Edtech

Technology is transforming what happens when a child goes to school

Reformers are using new software to "personalise" learning

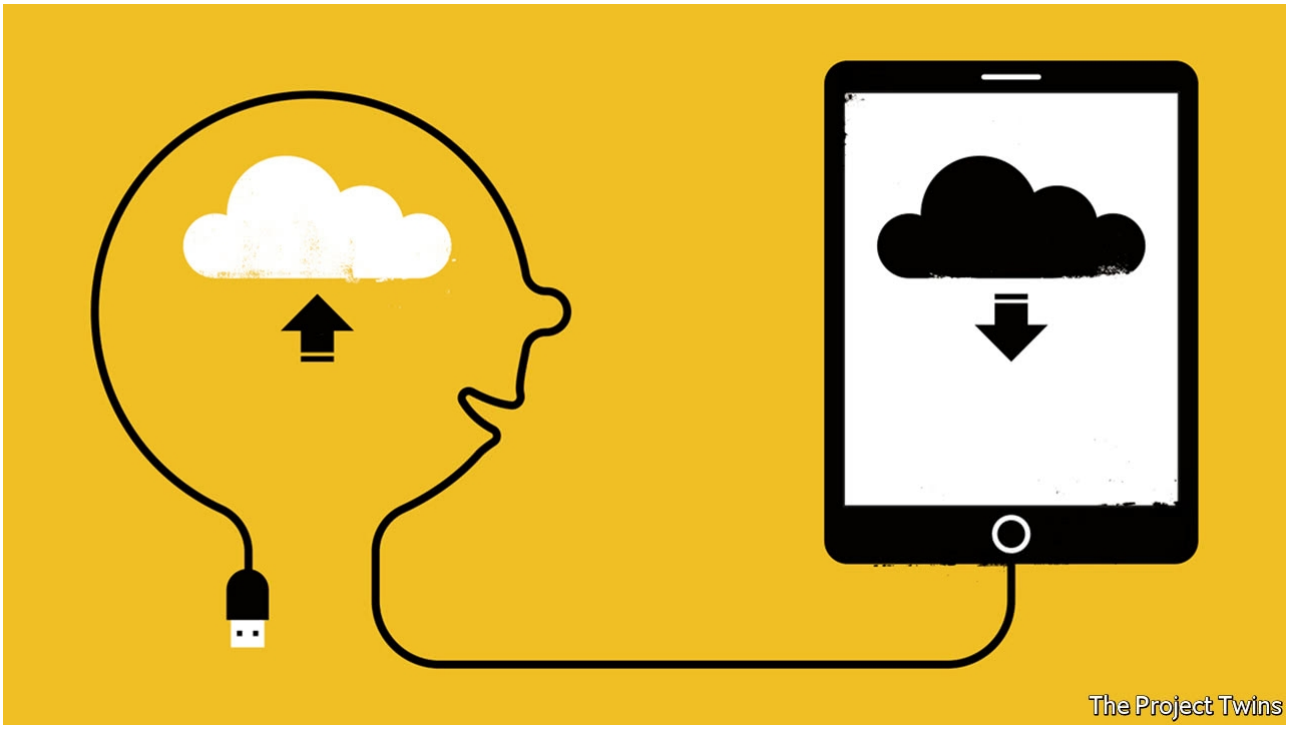

Print edition | Briefing Jul 22nd 2017 | BANGALORE AND THE SAN FRANCISCO BAY AREA

FOR a ten-year-old, Amartya is a thoughtful chap. One Monday morning at the Khan Lab School (KLS) in Mountain View, California, he explains that his maths is "pretty strong" but he needs to work on his writing. Not to worry, though; Amartya has a plan. He will practise grammar online, book a slot with an English teacher and consult his mentor. Later he will e-mail your correspondent to ask for help, too.

This is the sort of pluck KLS produces. Its pupils do not have homework or report cards or spend all day in classrooms. They are not stratified by age; they share common spaces as they pursue individual goals and schedules, using software built by in-house developers to take tests and watch video lessons from the school's sister organisation, Khan Academy, which makes online tutorials. Half the teachers act like tutors, helping with academic work. The rest mentor pupils in character traits such as curiosity and self-awareness.

Latest updates

An end at last to the Charlie Gard case
BRITAIN

The revival of the great British lido
PROSPERO

The euro's obituaries were premature
BUTTONWOOD'S NOTEBOOK

See all updates

The idea of using technology to revamp education is not new. In 1928, Sidney Pressey, a psychologist, invented a "teaching machine" which he imagined "freeing ...teacher and pupil from educational drudgery". The automaton had a paper drum displaying multiple-choice questions. Pressing the right key moved the drum on, yielding sweets for smarty-pants.

Despite its sugar-coated bait, Pressey's teaching machine went the way of most such technology. It did not live up to the hype. Since then a succession of inventions promising to overhaul schools has done no such thing. Information technology has reshaped other sectors; it has had little impact on education.

This has not been for want of hardware. In 1984, the year the first Macintosh was launched, American schools averaged one computer for every 125 pupils. By 2012 there were five for every nine. But this big bang in access to IT had "little or no positive effect" on outcomes such as test scores, according to an analysis of trials from around the world published last year by George Bulman and Robert Fairlie of the University of California. In 2015 the OECD found no link between what countries spend on IT in schools and their 15-year-olds' abilities in maths, science and reading.

Now, though, the stasis is finally starting to shift, for two reasons. The first is that "edtech" is increasingly able to interact with students in sophisticated ways. 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. The second reason is the experience of a growing number of schools, like KLS,

which are not just bolting edtech onto the existing way of doing things but using the new software to change how pupils and teachers spend their time. Both, it seems, get more productive. For many decades educational innovators have happily anticipated the end of "factory model", whereby children of the same age learn from the same teacher in broadly the same way, yet the model endures. Now, at least in some places, its days seem numbered.

Investors, both philanthropic and otherwise, are excited. Edtech is one of the priorities of the investment fund set up by Mark Zuckerberg and his wife, the Chan Zuckerberg Initiative (CZI). He wants most American schools to adopt the new sorts of education it promises within a decade—and then help spread it worldwide. The combined value of the North American and European edtech markets (including further and higher education as well as schools) is set to grow from $75bn in 2014 to $120bn in 2019, according to Technavio, a research firm.

Research in two fields is shaping the new technology. Artificial intelligence (AI) is letting machines learn about the pupils using them by studying the data produced in the process. And research drawing on psychology, cognitive science and other disciplines is providing practical insight into the "science of learning".

The late American psychologist Benjamin Bloom convinced many educationalists that overcoming the failings of the factory model required making group instruction more like personal tuition—which his studies showed to be the most effective form of teaching. "Adaptive learning" software, first developed by computer scientists in the 1970s, aspires to mimic tuition's one-on-one strengths. Such programs use pupils' answers to inform their choice of subsequent questions, adjusting the difficulty as they go along.

Machine learning, a branch of AI that allows computers to pick up on patterns they were not explicitly programmed to perceive, lends itself well to this approach. But it is not essential. Mindspark, developed by Educational Initiatives, an Indian company, simply draws on a bank of 45,000 questions and the 2m answers generated every day. Its developers have anticipated common mistakes, using more than a decade's worth of pupil data and written code to diagnose the errors. For example, children often say that 3.27 is greater than 3.3, or 4.56 is greater than 4.9; the reason is that they are seeing the "27" and the "56" after the decimal points as

being larger than the "3" and the "9", an error known as "whole number thinking". Mindspark will pick up on this pattern of error and recommend specific remedial exercises.

Newer programmes being developed around the world use machine learning to find pupil-specific patterns of error and strength. Leading American brands include ALEKS, Knewton and DreamBox Learning. Siyavula Practice, a South African product, is used by more than 32,000 pupils in 388 schools to teach maths and science. Geekie has been used by 415,000 pupils in São Paulo's public schools, and by many more at home. Byju's, another Indian education company, received $50m in an investment round led by CZI in 2016. In China 17zuoye ("homework together") uses voice-recognition software to help students learn English. If a child says "seven potato", or "nine apple", 17zuoye will offer help with plural nouns.

Rapid progress in speech recognition and generation may take such ideas further. Researchers at the ArticuLab at Carnegie Mellon University have used voice-recognition technology to develop Alex, a "virtual peer", who talks to children in a vernacular that makes them feel more comfortable in class. Their findings suggest that some black children learn science quicker when they interact with a virtual peer using African-American vernacular than one speaking with a standard dialect.

Some of these companies pay close attention to the science of learning. Siyavula's algorithms adjust its questioning so that users get the right answer about 70% of the time. That is roughly the success rate, it says, that neither bores nor deflates learners. ALEKS, meanwhile, eschews multiple-choice questions. Instead it requires users to type responses—a more taxing method. Both products periodically return to topics; studies suggest "interleaved" practice helps facts stick.

A forthcoming paper by Philip Oreopoulos and Andre Nickow for J-PAL, a group at MIT which looks for evidence about what actually works when it comes to alleviating poverty, reviews dozens of randomised controlled trials involving edtech. In nearly all the 41 studies which compared pupils using adaptive software with peers who were taught by conventional means the software-assisted branch got higher scores. In most studies, language scores were higher, too. "There are not many other interventions with credible evidence showing these kinds of effects," says Mr Oreopoulos (see chart).

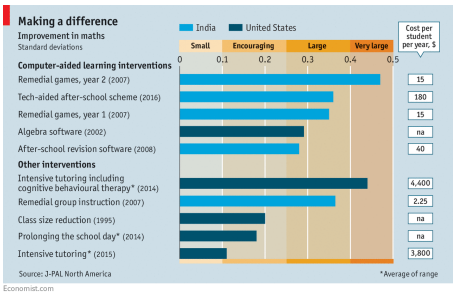

One study in the J-PAL review (http://bit.ly/2eAuvxE) is a paper by Karthik Muralidharan, Alejandro Ganimian and Abhijeet Singh, which looks at an Indian after-school scheme where children used Mindspark for 4.5 months. They found that the progress made in language and maths by those pupils was greater than in almost any study of education in poor countries—and for a fraction of the cost of attending a government-run school.

In part this is a function of a low baseline. Indian curriculums are far too ambitious, artefacts of an era when schools were the preserve of the elite, and at any given time a quarter of the teachers will be absent. About half of India's ten-year-olds cannot read a paragraph meant for seven-year-olds. One particularly encouraging aspect of the study was that it seemed to show those least-well-served by the current dispensation benefiting most—the poorest performers saw larger improvements than those who had previously been getting by.

Analysing published studies may not give a full picture of the field's progress: as in many areas of research, studies with ambiguous or negative results may never make it to publication. It is also much harder to judge the technology in softer subjects—fields where mimicking a tutor is undoubtedly harder. How to improve the argument of a history essay is not something edtech easily grasps, any more than it could advise on the use of humour in a drama class. But it can still help teachers' assessments in these fields. No More Marking, a British company, shows teachers paired excerpts from pupils' essays and asks them to decide which is better; with enough such comparisons its "comparative judgment" algorithms can then rank the pupils. The method saves teachers' time and helps pupils, too. They are less likely to suffer because a teacher is hungry, or tired, by the time of the last essay.

No dark sarcasm

It is also worth noting that the same system can show different effects in different trials. A study published in 2014 found that pupils using Teach to One: Math learned faster than the national average, according to a standardised test. But research that came out a year later could reach no conclusions as to its impact. A

study of another system, DreamBox Learning software, found that its impact differed from school to school. When it was used for 60 to 90 minutes a week, as its producers intended, and their suggestions as to how to get the most out of it were followed, it had much better effects.

Seeing Teach to One: Math in action underlines how much change is needed to make it work—which may explain why it does less well in some studies than others. When pupils at the Ascend School in Oakland arrive for their daily hour and a half of maths, they look up at monitors resembling airport information screens which tell them what and how they will learn today. One child is to work on geometry in a group; another will take algebra questions on his laptop. Three teachers walk around the open space, checking on pupils' progress. At the end of the session pupils take a short test, which is used by developers at New Classrooms, the charity behind Teach to One, to set children's schedules for the next day. Wendy Baty, the school's head of maths, is an enthusiast; she says that pupils receive feedback that "even the best teacher could not provide to all of the class". Several pupils say they like that they can learn at their own pace. But others admit to finding the experience confusing.

Rather than working for a few hours in a conventional school, other reformers are opening their own. AltSchool is one of various comprehensive attempts to use edtech to provide a form of "personalised learning"—thus slotting it into a tradition that reaches back to Jean-Jacques Rousseau and Maria Montessori. Founded by Max Ventilla, a former Google engineer, it is backed by, among others, CZI and the Omidyar Network, set up by Pierre Omidyar, the eBay founder. At each of AltSchool's seven "lab schools" in California and New York, pupils consult two pieces of software on their tablets. The first is the "portrait": a record of a child's progress in academic subjects and social skills. (One measure is whether children can "respond with wonderment and awe".) The second is the "playlist", which is where pupils gain access to material and complete work.

Perhaps surprisingly, and reassuringly, for a school so dependent on software, screen time is limited to no more than 20-30% of the day. The emphasis on project work means pupils collaborate with each other. At the Yerba Buena AltSchool, in San Francisco, Hugo, 12, explains that he learns more from his peers here than at his old school. Teachers at AltSchool say they save time by not marking or planning

lessons. Instead they analyse data on pupils' portraits and tutor them on individual problems. Hugo says "I feel like the teachers here really know me."

Giving children such attention is not cheap. Hugo's parents pay $27,000 per year, more than twice the average spending per pupil in OECD countries. That does not mean that the software AltSchool is developing will be particularly expensive. But overall cost is definitely an issue. Many of the public schools trying to combine edtech and personalised learning are supported by philanthropic organisations such as the Gates Foundation. A study last year of early adopters by the Centre on Reinventing Public Education at the University of Washington, also partly funded by the Gates Foundation, concluded that those schools' "long-term financial stability is still unclear".

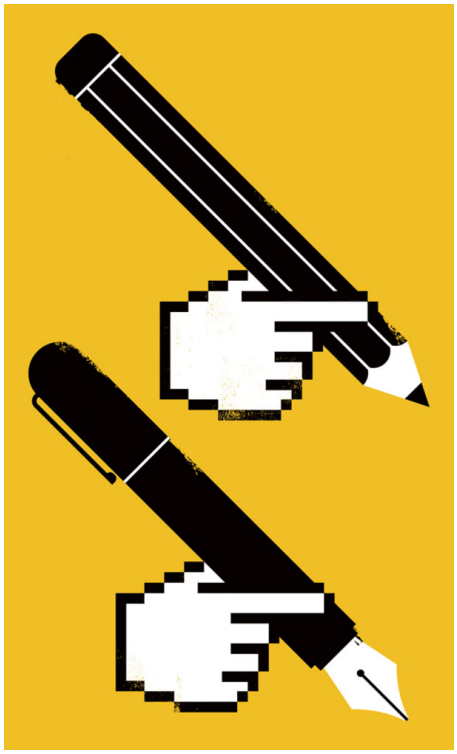

Arguably the most influential attempt to find out whether high-tech personalised learning can both work and be afforded at scale is that of Summit Public Schools, a publicly funded network of 11 schools in California and Washington which serve mostly poor, often Latino students; 130 more "partner schools" across 27 states use Summit's software and get training from Summit staff. Its platform was built *pro bono* by Facebook engineers.

Andrew Goldin, Summit's chief of schools, argues that the Summit Learning Platform lets pupils learn more efficiently than they do when led through every lesson by a teacher: "Children don't need to be walked through every step." That gives them more time to spend on projects, which take up half of the school day, and to be mentored by teachers.

Some information first

This sort of personalised learning has its critics. Putting students in charge of how fast they learn worries some cognitive scientists. "Our minds are not built to think," argues Benjamin Riley of Deans for Impact, a charity championing the science of learning. Thinking hard about things does not come naturally, and if schools make

it easy to avoid thinking, some children will do so. Another criticism is that people need a ready store of facts if they are to develop many forms of creativity and critical thinking (an insight championed by one of the early giants of AI, Herbert Simon). As Daniel Willingham of the University of Virginia puts it: "knowledge is cumulative". In the always Googleable world of tablet and phone, it could be tempting for children not to fill that store, and for their teachers not to worry too much.

Giving children more control over their learning, Mr Goldin argues, motivates them; if pupils do not grasp the basics they cannot participate in projects. He also points to Summit's results. About two-thirds of pupils score as well as or better than demography would predict in a nationwide maths test. In 2015, 93% of pupils who entered Summit went on to graduate, ten percentage points more than in comparable neighbouring schools. Of those graduates 99% got to university.

Achievement First, a group of 34 schools on America's east coast which is famed for tough discipline, is testing a similar model. So too are schools in cities like Chicago, New York and Boston. More than 3,000 superintendents (the officials who run America's school districts), representing about one-third of pupils at public schools, have signed a pledge to "transition" to "personalised, digital learning".

How well the model will work when it spreads is unclear. In 2015, the RAND Corporation, a think-tank, published the most thorough study yet of schools using high-tech personalised learning. It compared test results of pupils at 62 such schools with those of similar pupils at ordinary schools. The former made greater progress, especially those who started near the bottom of the class.

The report is widely cited by advocates of personalised learning. Mr Zuckerberg uses it to claim that: "We know that personalised learning is way better." That is a stretch of yogic proportions. The results are from early adopters of the model, with highly motivated teachers. And the RAND researchers were not able to work out what it was schools were doing to gain their results. Without that understanding expanding the model will be tricky. A further RAND report, released on July 11th, reiterated these concerns.

Teachers may be more sceptical away from Silicon Valley. And parents may be more concerned about privacy. Machine-learning software has an incentive to accrue data; they make predictions more accurate. New platforms contain accounts of a child's abilities far more detailed than any report card.

Supporters and sceptics of the new model will continue to argue. But both sides are guilty of caricaturing the other. Techies can make it seem as if teachers in ordinary schools talk to every pupil in the exact same way. They do not; studies repeatedly show that teachers use "differentiated instruction" among pupils of different abilities, even if they cannot offer one-on-one attention.

But schools using personalised learning are not anarchic playgrounds. Pupils may have more power but they do not have complete control. "Unadulterated choice is not good," says Aylon Samouha of Transcend Education, a charity. "You need standards and structure."

If schools can combine personalisation and rigour it is hard to imagine pupils failing to benefit. Education software is not making teaching obsolete. If anything it is making the craft of teaching more important. That would be good news for the staffroom and the classroom. For as 12-year-old Hugo observes, "too many teachers are just trying to get to the end of the day".

This article appeared in the Briefing section of the print edition under the headline "Machine learning"