The ‘No BS’ Guide to AI

Co-written by Professor Rose Luckin, Director of EDUCATE at UCL and Priya Lakhani, OBE, Founder-CEO of CENTURY.Tech

What is AI?

Artificial intelligence refers to technologies which are capable of performing tasks as well as, if not better than, humans.

It's useful to distinguish between what people call Strong AI and its opposite, Weak AI. Strong AI (also full AI or generalised AI) is what we see in sci-fi depictions: sentient robots able to perform any task much the way that humans can.

On the other hand, Weak AI (also narrow AI) is an AI system that can do usually just one task at least as well as a human can. Although movies and media might lead us to believe the robot invasion is just around the corner, in reality, we've only got as far as developing Weak AI systems.

Some examples of the Weak AI systems that we have are:

IBM's Watson: Beat a human player at a game of Jeopardy.

Deep Blue: Beat the current chess champion.

AlphaGo: Beat the world champion at the game of Go.

Siri/Alexa: Voice-activated personal assistants that can "understand" natural language and reply in kind.

Amazon/Netflix/Spotify/Pandora: Recommendation engines that can predict what we will want to buy/watch/listen to based on what we have done.

Tesla/self-driving cars: Cars use AI to learn how to safely navigate and drive along roads.

So, is AI the same as Machine Learning?

Machine Learning (ML) is a way of achieving artificial intelligence. Rather than a software developer writing the specific rules and logic steps for a program, which can be hugely intricate and time-consuming, instead the program is "trained" using large quantities of data and basic algorithms are written to allow it to "learn" from this training to perform a specific task.

This means that **artificially intelligent programs are far better able to cope with unusual cases, extreme examples or incomplete data.** Rather than a programmer thinking of every single possible scenario, the program teaches itself and then draws conclusions. We are using these technologies to understand how people learn and provide them with a personalised education.

I've heard people mention neural networks, how do they fit into all this?

Neural networks are a special form of machine learning that that are inspired by the design of the neurons in the human brain. Neural networks are small networks of these 'artificial neurons'. Similar in processing to real neurons, they can form connections and use those connections to detect patterns. This is widely used for image detection and natural language processing.

Deep learning is usually mentioned together with neural networks; it is a special variant of neural networks.

You can think of artificial intelligence, machine learning and deep learning as nested within one another: deep learning is a type of machine learning; machine learning is a type of artificial intelligence.

What differentiates AI from other technology?

AI and ML generate an extremely fast feedback loop between result, feedback and learning. With conventional algorithms, results are produced and feedback collected from the user, but the learning is done by a human who will adjust the algorithm. That can take hours, days or months. AI and ML remove the human element and can, therefore, learn from user feedback within milliseconds.

They say you need a lot of data for AI to be developed. How much?

How much data is needed depends very much on the ML algorithm used. Deep learning or other neural networks are the most data hungry, needing tens of millions, up to billions of data points. On the other end, algorithms for ranking and recommendations can get by with hundreds of thousands to millions of data points.

Humans have been constantly processing information drawn from our 5 senses from the day that we were born (and from some senses while we were still in the womb). If we just take visual images and conservatively estimate that a person will see a distinct visual image every 30 seconds, this results in 17,531,520 distinct images by the time someone is 25. And this is just from one of our senses.

We said above that AI is named because it mimics - somewhat - human learning. This is why huge data sets have been associated with AI for as long as it's been around. Most ML systems require millions, if not billions, of data points to be able to start making sound judgements.

A recent (2017) experiment led by Google and Carnegie Mellon University showed this with an AI trained in image detection. Most image detection AIs are trained on a standard data set of 1 million labelled images. This is the raw data that equates to the machine's experience. In this experiment, Google wanted to investigate whether access to more data would improve the accuracy of the AI image detection. Because they're Google, they had access to a data set of 300 million labelled images, which they used to train their AI. They found that as the amount of images in the training set increased, so did the accuracy of the AI.

Recently, huge advances have been made in training AIs with small data sets.

5

In the Google experiment, although the data increase was 300%, the observed increase in performance was just 3%. So, is it really worth going to the trouble of gathering that much more data? The answer may well be 'no' - although the jury is most definitely still out on this one!

Some companies, particularly those without the vast quantity of data that a company like Google has at its disposal, are starting to make use of something known as 'transfer learning'. With transfer learning, an AI that has been trained using a set of data can effectively "give" its knowledge to another AI that is going to use the same dataset, even if the outputs that the two AIs are generating are completely different. For example, an AI learning to recognise cars while driving can transfer the knowledge it has learnt to another AI that is trying to recognise trucks. It effect, one AI becomes the teacher of another.

When the terms 'Personalised', 'Adaptive' and 'Differentiated' learning are used by technology companies, what does this mean? Is this all AI?

The short answer is no. Personalised, adaptive and differentiated learning are outcomes that can be achieved in lots of different ways; AI is just one way to achieve them.

Most existing educational technology (EdTech) companies rely on rules-based technology to create personalised, adaptive or differentiated learning. This means that developers, usually in conjunction with teachers and pedagogical experts, explicitly tell the program what different routes are for students, based on what they've done previously.

Although the developers might write algorithms, create complex routes and base this on what the user has already done, the algorithms and the machine cannot move past this: they aren't doing any learning at all.

Rules-based adaptive learning platforms were beneficial because they were the first step along the road away from a 'one-size-fits-all' education. However, it is not genuine personalisation. Because the rules are human-derived and have to be explicitly programmed into the software, the number of routes available is finite. In practice, this means that learners are grouped into cohorts and each cohort has its own route. We can think of rules-based adaptivity like the standard clothes sizing: some parts of it will fit well, other areas not so much and you probably have to compromise on fit in one area to get a good fit in another. It was good when it was all we had, but it's not good enough anymore. AI can provide the equivalent of made-to-measure: genuine, full adaptivity and personalisation for a student.

Rather than developing specific rules for different routes, artificially intelligent learning platforms learn for themselves what the best route through the content is for a specific student in a specific moment. The ML algorithms will constantly learn, never leaving their training phase so that the intelligent insights they deliver are constantly improving, much like the best teachers are. This makes AI learning platforms significantly better for students than rules-based ones and allows education to fully move from the 'one-size-fits-all' model into genuine 'one-size-fits-one' personalisation.

CENTURY uses artificial intelligence to generate unique, truly personalised learning pathways for each and every student. As students complete diagnostics, learn and answer formative assessments on CENTURY, the AI picks up on their strengths, weaknesses and gaps in knowledge and immediately reacts to build on and scaffold the student as required. Not only does the AI determine what topic a student should study, it also learns what

material is most appropriate for a student and automatically differentiates to their learning needs.

What happens to my data?

This very much depends on what the AI company is trying to do with your data. But, in general, data will be stored in databases depending on the type of data it is. Most companies will keep the data about what you do on their site separate from the data about who you are. This is to ensure that private data, such as your name and email address, is kept secure, whilst data that is going to be used by the AI system can be processed and analysed.

Data collected by CENTURY Tech is split into three distinct categories: personal data, learning data and content. These three sets of data are stored separately from each other. Learning data is anonymised and used to improve our algorithms. Content is stored in our CMS database to allow learning. Personal data is, of course, private, stored securely and only used when needed.

What can AI do for me?

AI is a tool, just like any other. There is no point trying to use any technology just for the sake of it, it needs to be used where it can make a difference to teaching and learning.

Artificially intelligent learning platforms, like CENTURY, provide genuine adaptivity for learners: Students can learn at their own pace, while teachers and parents can rest assured that their child is adequately supported and challenged and making the progress they should be. See this quote from the economist about personalised learning if you need convincing:

> *"Recent studies show that software which imitates the responsive role of a tutor rather than just cranking out questions and answers can indeed accelerate children's learning."*

https://www.economist.com/news/briefing/21725285-reformers-are-using-new-software-personalise-learning-technology-transforming-what-happens

Teachers are also supported and enabled by AI. Take flipped learning, where students self-study a topic before it is taught by the teacher, allowing classroom time to be spent on higher-yield analysis, evaluation and problem solving tasks, rather than on basic knowledge acquisition. The EEF funded a study in conjunction with Shireland Collegiate Academy (Sir Mark Grundy's pioneering school) which found that flipped learning can add up to 2 months of progress each year. An AI platform like CENTURY can enable and support flipping the classroom: teachers can access teacher-created learning material and assign it to their students. In their own time, students complete this work, getting immediate feedback from the platform and supported by the AI insights. The teacher can access rich learning data in realtime without having to do any marking or data entry. At the beginning of the lesson following the flipped homework, a teacher knows exactly where each of their students has struggled and excelled and so can pitch their lesson to meet the needs of the students.

It might sound too good to be true, but **AI is here to change education.**