European Education Area Strategic Framework

Working Group on Digital Education: Learning, Teaching and Assessment & Working Group on Schools - Pathways to School Success

Joint Online meeting on AI Systems in Education

4th May 2023

Input paper: Generative AI in Education

Generative AI in Education

The following input paper was prepared by the Consultants of the DELTA WG in the frame of the online Joint Event on the use of Generative AI in Education, taking place on the 4[th] May.

• Introduction

In late November 2022, the world was introduced to ChatGPT, an Artificial Intelligence (AI)-based chatbot system based on GPT-3.5 that uses natural language processing (NLP) to generate conversations[1]. ChatGPT was the latest in a series of such tools but, unlike previous tools, it has captured the public interest and imagination. It had more than a million users within a week of its launch because of its ability to generate human-like text and its perceived implications and potential use in education, the workplace and everyday life. ChatGPT can answer questions and assist you with tasks like composing emails, essays, and code[2].

GPT-3 (Generative Pre-trained Transformer-3) is a Large Language Model (LLM) that was trained through pattern recognition on a vast quantity of data (499 billion data points – 800GB of data), a model some one hundred times larger than any previous models[3], and ChatGPT enables human-like conversations with this model. Put very simply, a Large Language Model is trained to predict the next word in a sentence, in somewhat the same way that a text message autocomplete feature works, and then to continue this process to generate larger quantities of text.

Users need to be aware that Large Language Models such as ChatGPT are weak on logical and mathematical tasks and have a tendency to 'hallucinate' facts and make reasoning errors (though the use of language, assigning such human like attributes to ChatGPT, rather hides what is really happening in terms of the processing). However, Large Language Models can be combined with standard data search in order to improve accuracy (and this has been done by Microsoft in incorporating GPT-4 with Bing[4], and by Google in developing Bard[5]) and can also be combined with quite different kinds of AI systems such as Wolfram Alpha[6] (which is based on computational approaches, rather than on pattern recognition) in order to better handle mathematical tasks.

ChatGPT is just one example of Large Language Models, which in turn are examples of 'Generative AI', which is a category of artificial intelligence (AI) algorithms that generate new outputs based on the data they have been trained on[7].

• What is Generative AI?

Generative AI systems create new content in the form of images, text, audio, video, and more, for example:

- **Images**: Generative AI can create new images based on existing ones, such as creating a new portrait based on a person's face or a new landscape based on existing scenery

---

[1] https://www.forbes.com/sites/bernardmarr/2022/12/21/chatgpt-everything-you-really-need-to-know-in-simple-terms/?sh=16674aadcbca

[2] https://www.zdnet.com/article/what-is-chatgpt-and-why-does-it-matter-heres-everything-you-need-to-know/

[3] https://hai.stanford.edu/news/how-large-language-models-will-transform-science-society-and-ai

[4] https://www.bbc.com/news/business-64562672

[5] https://www.bbc.com/news/technology-64546299

[6] https://writings.stephenwolfram.com/2023/03/chatgpt-gets-its-wolfram-superpowers/

[7] https://www.weforum.org/agenda/2023/02/generative-ai-explain-algorithms-work/

- **Text**: Generative AI can be used to write news articles, poetry, and even scripts. It can also be used to translate text from one language to another
- **Audio**: Generative AI can generate new music tracks, sound effects, and even voice acting[8]

These tools have captured the public imagination and people, in all walks of life, are exploring their possibilities. One of the reasons for the popularity of ChatGPT is that you can type natural-language prompts, such as "create an essay" and the system can create conversational, if somewhat stilted, responses[9].

Great claims have been made for the potential of AI in general and of Generative AI in particular – Bill Gates summarises these in his blog, *The Age of AI has begun*[10].

Sometimes these systems are able to perform tasks they were not explicitly trained on, for example ChatGPT is able to translate between some languages even though it was not trained on translation data. However, this may only be the tip of the iceberg, and some believe that these Large Language Models have much greater capabilities and these will emerge over time as the models grow in size and sophistication. The surprising abilities of ChatGPT have led to conflicting arguments about the implications of its abilities to generate very human sounding output, which have even (mis)led some to think of it as alive[11].

The linguist Emily Bender famously described ChatGPT as a stochastic parrot[12], and Chomsky similarly argues: "The human mind is not, like ChatGPT and its ilk, a lumbering statistical engine for pattern matching, gorging on hundreds of terabytes of data and extrapolating the most likely conversational response or most probable answer to a scientific question. On the contrary, the human mind is a surprisingly efficient and even elegant system that operates with small amounts of information; it seeks not to infer brute correlations among data points but to create explanations."[13]

An example of experimental work supporting these perspectives is the work of Trotte et al.[14] who provide empirical evidence from comparing the responses to the False Belief task to human participants to and GPT-3and that that "while statistical learning from language exposure may in part explain how humans develop Theory of Mind, other mechanisms are also responsible."

Some authors (but not many) have gone the other way, suggesting that the success of ChatGPT may actually tell us something about human language, indicating that it is not as complex as we think it is and that humans may actually produce language in ways not too dissimilar from large language models, and that that we may need to move to a position in which " the distinction between true understanding and mere pattern recognition gives way to something new".[15]

Another perspective positions humans and AI systems as working in dialogue, arguing that "we are intelligent primarily in networks and systems rather than as lone individuals" and that "the best model

[8] https://www.weforum.org/agenda/2023/02/generative-ai-explain-algorithms-work/

[9] https://www.cnet.com/tech/computing/why-were-all-obsessed-with-the-mind-blowing-chatgpt-ai-chatbot/

[10] https://www.gatesnotes.com/The-Age-of-AI-Has-Begun

[11] https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/

[12] https://nymag.com/intelligencer/article/ai-artificial-intelligence-chatbots-emily-m-bender.html

[13] https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html

[14] https://arxiv.org/abs/2209.01515

[15] https://www.noemamag.com/gpts-very-inhuman-mind/

for interacting with AI may well be one of human/machine pairing – where our own intelligence is augmented, not replaced by AI.[16]

• Some examples to try

Generative AI tools are becoming more and more popular and some such as ChatGPT, are becoming so popular that it is impossible to access the free service at certain times of the day. However, there is nothing like trying out these tools to see what they can create in terms of text, images, music and much more. A good place to start is to use [ChatGPT](16)[17], [Perplexity](16)[18] or Bing Chat or Google's Bard (as of time of writing you must join a waiting list for both of these, and Bard is only available in UK and US) and start to engage in a conversation. Pick a topic you are interested in and ask the chatbot a direct question such as "*Tell me about*" or "*Create a lesson plan or a quiz about*". You might also ask the chatbot to proofread a paragraph of text and ask it to show you what changes it made. The better your prompts the better the responses will be, and you should strive to provide as much context as possible and use specific and detailed language. For example, you might include information about:

- your desired focus, format, style, intended audience and text length.
- a list of points you want addressed
- what perspective you want the text written from, if applicable
- and specific requirements, such as no jargon[19].

Moving away from text to images you can use prompts to create images generated from text descriptions, such as "the face of a shiba inu[20] blended into the side of a loaf of bread on a kitchen bench, digital art"[21]. Some tools you can try are: DALL.E 2[22], Stable Diffusion[23], Google's Imagen[24] and Bing Image Creator[25] that can generate images from text. You might also like to try out Google's Quick, Draw![26], which asks you to draw a doodle and it then attempts to recognise it.

• How can Generative AI be used in education?

Generative AI tools have the potential to transform how we work, how we teach, how we learn and how we assess. In the world of work these tools are now being embedded into existing applications such as email, word processing, spreadsheets, presentation tools etc. to increase productivity and support us in generating and analysing information. We are already using AI tools as part of our word processing software to enhance our expression and flow, and there is an expectation that these features will become more sophisticated and powerful over time[27].

---

[16] https://theconversation.com/has-gpt-4-really-passed-the-startling-threshold-of-human-level-artificial-intelligence-well-it-depends-202856

[17] https://openai.com/blog/chatgpt

[18] https://www.perplexity.ai/

[19] https://theconversation.com/how-to-perfect-your-prompt-writing-for-chatgpt-midjourney-and-other-ai-generators-198776

[20] **Shibas** are currently one of the most popular companion dogs in Japan.

[21] https://theconversation.com/ai-art-is-everywhere-right-now-even-experts-dont-know-what-it-will-mean-189800

[22] https://openai.com/product/dall-e-2

[23] https://stablediffusionweb.com/#demo

[24] https://imagen.research.google/

[25] https://www.bing.com/create

[26] https://quickdraw.withgoogle.com/

[27] https://www.nytimes.com/2023/02/01/technology/personaltech/chatgpt-ai-bots-editing.html

In education these tools have the potential to help teachers manage their administrative work, to generate lesson plans, to provide more personalised responses to student inquiries or provide students with instant feedback and support.

These tools can eliminate '*writers block*' for both the teacher and the student. It can assist teachers and students in the following ways:

- Drafting and brainstorming for lesson plans and other activities
- Help with the design of quiz questions or other exercises
- Providing grammatical or structural feedback on portions of writing
- Customising materials for different preferences (simplifying language, adjusting to different reading levels, creating tailored activities for different interests)[28]

These tools can very quickly "mindlessly generate text"[29], which has led to concerns over plagiarism. This has led to some jurisdictions and institutions seeking to ban these tools, though other are seeing this as an opportunity to rethink how they teach and assess student learning[30]. With the upcoming incorporation of tools based on these technologies by Microsoft and Google into their search engines, and office applications the prospect of effective bans on their use appear likely to be futile. Perhaps a more effective way forward is to see these new tools as partners and helpers rather than replacements for either teachers or students work. For example, these tools can quickly synthesise large datasets and identify patterns and trends that could inform teaching strategies and curriculum development.

The good news is that there are many groups of teachers getting together online to develop approaches to using these new technologies – join one!

• The Risks associated with Generative AI

The Technical Report[31] which accompanied the release of GPT-4 (the latest version of GPT) in March 2023, provides some insights into potential risks. The first indicator is that unlike in reports for earlier versions of GPT they provide no data on the training set or training methods for this version – citing both competition and safety concerns. Some specific indicators in the report are:

- GPT-4 was able to do things the designers did not expect and did not predict – an example given relates to 'Hindsight neglect' (p.4)
- When given unsafe inputs, the model may generate undesirable content, such as giving advice on committing crimes (p. 12)
- Indications of some proficiency in generating text that favours autocratic regimes when prompted to do so in multiple languages (p.51)
- The emergence of concerning novel capabilities "including the ability to create and act on long-term plans, to accrue power and resources ('power-seeking'), and to exhibit behaviour that is increasingly 'agentic'." (p. 54) One of the studies referenced in the Technical Report highlights "the increasingly evident fact that ML systems are not fully under human control"[32]

Such risks have led to an open letter entitled "Pause Giant AI Experiments" signed by a large number of AI researchers and backers (including Elon Musk) calling for a pause in the development of these

---

[28] https://nationalcentreforai.jiscinvolve.org/wp/2023/03/14/getting-started-with-chatgpt/

[29] https://nymag.com/intelligencer/article/ai-artificial-intelligence-chatbots-emily-m-bender.html

[30] https://www.nytimes.com/2023/01/16/technology/chatgpt-artificial-intelligence-universities.html

[31] https://cdn.openai.com/papers/gpt-4.pdf

[32] https://arxiv.org/abs/2302.10329

models. The call talks of "AI labs locked in an out-of-control race to develop and deploy ever more powerful digital minds that no one – not even their creators – can understand, predict, or reliably control" and argues that there are "profound risks to society and humanity."[33]

A more wide-ranging discussion of potential risks of Large Language Models was set out in paper by Bender and her co-authors in 2021[34], and they responded to this open letter writing: "It is indeed time to act but the focus of our concern should not be imaginary "powerful digital minds." Instead, we should focus on the very real and very present exploitative practices of the companies claiming to build them, who are rapidly centralizing power and increasing social inequities."[35]

On the other side of this argument, see Donald Clark's response entitled 'Moral panic and AI regulation'[36] opposing bans on AI development, arguing that we already have sufficient regulation in place.

• What are the dilemmas Generative AI presents for teachers?

Generative AI will have an impact on all the media inputs around you, on education and on your students' everyday lives. The boundaries between human created and machine created media will be further blurred and this will necessarily impact on forms of teaching and assessment and provide further challenges for the development of digital competencies.

Whilst it is easy to get Generative AI systems to produce some output, the task of producing good prompts is not trivial and there is much to learn about how this is best done, though there are a number of resources available to support teachers in doing this[37].

In using Generative AI in ways suggested in the previous section, the limitations of the tools need to be kept in mind, with the user (whether teacher or student) constantly reviewing what is generated to ensure it is appropriate for the particular context. Large Language Models ultimately derive their responses from the data they were trained on and so can reflect all the errors that may be found there, and they do not have the ability to, for example, consult the internet to update their information (though increasingly they will be combined with other systems to enable this to be done – as we see in the new Bing search and in Bard for Google).

Reports giving empirical evidence of the impact on using Large Language Models are starting to be published, for example a paper by Noy and Zhang[38] provides evidence of the impact of using ChatGPT on occupation specific writing tasks of college educated professionals, demonstrating both productivity and quality gains.

There is a growing trend to embed Generative AI into existing educational products. Examples include Khan Academy and Duolingo which are actively working on incorporating these technologies into their products. Duolingo is using them to provide additional grammatical explanations, and to generate role play conversations. Khan Academy is exploring how ChatGPT-4 can act as a virtual tutor and help learners when they get stuck or have a question[39]. Such technologies will be likely to spread to many

---

[33] https://futureoflife.org/open-letter/pause-giant-ai-experiments/

[34] https://dl.acm.org/doi/10.1145/3442188.3445922

[35] https://www.dair-institute.org/blog/letter-statement-March2023

[36] http://donaldclarkplanb.blogspot.com/2023/03/moral-panic-and-ai-egulation-primer.html

[37] For example: https://drive.google.com/file/d/1q9exc7gm3DpRAygeV8hgZ-7sVnyrTq6b/view

[38] https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf

[39] https://www.youtube.com/watch?v=rnIgnS8Susg

other educational products and this will throw up challenges for teachers in integrating these into their teaching. The promises may however not be achieved easily, a recent small empirical study[40] examining the relative effectiveness of ChatGPT and human tutor generated algebra hints, found that the human created hints resulted in substantial and statistically significant better results than ChatGPT hints.

Hopes have been expressed that ChatGPT will be able to help with marking. In a number of small-scale studies[41] [42] Christodoulou has explored the value of ChatGPT-4 in marking children's essays with rather poor results – finding poor agreement between human judging and GPT marks; poor predictive validity for GPT marks; and clustering of marks given by GPT. Another study she carried out on using ChatGPT to provide written feedback[43] on essays found that ChatGPT provided feedback that was as good as that provided by teachers –though she was critical of the value of the feedback by both teachers and ChatGPT.

The success that ChatGPT-4 has had on taking public examinations has raised some questions about the value of these examinations and of the curriculum they are assessing. Some caution is needed here as the great success of ChatGPT on the Bar Exam, for example, has been questioned[44], raising the possibility that it was being tested on data it had been trained on.

The linking of Wolfram Alpha to ChatGPT leading to success in a UK Maths-A level exam has led to the claim that high school maths may need to be completely revamped[45] though Christodoulou argues that to some extent this is a misreading of what assessment is for[46].

The use of AI tools in computer programming[47] is raising questions about the nature of computer science education. Coupled with this is the observation that the use of systems like ChatGPT relies not on traditional programming but on scripting 'prompts' (and libraries of such prompts are now available online[48]). It just might be that at school level there should be a reconsideration of the role of coding and computational thinking in digital competences and increased emphasis on the rather different skill of writing prompts for Generative AI.

- ## Questions
The following questions are intended to stimulate further discussion and reflection on the implications of Generative AI for education. They are not meant to be exhaustive or definitive, but rather to invite different perspectives and opinions from various stakeholders.

1. What are the opportunities and challenges of using Generative AI into education, and how can we balance innovation and regulation? How can we ensure that Generative AI is used in a legal

---

[40] https://arxiv.org/abs/2302.06871

[41] https://substack.nomoremarking.com/p/can-gpt-3-mark-writing-the-data-is

[42] https://substack.nomoremarking.com/p/more-gpt-marking-data-is-it-better

[43] https://substack.nomoremarking.com/p/can-chatgpt-provide-feedback-our

[44] GPT-4 and professional benchmarks: the wrong answer to the wrong question (substack.com)

[45] https://www.conradwolfram.com/writings/game-over-for-maths-a-level

[46]   https://substack.nomoremarking.com/p/if-we-are-setting-assessments-that-a-robot-can-complete-what-does-that-say-about-our-assessments-cbc1871f502

[47] https://www.theregister.com/2022/10/20/ai_programming_tools_mean_rethinking/

[48] https://promptbase.com/

and ethical way in education across Member States, and that we are aware of the potential risks and benefits of this technology.

2. How can we prepare educators and learners for the changing educational landscape that Generative AI will create? Specifically, what are the skills and competencies that they will need to thrive in a world where Generative AI is ubiquitous and influential?

3. How have different countries and regions responded to the emergence of Generative AI in education?