

# Algorithm's gonna get you

Are robots powered by artificial intelligence going to steal your job? That's unlikely, says **Kat Arney**, but she argues that AI does have the potential to transform the way we view education, how teachers teach and how everyone in schools behaves – for the better and for the worse



Let's get one thing straight: robot teachers are not coming to your school and they are not going to steal your job.

Quite how this view of the future role of artificial intelligence (AI) in education became the prevailing one is difficult to pinpoint, but we need to strip it of validity: the paranoia it creates is stifling open conversation about the real possibilities of AI in schools – and that conversation is one that needs to happen.

You might not realise it, but you are already interacting with AI every day: on your smartphone, when you shop online, as you check out recommendations on streaming services such as Netflix or browse the web. Businesses all over the world, from banking to manufacturing, are also turning to AI to analyse data and produce more effective results. Unsurprisingly, researchers and companies are busily working out how to do the same for education, with potentially huge benefits for teachers. But some argue ethical

questions that should be posed about this work have, thus far, not been asked.

Meanwhile, some believe AI might begin to change how we act and think. Could it make us lazy, change how we perceive repetitive tasks, or adjust our levels of patience?

The rise of home AIs, such as Amazon's Echo, has led to a slew of stories raising concerns that these devices are encouraging children to lose their manners. Unlike flesh-and-blood humans, who expect a certain level of civility in their interactions, our silicon friends don't care if you don't mind your Ps and Qs, or whether you play nicely. It's even led one San Francisco tech guru, Hunter Walk, to worry in a public blog post that his family's device is "turning our daughter into a raging asshole".

So no, AI won't replace teachers, but it does have the potential to fundamentally change not just what teachers do, but the nature of those they teach. It's already happening, in fact. And yet AI's creep into schools has largely been ignored and the classroom doors have been left open. Those in education need to get clued up. Fast.

Simply put, an AI is a computer programme (algorithm) teamed with real-world data, which can be trained to perform a task. Some are relatively straightforward – little more than sophisticated spreadsheets – but others are much more complex. Neural networks, for example, modelled on the interconnecting cells of the human brain, can pick through and process huge amounts of data.

The most common AI you'll encounter is probably your recommended items that pop up during your online shopping experience. The AI has learned your behaviour and knows you well enough to suggest things to tempt you into spending more money.

Rose Luckin, professor of learning with digital technologies at University College London, says that, at its core, AI is about developing technologies that can exhibit what humans would consider to be intelligent behaviour. "In the case of AI in education, it's about using that intelligent behaviour to support learning," she says. "And that could be at any age, any learning."

She draws a distinction between general AI – universal intelligence capable of doing

anything that a human brain can, such as the role of a human classroom teacher – and domain-specific AI, trained to do one particular task and do it well.

While some believe that the goal of a general AI is feasible, Luckin is more sceptical. "Personally I'm not a believer, because we still don't know what it is about human intelligence that makes us able to be generally intelligent," she says. So no robot teachers any time soon, then. "But within education, there's a lot that can be done with domain-specific AI that can be very useful."

### Man or machine?

One example of an AI put to work in an educational setting comes from Ashok Goel, computing professor at Georgia Tech in the US. In 2016, he stunned his students, who were taking an online master's degree in computer science, by revealing that one of the teaching assistants (TAs) they'd been messaging all year was an AI chatbot.

"In a typical year my students might pose more than 10,000 messages on the online discussion forum, but I just don't have time

to read them," he admits. "I felt terrible not being able to respond to all of them. We thought perhaps we could automate the answering of routine messages, so that the TAs and I could devote our time and attention to the more difficult questions."

Goel programmed his chatbot system with the answers to simple but common questions such as requests for information about the syllabus, deadlines, assignment formats and so on. Named Jill Watson in homage to IBM's flagship Watson AI, the bot learned from its mistakes and improved as the semester wore on. By the end, "she" was answering student questions with 97 per cent certainty, while batting unanswerable questions over to human TAs to solve.

Impressively, none of the class twigged that their helpful TA wasn't all it appeared to be – and, perhaps predictably, one even asked "her" out on a date.

"I only told the students after their final examination that Watson was actually an AI teaching assistant, and the response was uniformly and extremely positive – they were spellbound!" he laughs.

The following year he developed two AI TAs, the ostensibly female Stacey Sisko and her apparently male counterpart, Ian Braun, with similar success. This year, there are three, although their identities are a closely guarded secret.

Since introducing Watson to the world, Goel has been inundated with requests from overwhelmed teachers (working in education settings across the student age ranges) who are desperate to adopt his technology.

"The problem is that it takes a lot of effort to build this kind of teaching assistant, because everyone will have their own content to put in it," he says. "Watson took us 1,500 hours to build; we're hoping the next one will take 150 hours, and when we can get it to 15 hours, we hope it will become worldwide. We need to make it so simple that any middle-school teacher can spend a few hours and build their own, but we're not there yet."

For time-pressed teachers wading through a workload crisis, you can see how appealing such a tool could be, freeing up both their time and TAs' time for more complex questions or letting them get on with other

### The cost and infrastructure problem

No developer is going to waste time building an AI tool that runs on an obsolete browser through a snail's-pace connection, yet many schools are still struggling with flaky broadband, ancient hardware and outdated software.

Computing infrastructure costs money, as does AI development. But Rose Luckin, professor of learning with digital technologies at University College London, believes that a whole-hearted move towards using this technology would actually save more than it costs.

In a recent paper in the journal *Nature Human Behaviour*, she estimates that a detailed AI assessment tool for a core curriculum in England would cost about £500 million a year to implement. That seems like a staggering sum at a time of ever-tighter school budgets until you realise the estimated costs of the exam system run upwards of £800 million per year.

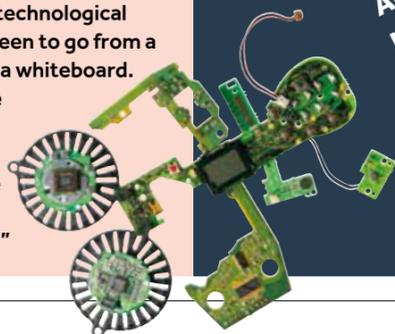
Luckin also suggests that switching to AI assessment

would free up a small army of exam setters and other workers to take up more useful roles – including teaching.

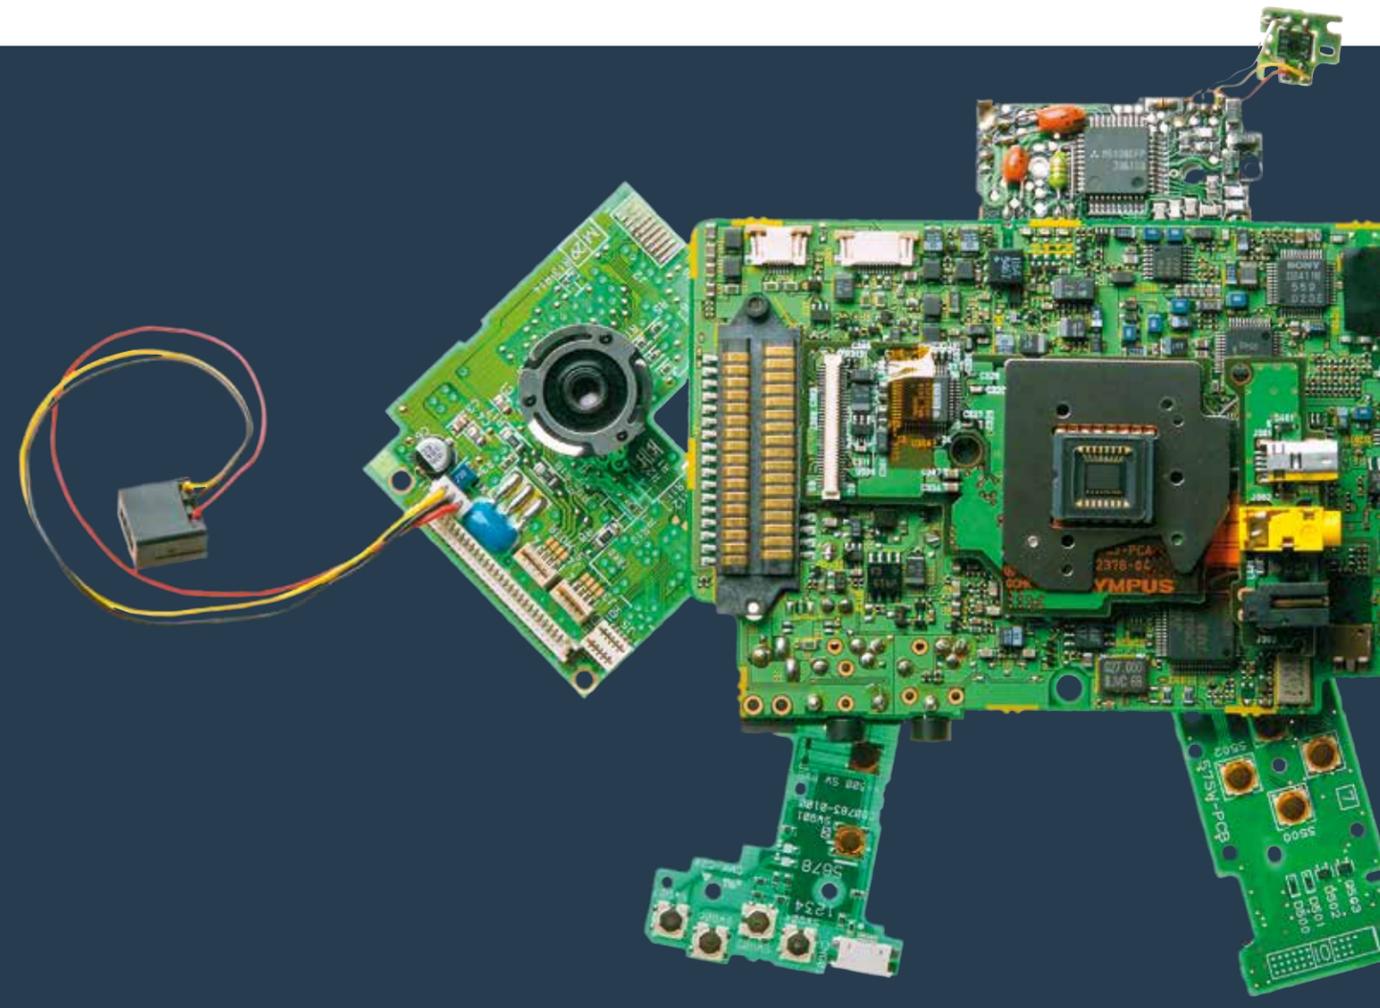
Yet the UK seems to be lagging behind with most of the work developing educational AI that is going on in the US and elsewhere in Europe. This is a major concern for Luckin and her colleagues in the field. "It would be nice to think there could be a magic moment when we buy into what AI can do, but I think it will be more gradual. It will happen, but it may not happen here post-Brexit," she says.

Priya Lakhani, founder and CEO of UK-based Century Tech, isn't giving up hope just yet, though. "Our education system was built during the industrial revolution, and since then the only technological advance has been to go from a blackboard to a whiteboard.

We should use technology to supercharge teachers – the status quo is unacceptable."



Artificial intelligence has the potential to fundamentally change what teachers do and the nature of those they teach



tasks. But there's more to educational AI than chatbots and online courses.

Some digital educational tools that are badged as AI are little more than recommender systems, pushing pupils towards relevant content. Others are more like Choose Your Own Adventure books, with a limited number of paths that students are shuffled down, depending on their aptitude and progress.

But an increasing number are offering something much more complex – a system that creates a truly personalised learning experience. A number of US schools have been trialling such systems, and some – such as the Rocketship network of schools serving low-income neighbourhoods – have integrated them fully into school life, with students spending time each day with tailored learning content on computers.

Priya Lakhani, founder and CEO of UK-based Century Tech, believes that these systems are where AI will really play a fundamental role in the future. Her company's system, built with input from neuroscientists and already operating in a small number of schools, is already proving what is possible, she claims.

"Students log in to our learning platform and the technology learns how each child is working," she says. "We're collecting every scrap of data, every mouse movement, each keystroke, every little thing, every nanosecond they're on there, anywhere in the world. It's aware of their skills, their gaps in knowledge, their strengths and weaknesses and their focus level – even when something goes from short-term to long-term memory. This will strengthen the personal pathways in their brains to get their recall quicker, and that helps with problem-solving, deduction and all the other skills they need to learn."

Century's AI looks for patterns and correlations in the data from the student, their year group, their school and even the entire population to offer a personalised learning journey for the student. It also feeds the whole lot back to the teacher in the form of a kind of dashboard, giving them a real-time snapshot of the learning status of every child in their class. The information can be used for swifter lesson planning, assessment and reporting back to the school management on progress against goals – freeing up time, says Lakhani, for teachers to concentrate on teaching.

"I have never met a teacher who wanted to go into teaching because of the data management," she argues, highlighting the current crisis in recruitment and retention in the profession. "[This system could save] an hour a week per class, so for one of my

maths teachers, that's six hours a week. We're giving him his weekends back and he's able to go in feeling energised. You want to go into a class knowing exactly where your kids are in terms of their understanding and knowledge. This technology can tell you immediately, and that instils confidence."

Those are some substantial claims that will raise eyebrows of doubt among many. And although AI might save time and reduce stress, the advent of this kind of technology raises fears that teachers themselves might be reduced to little more than crowd control and hand-holders, merely keeping watch while their students beaver away online. Those fears are unfounded, says Lakhani.

She argues: "Teachers pick teaching because they want to do the bit that only a human can do. None of them came in to do the micromarking, data management and analysis."

This type of role for AI fits with where Luckin seems to suggest the technology will be of most use to teachers.

"The benefits are much more around augmenting the skills of the teachers; they're absolutely not replacing them, but instead finding ways that teachers can use their uniquely human skills more effectively," she says. "It's the idea of artificial intelligence working together and augmenting human intelligence. For the teacher, for the learner, for the employer – it's about making it easy to address these big issues in education by blending human and artificial intelligence."

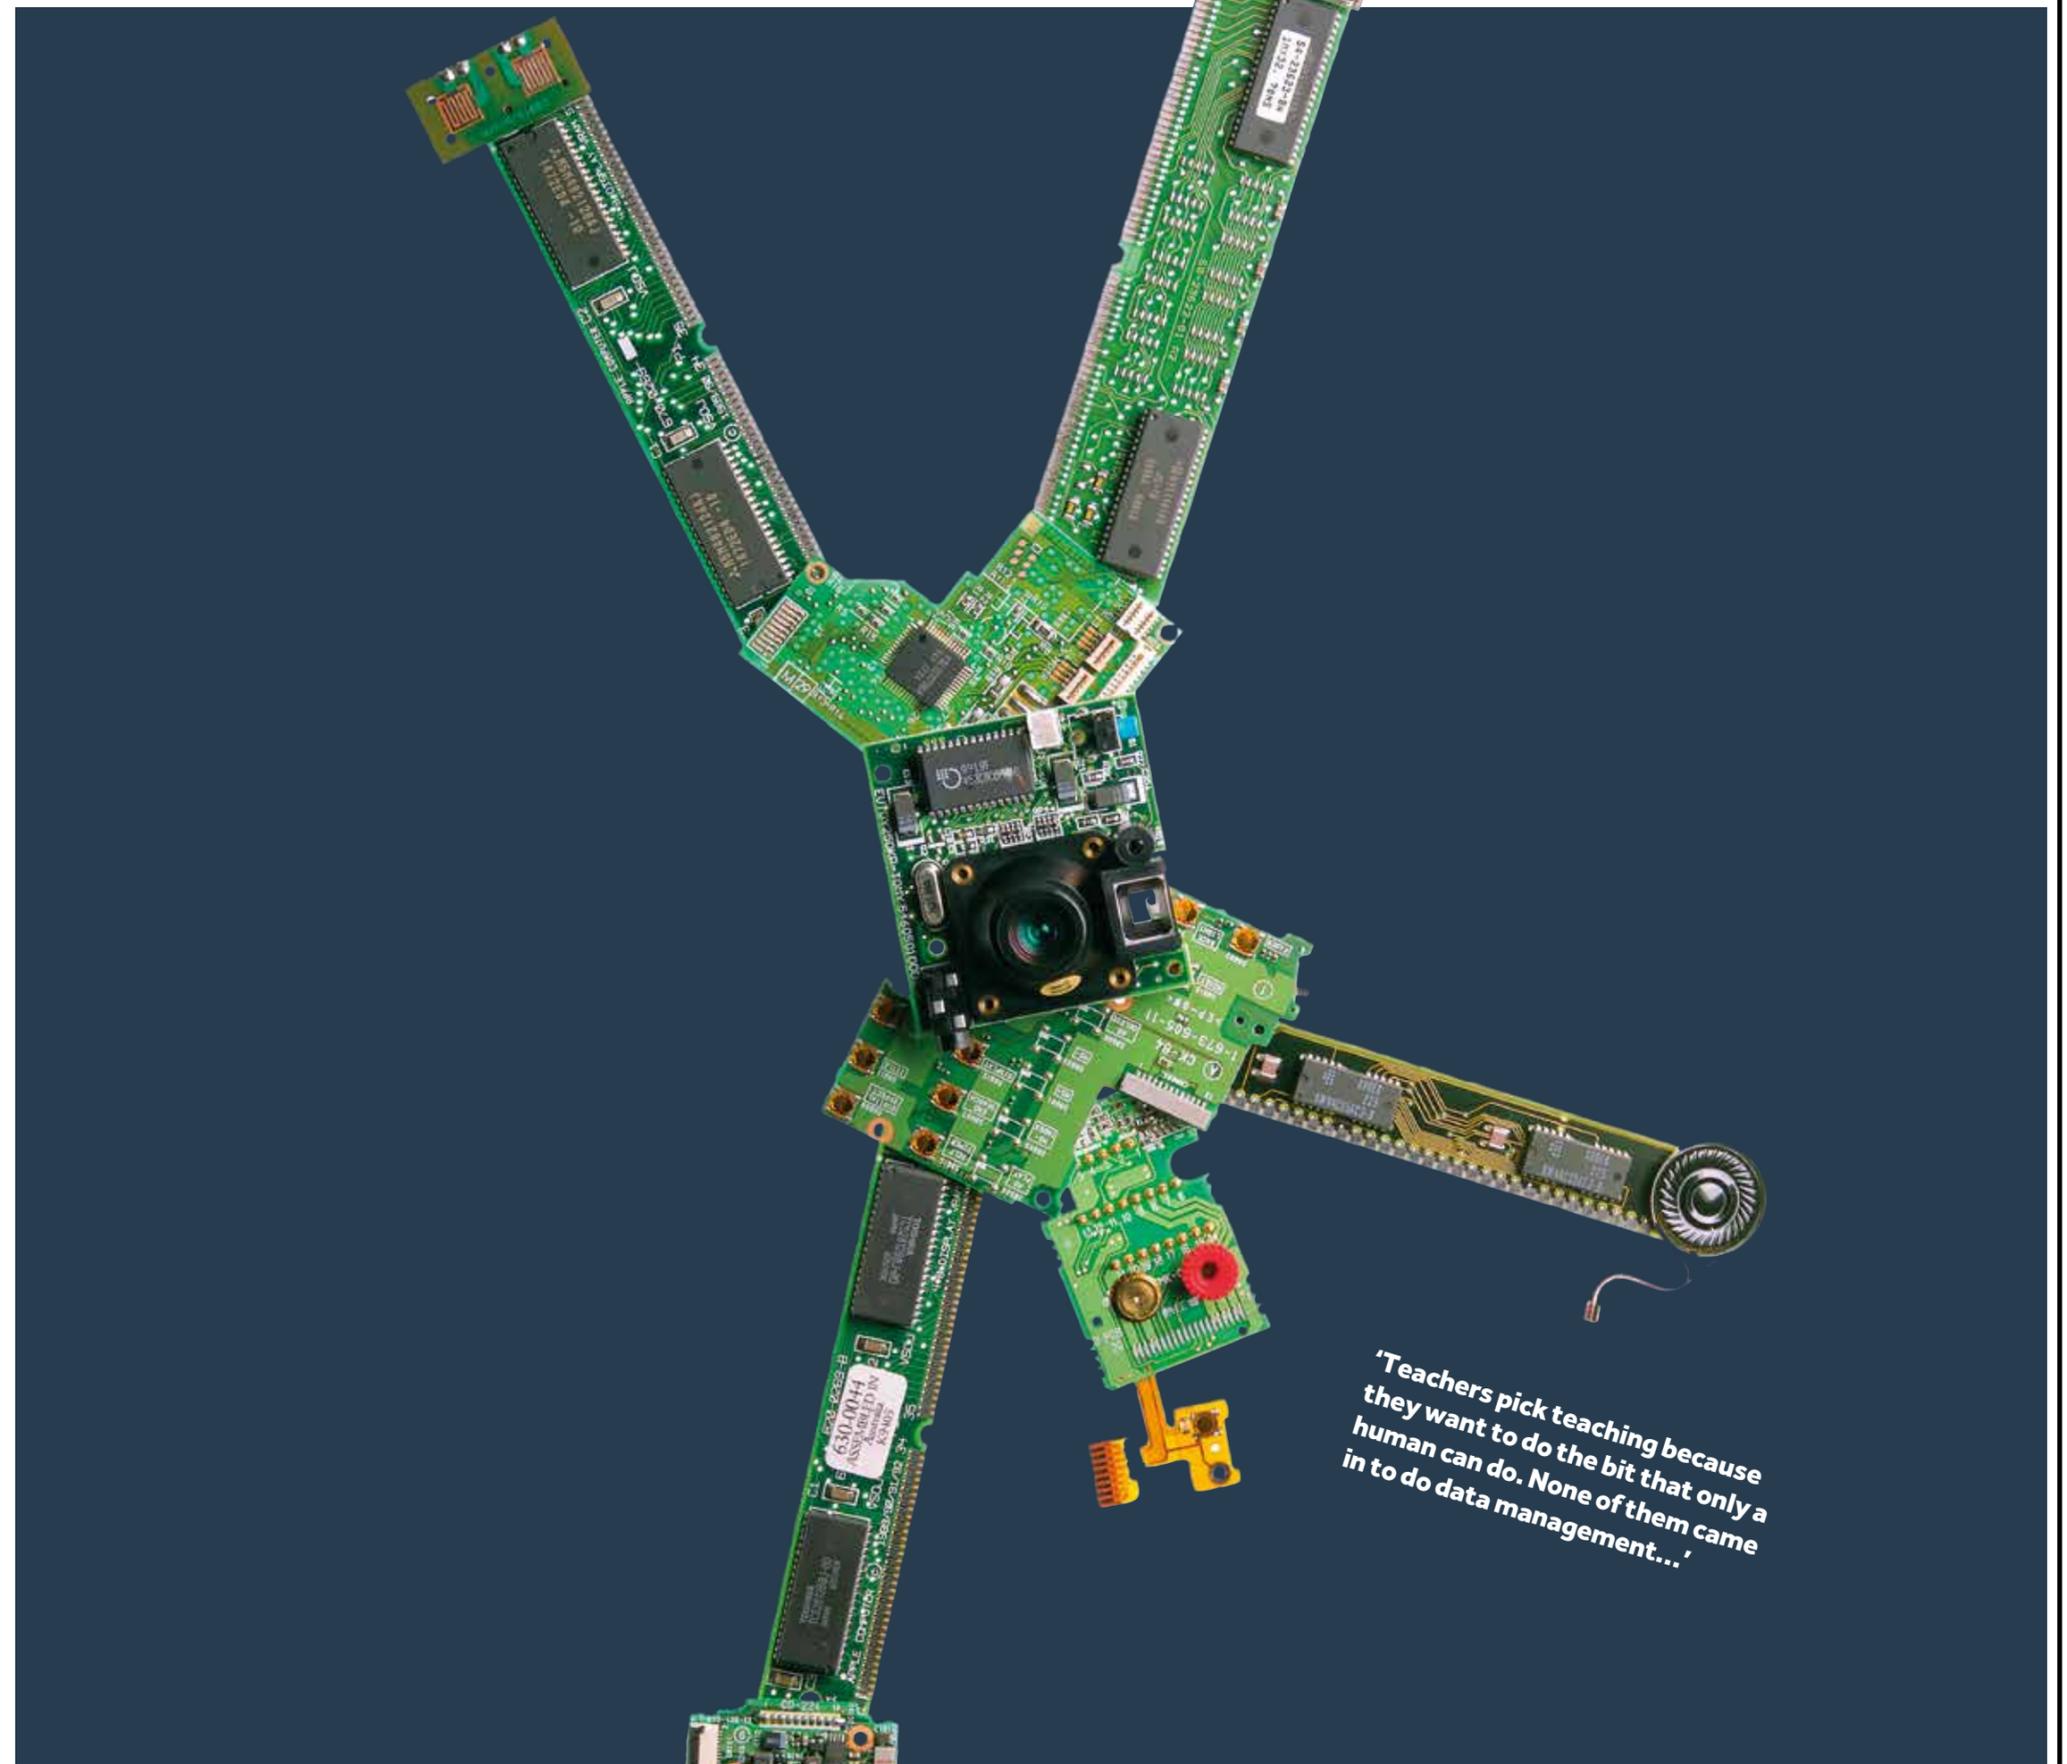
### Lacking the human touch

This future AI-integrated school might sound a little less gimmicky and a little more procedural than you might have imagined, but while it makes for much less dramatic headlines, the effect on students and teachers could be huge.

Of course, those effects could be negative as well as positive. Some believe that soon-to-be teenagers in the school system, who have grown up with AI, might be affected by the unsavoury side effects of the technology.

For this generation, interacting with computers and the AI behind them will be second nature. This may change how they interact with people: AI does not require politeness, small talk or indeed much in the way of social niceties. It could also change expectations of people's own role in their learning: if everything is usually offered on an AI-created plate in everyday life, what effect will that have in the classroom when teachers ask students to act on their own initiative, or to put in some hard graft?

Here, those companies developing AI tools perhaps need to be a little more aware of



the effect of their creations. They have the power to create tech that avoids such pitfalls, according to Alan Winfield, professor of robot ethics at the University of the West of England in Bristol, as computers are still constructed by humans and we have a choice in how they work. “This is basically an oversight on the part of the developers of conversational AIs,” he says, “because they could have built-in protocols that require you to, [for example], ask nicely and encourage or prompt a “thank you” at the end.”

He argues that this is part of a wider failure of ethical oversight within the tech sector: “Many of these AI developers do not have an ethics panel, or even just someone to say, ‘Have you thought about this?’ I believe this is an industry that is forging ahead with quite extraordinary AI developments without thinking of the societal consequences – and in this case the consequences for development of very young children, which is serious.”

As a solution, he wants to see ethical reviews and standards embedded deeply into the creation of AIs, particularly those used in an educational setting. The Institute of Electrical and Electronics Engineers has a Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems, and is developing guidelines for what it calls ethically aligned design – similar to the British Standards Institution’s Kitemark – to show that an AI product has been designed with these concerns in mind.

### Bad robots

It’s not just about how educational AIs might shape society in the future, though. It’s also important to turn the question the other way and think about how society shapes AI. When it comes to using machine learning, such as language-based tools like chatbots, an AI is only as good as the data that’s fed into it. As William Stewart’s feature highlighted this month in *Tes* (“Drowning in data”, 5 May), a lot of education data is highly suspect. And if it’s human data, it will come with human variables. How, for example, does the AI recognise why a student is performing in a certain way on any given day?

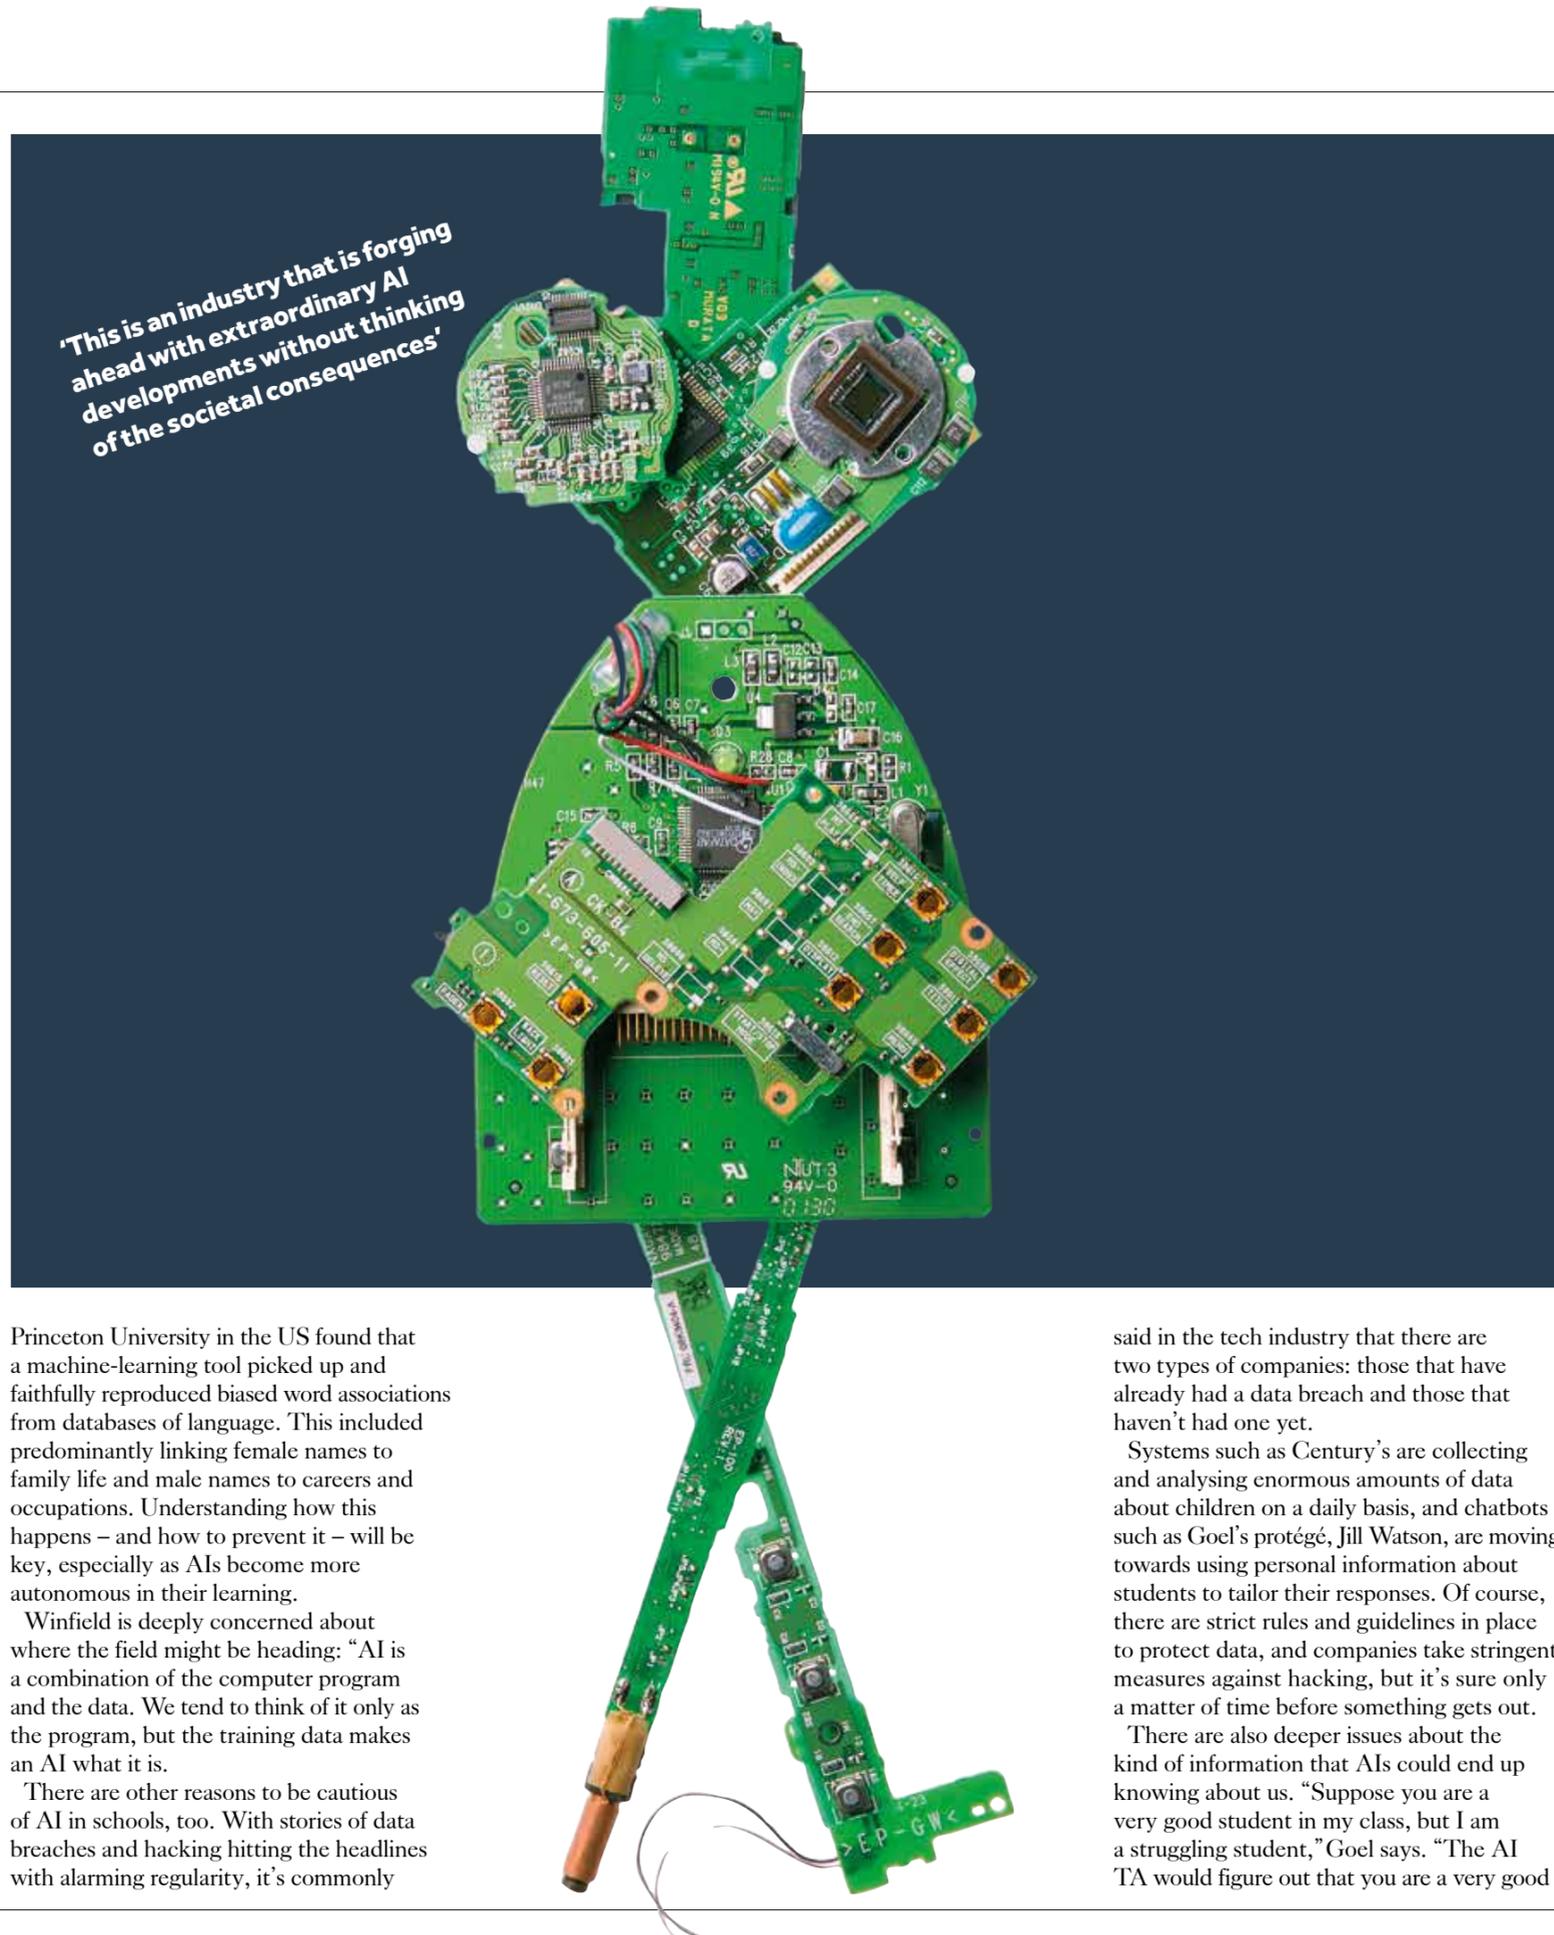
Any number of issues may impact performance, yet the data will only tell one story. Equally, human biases will be problematic. For example, after a mere 24 hours interacting with people on Twitter, Microsoft’s Tay chatbot started coming out with racist comments.

These biases might not be explicit, or even conscious, but they’re there in the data used to “train” AIs. Publishing in the journal *Science* earlier this year, researchers at

Princeton University in the US found that a machine-learning tool picked up and faithfully reproduced biased word associations from databases of language. This included predominantly linking female names to family life and male names to careers and occupations. Understanding how this happens – and how to prevent it – will be key, especially as AIs become more autonomous in their learning.

Winfield is deeply concerned about where the field might be heading: “AI is a combination of the computer program and the data. We tend to think of it only as the program, but the training data makes an AI what it is.”

There are other reasons to be cautious of AI in schools, too. With stories of data breaches and hacking hitting the headlines with alarming regularity, it’s commonly



said in the tech industry that there are two types of companies: those that have already had a data breach and those that haven’t had one yet.

Systems such as Century’s are collecting and analysing enormous amounts of data about children on a daily basis, and chatbots such as Goel’s protégé, Jill Watson, are moving towards using personal information about students to tailor their responses. Of course, there are strict rules and guidelines in place to protect data, and companies take stringent measures against hacking, but it’s sure only a matter of time before something gets out.

There are also deeper issues about the kind of information that AIs could end up knowing about us. “Suppose you are a very good student in my class, but I am a struggling student,” Goel says. “The AI TA would figure out that you are a very good

student, based on your academic history, and could also tell that I was not doing so well. So it might give a different response to your question than it would to me. I find this a bit creepy, because on the one hand it’s really positive if the AI is gathering data about all the students and learning from it, but on the other hand, it means that it is collecting [private] data about you [that is private]. Who gets access to the data? Right now, I cannot share that information with anyone, not even the student’s parents, because that is confidential.”

### Deal with the ‘data devil’?

Winfield describes this dilemma as a kind of Faustian pact that we make with technology in exchange for its advantages.

“What people need to understand is the personal risk they are opening themselves up to. There’s all this amazing technology that seems like it’s free, but it’s not, because you hand over your data,” he says. “The amount of information that AI companies are collecting from us is often more than they can justify to support their business model, and of course this data is very valuable and can potentially be monetised.”

Luckin admits she shares some of these concerns. Because algorithms are designed by people, she says, they “therefore run the risk of being biased”. She also agrees that the market for personal data is a “huge worry” and that data security is also a challenge.

But she thinks that educating the public about these issues is the best form of defence, rather than ignoring AI altogether.

When it comes to ethics, Luckin stresses again that valid fears about implications should not mean shunning the technology. “There is too little concern with ethical issues and implications,” she says. “Much of what I do is to try and prompt this sort of debate, so that we can make informed decisions about what type of AI we do want for society and education in particular.”

Lakhani stresses that with the right “vision, culture and safeguards”, AI can be a force for good. She says privacy is something that everyone who works in AI takes seriously and that, at Century UK, “not only have we self-certified with the DfE as compliant with their privacy standards for cloud services for schools, we’ve gone beyond these requirements with regards to data safeguarding”.

She adds that the only time data would be shared would be for research and development with “well established” education research bodies and the data would be anonymised.

Finally, on the “garbage in/garbage out” issue and ethics risks, Lakhani concedes that she is aware of the problem. “This is

something with which all AI companies should concern themselves,” she says. “Credible AI companies should never make assumptions and never make absolute claims about their users. We employ highly qualified data scientists who ensure that our machine learning algorithms never make assumptions, and rely on traditional, respected and rigorous statistical methodologies.”

“In any field, initial data is at risk of being misleading, so it is an ongoing challenge for all AI companies to collect meaningful data.”

But as AI creeps into the classroom, perhaps even becoming integral in the default educational system in the future, some parents may not believe the assurances of those developing the AI tools, or of those like Luckin who have researched it. They may decide that they wish to shun this deal with the “data devil” and take their children out of it. This could be seen as the technological equivalent of home schooling: potentially advantageous to those who want to keep their data to themselves, but a big disadvantage for a child’s future if AI-based educational tools turn out to equip them better for life than current teaching methods.

Of course, that’s not proven as yet. But if that does happen, then Luckin says AI’s future in education will depend on more than just that evidence. She describes the challenge as winning over “hearts and minds” – AI doesn’t just need people to believe it works, it needs people to trust it, to believe that it does not endanger that sacred relationship of trust between a student and their teacher. Parents need to trust that it is safe, too.

For advocates of AI in education, this will be the real battle if the potential negatives are to be overlooked for what she and others believe are the substantial, and education-changing, benefits.

“We need to persuade people that the benefits that can be gained through the combination of big-data AI processing outweigh the fears,” Luckin argues. “It is more poignant in education than in any other area – people are very careful when it comes to their children – and I can completely understand if parents say they don’t want that for their child.”

“But that will be such a great shame because it can do so much.” ●



Dr Kat Arney is a science author, broadcaster and co-presenter of the BBC Radio 5Live Science show. She tweets @harpistkat