

In September 2023, Learnus council member Michael Thomas attended a meeting of the *All Parliamentary Party Group (APPG) on Artificial Intelligence* at the House of Lords, convened to discuss the potential impact (for better or worse) of generative artificial intelligence on education.

**What is generative AI?** It's a computer programme capable of generating text, images, or other media which are very similar to those generated by people: you may have heard of 'ChatGPT' for example. Generative AI involves machines learning patterns from massive exposure to content on the internet. Generative pre-trained transformers (that's the GPT bit) are a type of large language model that has learned how to generate extended text on a given topic simply by exposure to sequences of words. The version tuned to talk with you is called Chat-GPT, the transformer bit concerning transforming a question or request into an answer (others like Bing and Bard do similar things). Generative AI can also learn to link text to images. For example, I asked one programme, DALL·E, to "draw a photorealistic picture of a university administrator thinking very hard about artificial intelligence". The output wasn't perfect, I had to edit the text in the bubble (initially gibberish), and the programme still struggles with hands, but it's not bad, is it?

**Generative AI stands to be immensely disruptive for education.** Why? Because it can automatically write answers to essay questions. This means, potentially, that written work generated outside of invigilation (e.g., homework), can never again be reliably used to assess student knowledge and learning. ChatGPT can generate a plausible answer that may be hard to detect as computer-generated. Does this innovation mark a disaster for education? Is there a way that generative AI can be used productively by students?

Here are four themes that emerged from the recent evidence meeting of the APPG.

**1. No one was panicking that AI robots were going to take over the world.** Everyone recognised the drawbacks of AI using information 'scraped' from the internet, and the implications of AI for employment (variously: the risk of data scraped from the internet reflecting biases of that content, such as inaccuracy, stereotyping, and discriminative views; of data privacy, particularly for minors because interactions with chatbots are owned by the provider; of how to ensure age-appropriate content; of data ownership because large language / image models exploit original content, especially from the arts; the risk that chatbots can be purposed to polarise users' beliefs; the risk of reliance on foreign big-tech companies to provide key educational tools; the concern that people will be made redundant by jobs that AI can now do...). So yes, some downside risk, but no one was panicking!

**2. There was a diverse range of views** expressed on what tools like ChatGPT mean for education – all the way from 'that don't impress me' to 'it's a steppingstone to utopia'. Some at the meeting viewed generative AI for education as on a par with the introduction of calculators to maths class, or of search engines for researching essays and projects: a helpful tool, necessitating some tweaking of teaching practice, but not much more. Others were prepared to embrace the fact that we can never again use independent written work to assess students' learning. That the likes of ChatGPT should mark the end of an education system based on cramming students full of knowledge and asking them to regurgitate it

using pencil and paper in hot exam halls in July. To shift instead to an education system based on critical thinking, systems thinking, analytic thinking, group-based problem solving – the skills young people require for the 21<sup>st</sup> century. Homework should now be about *research*, using all the tools available, so that the classroom is for group-based application of knowledge. Generative AI should be available to all, it can close the skills gap, democratize education, lead to a more equitable society, it can lead to ... utopia. One suspects that the reality falls somewhere in between these views!

**3. The kids currently know much more than the teachers** – pretty much everyone agreed that the most important first step is to improve teacher literacy on generative AI, to understand what these systems can (and can't) do, and to begin to think about how they may be used. Right now, students know more about ChatGPT than teachers (there was, indeed, a drop off in ChatGPT use in July... because the students went on their summer holidays). While there may be a knee-jerk response to ban use of generative AI, this would be like banning internet search engines like Google. Search engines have empowered students across the world to access a vast reservoir of world knowledge – even while we know that what is pulled up by a search has to be judiciously handled and verified.

Perhaps the most important take-home for teachers and students alike is that **you've got to know the technology's limitations**. These AI systems are not really intelligent (watch out for marketing from the Big Tech companies, they want you to think it is! You'll pay more!). Generative AI doesn't have any knowledge of the world, just what plausible-looking text looks like learned from massive exposure to text and the images that tend to go with text. They are souped up versions of the predictive texting facility on your phone. The AI generates plausible or 'high probability' text but not necessarily factually accurate text. That is, ChatGPT can *make stuff up*. This is tricky for teachers if they are using generative AI as a research tool, and risky for students if they are using it to write their essays. But just because a tool is flawed doesn't mean it can't be useful.

**4. Guidance is beginning to emerge** - institutions are thinking hard about the educational impact of generative AI, and some guidance is beginning to emerge (e.g., from the Department for Education and from the Russell Group of universities). I gave a lecture last week to advise postgraduate students on how they might use ChatGPT as a tool in their essay writing, what the chances are of getting caught if they simply use it to write their assessments (and the very mediocre mark they would likely receive even if they didn't get caught). There are many ways generative AI can be useful in education: to suggest initial ideas, to give feedback on text, to help second language learners improve their writing, for checking and recommending Excel formulae. There are inevitably pitfalls we need to avoid (mostly linked to ensuring that content is factually true, and that creativity is not stifled – ChatGPT will encourage you to write like everyone else on the internet!). One can hope that the involvement of commercial interests doesn't lead us as societies to mess up the opportunities of the technology (like we've done with social media and its impact on politics and on young people's mental health). One can hope that the new technology does not impoverish us, reducing the number of tasks we have to do so that we are less cognitively stimulated, but enriches us, enabling us to do new more complex or more rewarding and satisfying activities.

My forecast is that as with many new technologies, generative AI is likely to be gradually adopted into educational practices, as we increasingly recognise the opportunities it provides and the risks it carries. We will gain a better understanding of best practice and the new goals the technology will enable us to pursue. But such a future requires that we as educators engage and understand more about what the technology can and cannot do. Start playing with generative AI.