the EU 's AI Act is an important step towards regulating AI from a statutory perspective. More widely, governments are engaged with the leading tech companies in establishing Known risks in using AI in language assessment are associated with the *collection and use of personal data*, including concerns for the security and privacy of learners, and the potential for bias in outcomes. Other risks may occur when AI systems are not validated for specific uses, or the outcomes cannot be adequately explained to the users.

Ethical AI is emerging as a field in its own right and ethical frameworks are being developed in various educational contexts (e.g., for assessment) and these are providing useful information for leaders and practitioners. The *Institute for Ethical AI in Education* (2021), for example, promotes a learner-centric approach that seeks to empower educational leaders to make ethically informed decisions about AI on behalf of their learners (see http://www.buckingham.ac.uk/wp-content/uploads/2021/03/The-Institute-for-Ethical-AI-in-Education-The-Ethical-Framework-for-AI-in-Education.pdf) and the World Economic Forum has proposed a very practical 'toolkit' in considering 'Artificial Intelligence for Children' (2022) (see https://www3.weforum.org/docs/WEF_Artificial_Intelligence_for_Children_2022.pdf).

Conclusion

In conclusion, careful thought needs to be given to ways of working that enable a multidisciplinary approach that balances the benefits of AI with its potential risks. This will involve political leaders and policy-makers working collaboratively with educational practitioners and experts from a range of academic disciplines. In language education specifically, a better understanding of the technical and ethical issues will facilitate innovations that value a hybrid approach, bringing together the combined strengths of humans and machines working together to deliver better language teaching and learning outcomes.

See also:
- OECD. AI Principles Overview (2019).https://oecd.ai/en/ai-principles
- Office of Educational Technology (2023). Artificial Intelligence and the Future of Teaching and Learning. https://tech.ed.gov/ai-future-of-teaching-and-learning/
- *World Ethical Data Foundation (2023). https://worldethicaldata.org/*