

AI's Recent History

The Skinny on AI for Education, Late March 2024

Headlines

- [Who is in charge of educational AI?](#)
- [Cometh the hour, cometh the man... but which one?](#)
- [Celebrating being human](#)
- [Other articles you might find interesting](#)
- [Further reading: find out more from these resources](#)

Before I get into the detail, I thought you might like to know that 'The Skinny' is not written by an AI. AI certainly helps with some grammar, proof reading and spelling checks. But when it comes to actually writing 'The Skinny' for better or for worse, you, the reader, get my personal synthesis and analysis of the various AI developments and opinions that I have been reading about over the past few weeks. This comes with my bias and my understanding about AI and education built up over 30 years or so. I am sure that an AI could collate a selection of articles and produce a newsletter that summarised these articles and offered its perspective. But this wouldn't be the same and it wouldn't be 'The Skinny'.

One thing that has struck me about my recent reading is a lack of articles about the practical application of artificial intelligence, and the advantages and disadvantages that it's bringing to different communities and disciplines, particularly those that are relevant to education and training. Where is the evidence about what is happening with AI in education and training? Is it really reducing workload, increasing personalised learning, helping with professional development or delivering any of its other myriad potential benefits?

Who is in charge of educational AI?

Every day I hear about new Educational AI applications, and it is impossible to keep up. I am always keen therefore to find articles that try to keep track of the different developments in educational AI, such as [this one that looks at Generative AI products and the purposes that they can serve](#).

But Educational AI does not exist in a vacuum and what happens in the wider AI context has a direct impact on the AI that ends up in schools, colleges and universities. Education needs a much louder voice in conversations about AI and part of making that happen comes from understanding who has the power at the moment.

I worry about the lack of diversity at the heart of what is being developed. If you look under the hood of the educational AI products that are available, particularly the generative AI products, you see that most of them are connected to a small number of large technology companies. Sometimes this connection is through:

- The company producing the educational AI product using one of the Generative AI models that the big tech companies produce, such as OpenAI's ChatGPT;
- Sometimes the connection exists because the company selling the educational AI product (or the technology that it uses) has benefitted from investment from a big tech company, which usually means the big tech company has some ownership rights (Microsoft's investment in OpenAI, for example)
- On other occasions, the company producing the AI that powers the educational product has been founded by someone who previously worked for one of the big tech companies. For example, Anthropic, who produce the Large Language Model called [Claude](#) was founded by Daniela and Dario Amodei, who previously held senior positions at OpenAI. It is also the case that Anthropic has received investment from both Amazon and Google

An interesting aspect of what is happening with investment by the MANGs (Microsoft, Amazon, Nvidia, Google) is that their relationship with the AI company in which they are investing is not purely financial. The MANG often provides more than money. In the case of Nvidia, they provide scarce processing

power, and in the case of Microsoft their azure cloud computing platform. In return, [the MANG may acquire technical expertise and customers as well as some ownership rights](#). Of course there is also a competitive edge at stake. Obviously, the likes of Amazon and Google don't want Microsoft and OpenAI to 'rule the roost.' They therefore invest in companies like Hugging Face, who provide open source Generative AI models as an alternative to OpenAI. Other Hugging Face investors include Nvidia, Salesforce and Intel. And let's not forget Apple, who are quietly [increasing their AI capabilities with an eye on what it can do for the iPhone](#).

One thing is clear, when it comes to AI, and in particular Generative AI, most if not all roads lead back to the so-called MANGs, and I am not sure this is a healthy situation for learners, teachers or parents. In fact, I am quite sure that it is not. As more educators start using AI and learning about it, I hope that we can give education a louder voice in what AI is and is not developed.

Who, or what, might stop the march of the MANGs?

Regulators

The Federal Trade Commission (FTC), which regulates competition in the US is investigating these partnerships between big tech companies and is [demanding more information about the deals](#).

The hunger for chips

The 'chips are always down' for AI, because it requires huge processing power, and that power depends upon a plentiful supply of 'chips.' And there are challenges to the plentiful supply of chips. A computer chip is a tiny electronic device that is responsible for processing and storing data, performing calculations, and managing the flow of information within a device, such as a phone or computer. AI requires sophisticated chips with a lot of computational power to function effectively, and specialised AI chips are designed to handle these demanding tasks. These chips can process vast amounts of data quickly and efficiently, making it possible to train and run complex AI models.

Many of these chips are produced by Nvidia, a US company, whose profits are soaring on the back of the AI gold rush. But building and expanding chip manufacturing facilities is a capital-intensive and time-consuming process, and some of the raw materials and equipment can be in short supply. Precious metals, like gold, silver, and palladium are required for some advanced chips particularly those used in high-performance computing and AI applications – these can be in short supply and expensive. It may therefore not always be possible to satiate AI's ferocious hunger for chips.

Money, money, money

The business of building advanced generative AI is extremely expensive, and there are concerns about the extent to which even huge companies like Microsoft can actually [resource their AI aspirations](#).

Water, water everywhere and not a drop to drink

As Samuel Taylor Coleridge's Ancient Mariner lamented, water is around us but is often not clean or safe enough to drink, and AI is super thirsty! The processing power required to make the AI work uses huge amounts of energy, and [water is needed for cooling the heat generated by this process](#).

Human labour shortages

We usually think about the machines that power our AI, but actually [these intelligent machines require a lot of humans too](#). Humans play a crucial role in annotating and labelling the data used for training AI algorithms, which can include tasks like identifying objects in images, transcribing speech, or classifying text sentiment. Generative AI models require vast amounts of human-created data to learn from. For example, language models are trained on large corpora of human-written text, such as books, articles, and websites.

Similarly, image generation models learn from huge datasets of human-created images. Human expertise is also often required to ensure that the AI system is learning meaningful patterns and avoiding biases or errors. Even after an AI model is trained, humans are needed to monitor its performance, identify errors or biases, and make necessary corrections. Generative AI models often rely on human-written prompts or instructions to guide their output. Crafting effective prompts that elicit

the desired responses from the model requires human creativity, domain knowledge, and an understanding of the model's capabilities and limitations.

Humans are also essential for quality assessment and filtering. The outputs of AI models can sometimes be inconsistent, irrelevant, or even inappropriate. Human judgement is needed to assess the quality of generated content, filter out unwanted outputs, and provide feedback to improve the model's performance.

And crucially, humans are vital for ethical considerations: AI, including generative AI models can potentially produce harmful, biased, or misleading content. Human oversight is crucial to ensure these models are used responsibly, and their outputs align with ethical guidelines and societal values. This includes monitoring for biases, filtering out inappropriate content, and establishing guidelines for the use of generative AI. Humans play a vital role in defining ethical guidelines, monitoring AI systems for potential biases, and making decisions about how and where AI should be applied.

And this is not easy. As the [recent challenges faced by Google's chatbot Gemini](#) demonstrate. Gemini has been generating biased, inaccurate and politically incorrect responses, such as including a Black man among the US Founding Fathers and depicting German soldiers from World War II with a Black man and an Asian woman. Gemini's text-based responses also faced criticism for being overly politically correct. For example, it stated that there was "no right or wrong answer" when comparing Elon Musk's meme-posting on X to Hitler's actions.

Google apologised and paused the tool, but there is no easy fix. Much of the problem lies in the biased data used to train AI tools, like Gemini. Much of this data comes from the internet, which contains various biases and inaccuracies. Google, along with other Generative AI companies actively try to offset these biases, but the nuances of human history and culture are things that machines struggle to understand, so the challenge remains considerable and will certainly need a lot of humans for any progress to be made.

Labour relations

There is no question that AI relies heavily on human involvement. And humans do not work 24/7 without needing a break – they are not machines. It is interesting to note that investment in AI development is 50 times higher in the US than in Europe and one of the reasons for this is that labour relations are much tighter in Europe and this impacts on the speed at which companies can grow and restructure in the way that organisations like OpenAI have done. [Social protection for workers is much stronger in Europe than in the US.](#)

However, the US is not immune from labour relations challenges. The Hollywood writers' industrial action demonstrated that people power can be wielded. Hollywood screenwriters represented by the Writers Guild of America (WGA), [brought Hollywood to a standstill](#) last year over their concerns about AI potentially replacing them.

It will be interesting to see how the release of OpenAI's latest Generative AI text-to-video generator: Sora, impacts on Hollywood. Currently, Sora's capabilities are limited. It can only generate short video clips and lacks a human understanding of physics, for example. However, the AI is learning and improving, raising concerns about the source of its training data and the potential for copyright infringement. The entertainment industry and media groups are trying to prevent AI companies from using [copyrighted material to enhance their models](#) without permission.

Will AI be just another step in the evolution of filmmaking technology, or will it precipitate an increased flow of creative work, or will it stymie the quality and character of film?

Cometh the hour, cometh the man... but which one?

The concerns about the small number of companies behind the advance of AI are considerable, but it is also worth looking at some of the people who drive the conversation. As many headlines have reported [Elon Musk is suing OpenAI](#) for abandoning its original non-profit mission for a for profit one. One could just see this as two giant egos (and let's be fair talent titans too): Sam Altman and Elon

Musk, locking horns, but it is much more significant than that. As John Thornhill points out it's about [the future of AI transparency](#).

OpenAI started as a non-profit entity that benefitted from funding provided by Elon Musk. But to compete with the likes of Google with its massive computing power and technical talent, OpenAI was sucked into a relationship with tech giant Microsoft in 2023. As a result the not-for-profit OpenAI entity with its charitable status evolved to a for-profit organisation – and let's face it that does sound 'fishy.'

We will have to wait and see what the outcome of the lawsuit brings, but in the meantime, it's interesting to look at the two men at the heart of this: Elon Musk and Sam Altman. Both have founded and are involved with multiple companies.

Sam, for example, has a cryptocurrency company called 'Worldcoin' that offers people cryptocurrency tokens in exchange for their consent to have their personal data collected using an eyeball-scanning "orb." But there is [little information about what happens to that personal data](#).

Elon has an even more invasive AI company in Neuralink that he hypes up in style, suggesting that it has [implanted an electrode into a human brain](#) for the first time. The point being that eventually this would enable people's brains to be directly connected to computers leading to enhanced processing powers. Other companies have also successfully inserted electrodes into people's brains and used them to collect and interpret human brain signals, so Elon's claim that he was first is not exactly true. It is also true that much work on computer brain interfaces has the aim of helping patients with paralysis, which is slightly different to Musk's mind enhancement imperative.

Altman has also come up with some incredible statements about AI. One I was particularly struck by was in response to the New York Times lawsuit filed in January claiming that OpenAI had "intentionally manipulated" ChatGPT, so that it [reproduced entire articles and sections of articles](#) that had been previously published in the newspaper.

But, according to Altman, when reproducing the articles from the NYT, ChatGPT had suffered from "inadvertent memorisation." I must confess, I'm not too clear what 'inadvertent memorisation is'. I remember many years ago when I was revising for examinations and had been working hard at remembering vast amounts of information, I would sometimes catch myself memorising random things that I happened to see, such as a billboard advert or a slogan on a product I was eating. Perhaps this was 'inadvertent memorisation'? But isn't one of the great advantages of LLMs such as ChatGPT that they can help us by doing some of the memorisation for us? This now leaves us with a dilemma: who decides which is inadvertent and which is advertent memorisation – is it the Chatbot or the user?

Sam also called on congress to create a new agency to [license AI capabilities](#). Which would of course be anti-competitive and favour companies just like OpenAI. It is not the answer to the problems we are facing with AI regulation as [this nice article from the Brookings Institute explains](#), although it could contain some helpful ways forward.

It's also important to remember that Musk's original motives for investing in OpenAI were perhaps more of a 'Jilted John' response than a 'for the good of mankind' one. Elon had a spat about the future of AI with Larry Page from Google in 2013 and when Google acquired Deepmind, Musk tried to block the deal. It was after this that he started to invest in OpenAI.

All in all, a world of AI driven by Sam and Elon, and others like them, is not something that I relish and I don't think it will be good for teachers and learners.

Celebrating being human

The importance of empathy

But let's not be too downbeat. Let's celebrate how fabulous humans are.

Back in the 1990's BT ran a series of advertisements using actors such as Bob Hoskins to encourage people to give a friend a call because "it's good to talk." Obviously, this was to increase the profits that BT made from phone calls, but putting the commercial imperative aside, the importance of social interaction at the heart of the advertisement's message was well made. At a time when there is [an epidemic of loneliness](#) within human society, which our social media technology has not helped, it's important to focus more on real human interaction, communication and connection. I worry about the stream of AI applications that focus on creating a 121 relationship, from chatbot tutors to AI that promises friendship – even a [relationship](#). I appreciate that not everyone has a friend they can call up for a chat or some advice about an assignment. And, AI certainly has a role to play, but let's not hand over too much of the social interaction to AI. It is not a human substitute.

Education is all about relationship building. As educators, we know that the relationships we build with students are crucial to the success of our teaching. And our schools and colleges are one of the most important places where young people learn about building successful human relationships, making friends, learning the written and unwritten rules of being a member of a community. And about important human capabilities, such as empathy.

One only has to observe how couples can build an empathetic relationship even when there is no shared spoken language between them. They may learn to speak the language of the other over time, but initially their relationship requires a very complex process of physical communication that is born out of sophisticated, cognitive processing and extreme self- and other awareness. This is something that AI systems simply cannot engage with, and we underestimate the importance of such empathetic communication at our peril.

The wonder of babies' intelligence

A [recent study by researchers at New York University](#) has shown that AI models can learn from the sort of data that babies process in order to learn. The researchers used video and audio data from a baby named Sam to train a neural network AI model, which exhibited a remarkable degree of learning. What I found particularly interesting about this article was the way that the authors stressed the impressive learning abilities of babies. These abilities enable them to respond to the signals in what they see and to develop their own learning hypotheses. Alison Gopnik, a psychology professor at University of California is quick to point out that babies possess unique and impressive learning abilities that current AI systems lack, such as [imaginative model building, active exploration, and social learning](#). So let's celebrate babies and their phenomenal learning abilities!

Other articles you might find interesting

Andrew Ng's newsletter is always worth reading:

<https://www.deeplearning.ai/the-batch/tag/letters/>

The Brookings Institute's Artificial Intelligence and Emerging Technology Initiative

<https://www.brookings.edu/projects/artificial-intelligence-and-emerging-technology-initiative/>

Ethan Mollick's substack:

<https://substack.com/@oneusefulthing>

The Bourne-Epsom Protocol

<https://www.ai-in-education.co.uk>