Chapter 3 - II

Artificial Intelligence in Education

: Finding Its Place[1]

Joël Boissière and Éric Bruillard*

Artificial intelligence, or at least advanced IT including what is called 'artificial intelligence' is topical: its impact on jobs and a great number of activities is apparently quite significant. That is why numerous reports have been written on it, including Cedric Villani's in France in 2018. As we have described in the second part, China has made it a priority, and aims to become the world leader before 2030.

It undoubtedly poses major challenges, especially in transportation, health and security, but what about education? Significant budgets are being allocated (1.5 bn Euros in France) and one of the first measures has been to reinforce higher education training on artificial intelligence and the domains it is transforming, notably Masters' degrees in IT and engineering diplomas. As for the rest, the share for education has not been above 12 million euros (less than 1 %), which shows a rather *secondary* priority. What can artificial intelligence change in school education, in more or less formal training and in higher education?

To answer that question, we are going first to examine a few historical points. Using artificial intelligence in training is not a new idea. The first research dates back to more than half a century and we are going to present a few historical landmarks and show some major realisations and different limits that have already been well described. To some extent, it is a history of non-fulfilled promises. Is it going to be played again the same way?

We will then discuss personalisation and its two main meanings in education, and then adaptative learning. The latter is based on systems which take into consideration the student's level and offer them things that are supposedly tailored to... quite a lot of things actually: their knowledge (on the issue at hand but also on other topics), their multiple preferences, their interests, the time they have, the level other learners have reached (which cannot be completely ignored), etc.

Next, we will study three main topics that China has chosen in its approach to AI in education: big data,

* **Joël BOISSIERE:** Project leader of "Enfance protégée" at Bank de territoire. ; **Éric BRUILLARD**: Corresponding author, Full professor at Paris-cité University, eric.bruillard@u-paris.fr

1 The chapter entitled "Intelligence artificielle dans l'éducation : une place à trouver", was originally published in *L'école digitale. Une éducation à construire et à vivre* By Joël BOISSIERE and Eric BRUILLARD © Armand Colin 2021, Malakoff. ARMAND COLIN is a trademark of DUNOD Editeur – 11, rue Paul Bert - 92240 MALAKOF.

intelligent learning and platforms. We will refer to the central topics of explainability and algorithm governance, which will allow us to situate challenges linked to artificial intelligence in education, and highlight that the objective in education is not that systems be clever, but that human beings may effectively better learn and in *good* conditions.

1 AI and education – an already old history with new ramifications

At the end of the 1990s, so-called programmed learning reached its limits – everything must be prepared in advance, and if the features of the population that must be trained evolve a little, the 'programme' must be changed. In addition, the 'dialogue' and the path that were offered to learners were quite rigid (see Bruillard, 1997, for a detailed history of that idea). Was it possible to design machines that may sustain a somewhat free educational interaction with a student, and possibly let them ask questions, present them with knowledge and ensure that they learn it thanks to an open dialogue? That was the challenge for AI, and as early as 1970 Siklóssy suggested that what he called tutors, which know what they are teaching, be designed. Integrating a solving module, the computer must be capable of resolving problems posed by the student, explain how it does so and thus teach its own methods to the student.

In line with machines that teach the programmed teaching, the word 'tutor' is used, for it is an individual interaction (one student, one machine), and thus a sort of tutorship and not teaching to a group of students. Dialogue is conducted in natural language and the learner can take the lead and ask questions to the machine (that is what has been called mixed initiative). Such dialogue is currently ensured by what we call a 'conversational agent', one of the developing uses, mainly in remote interactions.

In fields close to mathematics, another form of interaction has been offered: a step-by-step monitoring of the actions that have been realised especially in a complex calculation or the transformation of an expression. To use a medical metaphor, the interaction develops according to a cycle of diagnosis and intervention. Diagnosis is based on the collected data and techniques to analyse them. Intervention is based on teaching objectives and the techniques needed to reach them, taking into account the diagnosis. In other words, a model of the student and a model of the teaching are needed.

The building of the student models has prompted a lot of research. There are three main approaches which lead to different forms of intervention.

The first corresponds to what is called partial or overlay expertise (Carr and Goldstein, 1977). The knowledge of the student is considered as a subset of that which is aimed at. That model is quite 'handy' and compatible with fragmented approaches about skill acquisition: lists of skills are provided and intervention processes correspond to the discovery or in-depth learning of non-acquired skills (or non-acquired knowledge), which can often be worked upon independently.

The second corresponds to the differential models that integrate 'false knowledge', which corresponds to disturbances of expert knowledge. That is the case in mathematics, where a lot of students will abusively transform the square of a sum into a sum of the squares – $(x + y)^2 = x^2 + y^2$ – by generalising what works for multiplication – 2(a+b) = 2a + 2b. Similar deviations are to be found in the teaching of a foreign language, when constructions from the mother tongue which are not correct in the target language are used. Detecting erroneous knowledge is important and means must be found to establish a good diagnosis. Then, fast intervention is needed to correct mistakes, in order to prevent their stabilisation. The best is to design a teaching model which avoids their appearance, by providing knowledge at the right moment (what Van Lehn calls the 'felicity' conditions). Anyway, didactic studies are needed, to spot the mistakes as well as to eradicate them.

The last approach is that of misconceptions, which is now well-established in science learning. There are many children, but also adults, who have ideas that are completely erroneous, for example, that an electric current is like the flowing of water inside pipes. Everyday life perception of physical phenomena is often in contradiction with physics theories. Misconceptions may be an obstacle to learning and it is important to know them. That is far from skill-based approaches. As in the previous case, it is important to conduct didactic studies to spot erroneous conceptions but it is not easy to figure them out.

As one can imagine, the conception and maintenance of such models is no easy task, whether it relies on a subset of a reference model (partial expertise) or on a deviation of a reference model (differential model) or on a misconception far from expert or reference models. They have led to the realisation of a few systems, IT tutors providing individualised teaching which 'worked' in terms of learning gains for students.

(1) A few exemplary achievements in formalised domains

A much quoted example is the general programme *Cognitive Tutor* of Carnegie Mellon University, a teaching system which supports learning based on practise, used in different domains, especially in algebra (Koedinger *et al*., 1987; Ritter *et al.,* 2007).

That course, created at the beginning of the 1990s, has been continuously tested in a great number of classes, from 75 schools in 1998-1999 to more than 1,400 in 2003. It was selected in 2004 by the American Department of Education for its list of technologies that 'work'.[2] The software has been sold to many schools, with researchers collecting additional data and constantly refining their programme. Thanks to that continuous experimentation, according to Koedinger, research has allowed to acquire extensive knowledge on teaching and learning processes, on what works and what does not, to help students master algebra correctly.

In that case, the achievement results from the long-term collection and analysis of student works, in a discipline (mathematics and algebra in particular) where 'correct' solutions are easy to identify and the pedagogical objective clearly (and strictly) defined.

While the algebra tutor is an example of what has been technologically possible and implemented in schools for about three decades, another field of application of AI in education is IT tutors which are capable of conducting a 'mixed initiative' dialogue with students or pupils, i.e., asking them questions and answering theirs. Van Lehn (2011) has done a survey of studies comparing the efficiency of human tutorship, IT tutorship and

2 What Works Clearinghouse

no tutorship. He has concluded that intelligent tutorship systems were almost as efficient as human tutorship.

AutoTutor is an example of that approach. It was developed by a research group on tutorship at Memphis University, initially to support Newtonian qualitative physics and IT teaching. Its conception draws on constructivist theories based on explanation, within learning, of intelligent tutorship systems which adaptively answer to the knowledge of students, and on research on models of dialogue in tutorial discourse. AutoTutor simulates the discourse models and pedagogical strategies of a typical human tutor (Graesser *et al.,* 2001) by providing information feedback, asking for more information, giving advice, and identifying and correcting erroneous ideas (Graesser *et al.,* 2004). The experiments that have been conducted show that AutoTutor can result in learning gains in multiple fields (for example, IT culture, physics and critical thinking) (Nye *et al.,* 2014). An extension of AutoTutor has expanded learning with a tutor by adding a third entity, a student agent for a three-person conversation (Graesser, 2016). More than a dozen systems have been developed based on the original AutoTutor.

Though those are promising examples, they are tightly linked to the traditional teaching of disciplines and the individual acquisition of knowledge. They are hardly linked to more participative and collaborative forms of learning.

(2) Is building intelligent artificial tutors the right direction?

To realise fully adaptive tutors, knowledge must be implanted into machines for the three components interacting in the training – the discipline, the student and the teacher – answering the following questions: what, to whom, and how. That leads to articulating three models: the domain model (to be taught), that of the student, and that of the teaching. Research in AI for education, in the 1970s and 1980s, showed how complex managing teaching and learning processes are (Bruillard, 1997), revealing serious obstacles to machine-driven management.

Indeed, a first limit is the interaction model – a student facing a machine alone – which does not correspond to collective teaching as it most often happens in classes. The teaching model – tutorship – has prevailed, even though some systems have simulated the existence of co-learners (as in some extensions of AutoTutor). Trying to better understand what a teacher does in a class has revealed the great number of decisions they have to take in action, and the great difficulty in modelling them. Thus, the intelligent systems which have endured have mainly been limited to individualised learning.

Beyond important designing difficulties in the automated management of learning, the most serious obstacle is a purely pragmatic one. In order to be able to teach, machines must be able to carry out the tasks they are supposed to teach or to solve the problems they pose, in a fairly similar way to what is expected of human beings. In less formalised fields than mathematics or IT, this is really difficult to achieve, and if machines perfectly know how to perform those tasks, what interest do human beings have in learning how to master them too? Why learn what machines perfectly do? Shouldn't one learn how to do, not alone, but with machines, how to develop computationally instrumentalised activities and interact to realise complex tasks?

A last point is about student models. As we have already described, in most fields, students are not empty receptacles that should be filled by education. They already have some knowledge, which is correct or erroneous, and sometimes fragile, and knowing what they know or think they know is important to ensure effective learning. Some mistakes are very common and studies may detect their origin, based on the features of the knowledge in question. They may be spotted in productions. AI techniques may automatically do that, but they cannot explain them without detailed analysis, and even less overcome them.

Thus research has shown the complexity of elaborating different models (of the field, the student and the teaching) and even more of articulating them. In fact, the promises of AI have dominated research on human being and IT, but it was the studies on hypertext and human-machine interaction which shaped the IT environment at the end of the 20th century. Pioneers like Douglas Engelbart and Alan Kay imagined machines that could extend human capacities, helping them in all sorts of activities (working, playing, etc.), which have now become common, especially for the young generations of developed countries.

The debate is still quite topical: should machines be made intelligent or should they be made to help human beings and be adapted to the ways of thinking and working of the latter, helping them to be more intelligent thanks to machines?

Should machines be designed so that they will automatically solve issues or provide human beings with documents and data in order for them to read and interpret them and decide what should be done? There is a similar debate about personalisation and individualised learning.

2 Personalisation – two opposite visions, but that of control always prevails!

For some years, the key word in education has been that of personalisation and AI seems to be the right technology to best ensure such personalisation. That word however hides very different interpretations, or even strong oppositions, which we are now going to present.

(1) Personalisation – *a self-evident fact* used by companies but hard to define

Many people today praise the personalisation of learning and teaching. First, the new expectations of the youth are to be taken into consideration – it is a known fact that they are bored at school, that activities and feedback should be offered to them to support them, and that they would love using so-called ‘modern’ technologies. However, on that last point, even though they are strong daily users of digital technologies, studies have shown that they fairly seldom use them to learn (see part I).

Second, educative authorities state that it is necessary to put the student at the centre, to implement a differentiated instruction, and to take into account the needs of each student seen as a singular individual. Indeed, all the children are different (disabilities, cognitive styles, learning styles, interests, aims, etc.), but how could teachers implement enough differentiation?

In order to help them, many small and big EdTech companies, as may be seen on websites, offer their help.

According to the Aurora Institute,[3] 'identify students' unique needs and address them [...] It is about optimizing learning every day and maximizing the amount of learning per unit of time', or according to Century, 'our IA technology understands how an individual learns best, and constantly adapts to provide the support or the challenge each student needs.'[4] We will come back to those statements which are exaggerated, to say the least, by trying to define the notion of personalisation.

What is personalised learning? There is not one but many definitions: as is specified in *The Glossary of Education Reform*, online resources were created by the Great Schools Partnership, 'Because personalized learning has such broad implications, and the term encompasses such a wide variety of potential programs and strategies, it may be difficult to determine precisely what the term is referring to when it is used without qualification, specific examples, or additional explanation.'[5]

Not only are there multiple definitions,[6] but uses in class are also very diverse. However, Larry Cuban says that 'wherever these classrooms, programs, schools, and districts fall on the continuum of personalized learning with their playlists, self-assessment software, and tailored lessons all of them work within the traditional age-graded school structure. No public school in Silicon Valley that I visited departed from that century-old school organization.'[7]

In any case, it is possible to list the multiple forms of personalisation, linked to the learning objectives, content, modalities, rhythms and places. The question is also who has to adapt – is it the learner or the learning system? Should a learner be provided with the adapted resources depending on their individual characteristics, objective and needs, or should they have the possibility to choose the resources that they think are the most relevant for their own learning?

(2) Two opposite visions of personalisation

According to Justin Reich,[8] everybody seems to agree on the fact that learning should be personalised, that learning experiments should be adapted to each student, and that personalisation is made possible by the new technologies. However, the meaning of such personalisation is different. Is it

- 'Using technology to give an individual diagnosis of the students' skills on standardised tests and then apply algorithms to adaptively provide stimulating content appropriate to each student in order to help them better 'pass these tests'?
- Or is it giving access 'to a world of information and expertise to each student and give them the power to explore and create, and allow them to follow their interests and passions in different directions'?

There are two possibilities: either adopt the industrial model of education and give each child an assembly line or destroy the manufacture and build something else (maybe creation agencies). Reich quotes

3 https://www.inacol.org/news/what-is-personalized-learning/ (last access in May 2024. Hereafter, the same date if not specified).
4 https://www.century.tech/the-platform/ (last access in January 2021)
5 https://www.edglossary.org/personalized-learning/
6 Benjamin Herold (2019). What Is Personalized Learning? https://www.edweek.org/technology/what-is-personalized-learning/2019/11
7 https://larrycuban.wordpress.com/2017/03/22/a-continuum-on-personalized-learning-first-draft/
8 https://www.edweek.org/education/opinion-battling-over-the-meaning-of-personalization/2012/06

the analyses of historian Ellen Lagemann who sees the history of education in the 20th century as a fight between Thorndike and Dewey, which Dewey lost. Indeed, while the latter had some influence in many scientific fields, it is the thought of Thorndike which had the most impact on the field of education and which contributed to shape the practise of public school (in the US). For Gibboney (2006), Thorndike saw human beings like machines while Dewey saw them like life. Thorndike supported educational science based on objective measure while Dewey wanted school to look like life, even though the results were difficult to measure. Reich concludes that it would be necessary to understand why Thorndike's vision has prevailed.

Why does control prevail over emancipation? One may see a combination of reasons here.

First, the growing industrialisation of education reinforces management methods which rely on indicators. It must of course be possible to calculate the values of those indicators, which implies, or is facilitated by, small independent tasks, small steps (as in programmed teaching), corresponding to reference frames. Next, the fact that activities are *evidence-based* leads to preferring tried and tested practices, (the *good* practices), which are easy to identify. That results in a pedagogy of mastery or of the illusion of mastery, which is deployed on paths which have been meticulously marked out. In addition, the fragmentation of tasks favours the development and provision of pedagogical resources that are tailored to those specific tasks, which may be a significant support for beginning teachers or teachers who have to teach fields they do not quite master yet. For education managers, all that allows visibility and the illusion they are managing the activity of teachers and students as closely as possible. To sum up, it is that industrialisation process of education (Moeglin, 2010; 2016) which is often presented as unavoidable and which reinforces everything that is linked to control.

There is that same opposition in eportfolios, between those which are called narrative and those which rely on databases. A learning eportfolio corresponds to a finalised and reasoned collection of documents which show the quality and progress of the work of a student through some of their productions. It integrates genetic aspects (showing the evolution over time) and reflexive aspects (showing the capacity to take a critical look at what has been done). However, structuring them as databases to satisfy the needs for uniformity in the assessment data of an institution risks erasing expression creativity, which has been the standard of portfolios for years. That is a change from 'authentic and reflexive' assessment to the aggregation of data for the accreditation of two competing processes, the first one being at the level of activities and projects and the other at the level of actions and tests (lists of basic skills). An artist who designs their portfolio as a list of 'skill evidence' and not as a set of productions expressing a vision and approaches, risks finding it difficult to convince they have artistic 'value'.

In any case, the companies in a way ride the wave of the two opposite visions that have just been discussed. Thus, the *Aurora Institute* suggests the following definition of personalisation: 'Tailoring learning for each student's strengths, needs and interests – including enabling student voice and choice in what, how, when and where they learn – to provide flexibility and supports to ensure mastery of the highest standards possible'.[9]

9 https://aurora-institute.org/blog/what-is-personalized-learning/

However, that translates into 'utilize real-time data for feedback to intervene exactly where each student needs it most.' In that technicist vision of personalisation, is it possible to really identify the unique needs of students and cater to them?

Changing from a teacher-centred class to a student-centred one, and optimally to a class that is in a way 'led' by the student, where students have the possibility to choose their rhythm, their tools, their learning objectives depending on their interests, is far from immediate. Other less technological means allow to develop personalised learning, for example in project-driven teaching.[10]

Machine-based learning is individualised, because, for the moment, machines cannot do anything else! In addition, a lot of data must be collected. Indeed, to have enough control, there must be some sort of fuel to make the AI engines work. Thus, a sort of vicious circle is created: in order for AI to work, there must be a lot of data, therefore machines, and regular and sustained use, but which in a way structure the students' experience. Personalisation may become standardisation.

Focus: AI – fantasies and approximations!

While discourses on AI proliferate, few of them are informed and balanced. To give an example, let us take recent media articles on the web. In August 2019, Siècle digital published an article entitled 'Denmark: AI is used to track students' behaviour and performance'.[11] That article is based on a short story broadcast by VB (Venturebeat) entitled 'Researchers use AI to track students' performance in online courses', [12] published on the day before. The latter quotes the source,[13] a scientific article from the EDM (Educational Data Mining) symposium of July 2018 (Lorenzen, Hjuler, Alstrup, 2018). In that article, there is no mention whatsoever of AI. It is an analysis of log files from an online application called Clio Online, used by Danish pupils, doing clustering using usual mathematical techniques (Markov chains). When comparing what is said in the French and English Internet articles, one sees the gap that separates them from the scientific article which they claim is their reference. Almost everything that is said is approximate or even false. Thus, there is nothing about 'tracking' pupils, but about studying recorded files *a posteriori*. According to the French article, 'researchers have examined the data of 14,810 Danish pupils and students.' While the figure, 14,810, is correct, it comes from Clio Online, the biggest provider of digital content for primary schools in Denmark. They are not students. 'The results show for example that the pupils who work on scientific projects spend a lot of their time reading. Those who learning foreign languages are usually rewarded with a good mark only if they work intensively.' But what the scientific article shows is that 'participation in quizzes seems to increase pupils' performance in languages more than in other subjects, where reading texts is more important.' Last, in order to avoid tedious comparisons, let us simply quote the conclusion: 'In a close future, they hope they will deploy a system which could be used to track changes at class scale over time, in order for teachers to adapt their lesson plan to each pupil.' However, as we have seen, there is no real-time tracking, and one can only hope

10 See for example https://siecledigital.fr/2019/08/27/danemark-lia-est-utilisee-pour-suivre-le-comportement-et-les-performances-des-eleves/

11 See for example https://www.gettingsmart.com/2018/03/personalized-learning-experiences-why-and-how/

12 https://venturebeat.com/2019/08/26/researchers-use-ai-to-track-student-performance-in-online-courses/

13 https://arxiv.org/pdf/1908.08937.pdf

to: 'Help teachers to encourage an optimal behaviour of pupils, for example by recommending training quizzes for pupils studying languages, or by making sure that pupils have more time to use the system at school.' It is difficult to find that many errors in such a short text! There are fewer of them than in the English article on the Internet, which shows degraded quotation.

3 Big data in education, what can collected data reveal and to whom?

Teachers and trainers should agree: it is important to better know the other in order to teach and learn better. Thus, it is useful to know what students or learners have done before, their history, what they know, what they can do, their environment, their constraints, their aim, their objectives... If the needed information is not available, in an in-person class, it is possible, before introducing a new notion, to ask them whether they have already come across it, what they know etc.. The idea is to ensure a sort of co-adaptation and to establish a trusting relation between teacher and students.

In today's vocabulary, that will translate into collecting data on learners. Having data is better than having none. But if one already has the data, is it better to have more of them, is it necessary, useful?

In a naive technical vision, one may think that the more data the better. However, that is far from self-evident and the popular expression 'too many data kills the data' could apply.

Indeed, collecting data is not a neutral process, it may change the pedagogical relation. It can never be complete – thinking it is complete is an illusion, and prevents thinking about the nature of the partial data that have been collected: what is revealed, what remains hidden? The very objective of the collection must be questioned: is it control, full transparency (panopticon) that is aimed at? Then several capacities are at stake, those of data processing, those of the interpretation (of the data and the results of the processing) and those of the intervention.

We are now going to specify what *big data* means for us, and give a few examples of what their collection and processing may yield in education.

(1) *Big data* need specific infrastructures

What are *big data*? They are the flow of data which have three main characteristics: they are really bulky, very diverse and real-time. The expression used is 3Vs: *Volume*, for how big the data are; *Velocity* for on-the-fly speed, collection and processing; *Variety*, for structured, or not structured, data.

The progressive addition of other characteristics has resulted in new words being added which all start with the letter V. There are the 5Vs, including *Visibility*, for easy-to-access-and-use data, and *Veracity* (data quality): visible data allowing to conduct good analyses (taking the right decisions). Then the 7 Vs, with *Variability* (context), differences in nature or judgement and *Value* (utility): changing data, interesting processing.

There are also the 10Vs with *Volatility* (freshness and preservation), *Vulnerability* and *Validity*.[14]

Strangely enough, on an American website,[15] the characteristics of the 10Vs are different, with *Venue* (where data come from, their owners, and access constraints), *Vocabulary* (data models, ontologies, taxonomies…), *Vagueness* (confusion even in the field of *big data*!).

In 2017, there was even the famous 42, a mythical number,[16] with a list of the *42 V's of Big Data and Data Science*.[17] A figure, from the same page, gives the list of the Vs used (in English) and the date of their first appearance. There are 15 of them, due to the different interpretations of the 10Vs depending on the commentators!

In any case, managing those *big data* necessitates big storage space, interfaces for real-time collection, organisation, processing and data visualisation software. No classic tool for the management of databases of information management can really work on those very large set of data, which has led to the creation of open-source specific infrastructures such as *Hadoop* or *Spark*.

It seems that the origin of reflections dates back to 2005, when people became aware of the quantity of data that users generated on Facebook, YouTube and other online services.[18] It is marketing that is at the origin of orientations and directs them. Is it really about education or behaviours that big companies are trying to regulate?

(2) Collecting many data in education – processes that are hardly convincing

If it is possible to collect a lot of data, what does it bring to education? We are going to discuss that possible input with several examples.

a. Passing a MOOC – clicking enough times?

The first example is linked to MOOCs, those massive open online courses, which have been briefly presented in the first part of this book. When they were launched in the media in 2011, they were presented as supports for a major evolution, or even a revolution. Massive teaching, of course, but with a promise – that of personalised massive teaching. Indeed, by recruiting dozens of thousands, or even hundreds of thousands of participants, who would leave big quantities of automatically collected data, the processing of those data would allow constant improvement of classes for each new edition, especially thanks to knowledge of mistakes and remediation strategies which would be efficient.

Ten years on, the results should be impressive. Justin Reich, who had access to the data of the MOOC of the MIT and Harvard (platform EdX), deducted from it major laws on learning, with a sort of irony. Roughly speaking, the more things a student or a pupil does, the better their chance at passing. Of course, it is a statistical 'law', since some pupils work a lot without success, while others succeed without working a lot.

14 https://le-datascientist.fr/les-10-v-du-big-data

15 https://mapr.com/blog/top-10-big-data-challenges-serious-look-10

16 42 is the answer to *The Ultimate Question of Life, the Universe, and Everything* in Douglas Adams's work entitled *The Hitchhiker's Guide to the Galaxy*, without what the question was being precisely known.

17 https://www.kdnuggets.com/2017/04/42-vs-big-data-data-science.html

18 https://www.oracle.com/fr/big-data/what-is-big-data.html

In MOOCs, that law (tested with the quantities of retrieved data), can be simply translated into: the more a student clicks the better their chance of passing.

In fact, a first difficulty with the large amounts of data from MOOCs is that the most tested and tried and used methods of data analysis, which are based on tests and regressions, are saturated by the volume of data. A second difficulty lies in the nature of the retrieved data: they are first and foremost behaviour data, which are linked to the interaction with numerical systems, much less learning data. Besides, there is a lack of interpretation models. A Mexican colleague (Tech de Monterey) thus reported the data analysis they conducted on the MOOCs launched in Mexico. They discovered that learners who eat chicken achieve better results! Behind such a result, which is quite difficult to interpret, there may be a hidden variable (a socio-economic one) or associated elements which are still to be clarified. In any case, that correlation is not a causal relationship, and one should not encourage pupils to eat chicken to improve their school results.

The very notion of achievement is different depending on the point of view. For a learner, everything depends on their real or supposed objectives. That may be getting a certification, or simply reaching the end of the course, or learning a few elements they are interested in. Those in particular who teach the topic of the MOOC may be quite happy to learn specific skills, to perfect their knowledge of elements they know little, etc. The teachers and designers of the MOOC are more interested in knowing what the participants have been able to learn. As to the platform managers, their achievement or its measure is first and foremost linked to the persistence of participants – the more they connect, the better.

We here confirm there is some uncertainty as to the objectives: is it about helping people learn or controlling their behaviour? Marketing or education? We will come back to that.

In his PhD on MOOCs, Cisel (2016) has managed to extract a very good indicator of MOOC training completion, by using the marketing concepts of entry key and registration behaviour. If someone registers for a MOOC by looking for a specific product (entry key product), they are more likely to finish it than if they have chosen it when surfing on a MOOC platform (entry key platform).

b. Better mastering of word processing and spreadsheets?

Having a lot of data is one thing but what is the aim of their processing.

The Microsoft company has hired ethnologists to study how people use their products (especially during air travel). They have also retrieved massive data from users from all over the world for many years. When one installs a product, the person in charge of that installation may simply tick a box allowing to upload user data. With that massive amount of information, one could think that the company could get quite an in-depth knowledge of the users' behaviours and find the keys allowing them to be more efficient in their uses, especially by identifying the concepts that have not been well understood or mastered, which renders their use far from optimal.

However, that is not what can be seen, far from it, and many users lose a lot of time because they do not understand how some things work. What has Microsoft done with all those data? Improved word processing and spreadsheet skills, or kept users in routines which are not always efficient but give the impression of something intuitive?

c. Who needs big amounts of data?

In the above examples, one can see that having massive amounts of data is far from enough to improve education, either because their very nature does not allow the wanted processing, or because the very aim of the processing is not directly oriented towards learning improvement.

At the national or international level, there are different campaigns of test-based data, like PISA (International programme for the monitoring of pupils' knowledge), TIMMS (Trends in Mathematics and Science Study) or NAEP (Programme of progress assessment in the American education system) in the United States. Though many data are collected in respect of a country, that is far from what is called *big data*. Well-mastered data collection and statistical processing by agencies yield fairly reliable results, which give information on the main trends of the education system. Simplified forms are also broadcast and feed national debates on achievements and difficulties encountered by the education of the people. Thus, one of the main results of the 2018 PISA inquiry was to confirm the very unequal feature of French education. 'France is the country where the socio-economic background best explains progress in scores.' [19] Data analysis, sometimes with AI approaches, may be very useful to find the characteristics of children who may become disengaged in order to implement prevention policies.

That type of inquiry does not necessarily give concrete elements to teachers in class, except if they are provided with information they can use. Researchers working on AI and learning are not always interested in big amounts of collected data, but rather by regular targeted collection in classes, like what has been done around the *Cognitive Tutor*.

A contrario, as we have seen, the companies which work on 'digital change' will more or less say the same thing. According to them, so-called adaptive learning is going to give a new impetus to education thanks to the volume of collected data and AI – it is going to be possible to precisely track progress and better apprehend individual needs: 'AI to personalise education depending on everyone's needs.' But 'That school should have digital tools that will allow it to identify the needs for support of each student' [20] (According to France, SCC).[21]

Indeed, the implementation of data collection and its analysis imply that students work on connected digital systems – thus, the necessity for companies to push school into that direction. An *industrial* vision of education leads them to offer solutions which can apply everywhere and the students' adaptation is always carried out within an imposed frame. Consequently, 'that personalisation is quite strangely even more uniformed than traditional teaching'. Indeed, there is significant fragmentation in education (teaching practices, disciplines, content, methods and student typologies, with language and cultural variants), which slows down automation (Olivier Ezratty).[22] But to develop personalised education, scale must be taken into

19 https://www.lemonde.fr/societe/article/2019/12/03/pisa-2018-les-eleves-francais-legerement-au-dessus-de-la-moyenne-de-l-ocde-dans-un-systeme-toujours-tres-inegalitaire_6021440_3224.html

20 https://france.scc.com/news/lia-pour-personnaliser-leducation-en-fonction-des-besoins-de-chacun/

21 https://france.scc.com/a-propos/

22 http://www.magrh.reconquete-rh.org/index.php/articles/formation/455-les-applications-de-l-intelligence-artificielle-dans-l-education

account. The best way to do so is normalisation, and the latter is the enemy of personalisation (Peter Greene).[23]

Who can then wish for personalised education? Isn't it subject to commercial logics (Alexandre Roberge)?[24]

4 What does AI bring to education actors?

How can those technologies work in class? Using digital personalisation programmes necessitates working a good part of the time alone with a computer, and, according to companies, the longer a student spends on the platform the better AI manages to refine its propositions based on what it knows about them. All that time on machines may be problematic. But AI may also, as is often said, help teachers diversify the activities with their students.

(1) Debated uses in class

Indeed, some implementations in schools spark off debates. Thus, an article from the *New York Times* on the use of a Facebook programme named *Summit*, is entitled 'Silicon Valley Came to Kansas Schools. That Started a Rebellion':[25] headaches, hand cramps, anxiety. 'We allow computers to teach and our children all look like zombies'. More than three quarters of the persons interviewed in an inquiry have declared they preferred their children not to be in a class that uses the *Summit* programme.

A more detailed report (Boninger, Molnar and Saldaña, 2020) confirms those concerns and explains that significant funding (close to 200 million dollars from the Chan-Zuckerberg Initiative, the Gates Foundation, and others) has led to one of the most prominent personalised digital programmes in the US. 'Its rapid spread, despite a lack of transparency and no convincing evidence that it may keep its promises, provides a powerful example of the way political decision-makers are challenged when they are confronted with a well-funded and self-serving pressure for schools to adopt digital programmes of personalised learning.'

Some schools use 'Playlists' curricula, which are some sort of reading lists, given to each pupil each morning (the list of activities they have to perform on their computer) in order to personalise the pupils' learning, but that technology is neither cheap nor tried and tested.[26]

The question really is that of the role of the teacher and of the control they are capable of exercising, or not: what are the data flows which reach them and the decisions they delegate to a system, or not?[27]

(2) Giving teachers control – the Villani report (2018)

The Villani report is clearly against approaches that put the teacher aside. It is not a question of yielding

23 Scaling Up Personalized Education, Peter Greene, https://www.forbes.com/sites/petergreene/2018/09/10/scaling-up-personalized-education/#127e0c3f735e

24 https://cursus.edu/articles/42761/qui-veut-dune-education-personnalisee#.XYDSq2bgqUk

25 https://www.nytimes.com/2019/04/21/technology/silicon-valley-kansas-schools.html

26 http://www.edweek.org/ew/articles/2017/03/29/curriculum-playlists-a-take-on-personalized-learning.html

27 Ed Week, Michelle R. Davis, November 5, 2019 : https://www.edweek.org/technology/q-a-the-promise-and-pitfalls-of-artificial-intelligence-and-personalized-learning/2019/11

to the control of some AI and entrusting it with decisions, but of developing what is called an ‘empowering complementarity with AI by reinforcing the place of creativity in teaching’.

To this end, it is important to promote the teachers’ mastering of dashboards presenting the data processing of their students and the student’s mastering of their own learning data. It is also important to ensure that AI is not used for surveillance or increased optimisation of performance purposes, but to increase the teachers’ power to act in their teaching freedom and dialogue with learners.

More concretely, in order to facilitate the exploitation of learning data by actors (learners and teaching team), the Villani report suggests identifying the relevant data, facilitating access to them and their enrichment, while respecting the rights and interests of learners. Let us note that, according to a study by Barabara Means (2010) on the use of technologies in class, it is the capacity to use the reports on the data generated by software which makes a difference, in terms of learning gains and efficient class management.

However, if teachers are expected to use the students’ data to improve the efficiency of their practice, they will have to be helped to do so. That issue is not tackled in the training programmes (Means, 2010), and what seems to be true in the USA is also true in France. The nature of the teachers’ skills and difficulties in terms of data use should be understood to provide the training and support they need. Giving teachers the means to use the data appropriately and ethically is a responsibility which must be borne by all those who prepare and support teachers and future teachers.[28]

Indeed, personalisation in education is first and foremost local, it is an *end-of-the-chain* adaptation, made by teachers. The question is how to help teachers to do so. It is possible to think that the (locally) processed data allow the teachers to identify the strengths and weaknesses of a group of students, the parts of a lesson that have been well assimilated, or not, and their level of attention, in particular for the integration of students with disabilities and disorders. Are there enough data in education?

Help them save time? A study by McKinsey suggests finding means to reduce preparation time.[29] In particular, AI could give them access to educational resources and suggestions of activities, adapted to their students. We will see in the next chapter on teachers that it is not that simple. If systems of recommendations are put in place, commercial interests may disturb the system. Then, if teachers do not have the possibility to adapt the resources themselves, they may not appropriate them enough.

However, the teachers’ working environment and the circulation of data on students should be thought anew for teachers to take advantage of it.

(3) Explainability – an essential requirement for educational resources

Health is one of the flagship domains of AI applications. AI is a bearer of improvement promises, especially about the automated analysis of photos and pictures to spot the signs of pathologies (see for example the website of INSERM[30] which provides an interesting overview). However, for them to be acceptable, or

28 https://datafordecisions.wested.org/wp-content/uploads/2016/08/2016_Teachers-Learning-How-to-Use-Data.pdf
29 https://www.mckinsey.com/industries/social-sector/our-insights/how-artificial-intelligence-will-impact-k-12-teachers
30 https://www.inserm.fr/information-en-sante/dossiers-information/intelligence-artificielle-et-sante

even set aside for not being relevant, the solutions suggested by the machine or the algorithm must be understood and, for that, there must be access to the machine's 'reasoning' or the machine must be able to provide explanations. That is especially necessary to allow the doctor to discuss with their patient and give them the possible alternatives.

The issue of the explainability of 'AIs'' proposals is an increasingly strong requirement. However, while digital approaches have been skyrocketing for dozens of years, what is called *deep learning*, i.e., networks organised in different layers, exchanging with one another, learning in a more or less supervised way, in most cases the 'machine' cannot justify its decisions: nobody knows what the algorithm does. How is it possible then to be responsible for a medical decision?

There are many publications on it, which first shows the mistakes that have been made (Heaven, 2019), the biases that have been noticed, the prejudice of an era, those of the designers, the overrepresentation of a category of persons, etc. Standard deep learning is not enough and researchers think classical approaches of AI and neuron network ones should be combined, for example by integrating causal models (Vasudevan *et al.,* 2021).

In any case, that requirement of explainability has led to a line of important inquiries, which are difficult to conduct, since beyond the strictly scientific aspects, identifying the 'right' level of explainability for a given situation depends on technical, legal and financial considerations (Beaudouin *et al.,* 2020).

For training purposes, it is central. The already old example of Mycin is characteristic. The Mycin system, created at the beginning of the 1970s, is an expert system based on rules, used to diagnose blood-clotting related illnesses. The use of Guidon (Clancey, 1983), a teaching system based on Mycin, has allowed to show that expert systems, which are above all designed for their expertise capacities, are not *a priori* good teaching systems. That is not linked to shortcomings in the teaching strategies, nor even to the very types of knowledge that are impacted in them, which are enough for experts but not for beginners. The latter need models (of a causal type) which explain the different rules of expertise that have been acquired through experience. Indeed, the compilation of the expert's strategy leads to an extremely condense reasoning, hardly understandable by a student, who thus has difficulties making sense out of those very heterogeneous rules.[31]

That shows how necessary explainability is and that black boxes are not adapted to teaching. It is also what Rosé *et al.* claim (2019), when they see that models are increasingly complex, often involving thousands or more parameters. With 'black-box' systems, it is not possible to look inside to try and understand how, why or whether those systems will work when they are applied beyond the data on which they have been built. Thus, explicative learning models, which, in addition to precise predictions, provide interpretable and usable information, should be developed.

However, the question can legitimately be asked whether some new trends in training are not going in the opposite direction of that requirement of explainability – the objective is not to lead the learners to discover and understand new things, but rather to adopt new behaviours, even though they may not be clearly aware of that.

31 http://tecfaetu.unige.ch/staf/staf-d/joye/staf11/IA/clancey.html. Joints quelques écrans (Boundless.ai et Century.tech)

(4) Is changing behaviours teaching's purpose?

'Imagine that I could make you do what I wanted you to do without you realising that I was even involved. All I would have to do is to rearrange the information around you in ways I know would lead you in the direction I desired. I could change the sequence of the choices you have to make, and use my knowledge of your susceptibilities and weaknesses to choose the appropriate time and method of delivering my *nudge*.' (Sætra, 2019).

That quotation comes from the introduction to an article entitled '*When nudge comes to shove: Liberty and nudging in the era of big data*'. The nudge theory, developed by behavioural economist Richard Thaler and jurist Cass Sunstein,[32] refers to a method of influence trying to change human behaviour, especially as to decisions or choices to make, without constraint, obligation or sanction. According to Ambrosino *et al.* (2018), that theory is 'based on behavioural economics and the idea that people have a limited rationality, that they often do not have very defined preferences and that they are under a certain number of prejudices, which lead them to make choices that are not in their own interest, and even sometimes against their will.'

This is very close to so-called persuasive technologies (Fogg, 2009), conceived 'to change the users' attitudes or behaviours, through persuasion and social influence, but not coercion', presented in the first part of this book.

An increasing number of scientific publications analyse that notion of nudge, especially from the perspective of ethics. Indeed, *nudging* is more and more efficient, for three reasons: (1) data more numerous about persons, which makes it possible to *nudge* them more efficiently; (2) increasingly sophisticated theories on human behaviour, targeting human vulnerabilities; (3) channels (social networks, online advertisement, geolocation, etc.) giving means to target each person individually. To use a military metaphor, nudge resembles surgical strikes more than the carpet bombings of the past, which makes it more efficient.

The notion of behaviour is at the heart of the new, neobehaviourist approach of education – new data on behaviour, and new ways of modifying or controlling it. Indeed, we have seen that many of the data automatically retrieved from platforms were behavioural ones.

Companies clearly aim to modify behaviours. In the previous chapter, we presented memory anchorage and its link with the adoption of new automatisms.[33] Similarly, the Boundless Minds company uses AI and neuroscience to shape, predict and analyse human behaviour at a neurological level. It was bought in 2019 by Thrive Global. One of the slogans of its presentation was: 'Brains can be programmed; you just need the code'. '*Boundless technology changes peoples' behaviors, beliefs and being.*'

> 'Modify your users' behaviour in a predicable way thanks to AI. We are changing the way human beings interact with their devices and we are proud of being leaders in the field of persuasive and behavioural technology.'[34]

As we have seen with MOOCs and Justin Reich's analyses, 'We have terabytes of data on what students

32 See https://fr.wikipedia.org/wiki/Nudge_(livre)

33 https://www.woonoz.com/blog/attention-ancrage-reflexe-formation/

34 La firme Boundless.ai a été rachetée par Thrive Global. On trouve encore https://www.linkedin.com/company/boundlessai/

click, but little understanding of what changes in their heads'. However, if the objective is to shape behaviours, that lack of understanding is not prohibitive.

The combination of principles from neuroscience (repetition, distributed learning) – which are compatible with modes of working made up of successions of short fragmented periods of time – with personal technology, i.e., the smartphone – which is close to the body, constantly at hand, communicating with a data-receiving platform that conducts analyses via algorithms (based, or not, on AI) and sends nudges back – may produce systems that can 'predict and shape human behaviour', at least partly. Users, confident with interfaces they think are intuitive, will become consenting and obeying partners.

In a book dissecting the links between the neuron networks and the neoliberal theories developed by Hayek, Pablo Jensen (2021) explains that their convergence leads to a '*manageable man that is highly governable* by market signals or other nudges'.

Is this taming or education? What are the objectives of education: simple modification or adaptation of behaviours or deeper form of emancipation?

5 Should platforms without intermediation but with significant funding be regulated?

The notion of platform is crucial in the current digital economy and education hardly escapes it. If we go back to the example of China described in the second part, we can have an idea of some characteristics of the AI market in education, which we are going to examine after a few ethical and political considerations.

(1) Debated ethical and political questions

Indeed, collecting and processing big amounts of data allows to feed predictive, but not explicative, statistical models: such configuration is likely to yield such result, without it being possible to provide reasons for such a link. As Antoinette Rouvroy underlines (2009, 2011), the myth of foreseeability may lead to new forms of government which may be quite oppressive under cover of transparency. Pending causal, explainable models, one may wonder what regulations to implement.

Ethical issues are mentioned a lot, especially because AI may lead to the loss of many jobs and to problematic forms of governance by machines. Thus, the Declaration of Montréal for a responsible development of AI pursues[35] three objectives:

1. Elaborating an ethical framework for the development and deployment of AI;
2. Orienting the digital transition so that everyone may benefit from that technological revolution;
3. Opening a space for national and international dialogue in order to collectively achieve an inclusive, equitable and environmentally sustainable development of AI.

35 https://www.declarationmontreal-iaresponsable.com/la-declaration

More specifically about the ethical aspect of AI in education, let us note that reflections are quite old already, and one of the first text about it (‘*Ethical Guidelines for AI in Education: Starting a Conversation*)[36] is dedicated by the authors (Aiken and Epstein) to a French researcher, Martial Vivet, one of the pioneers in that field. An institute for ethical AI has been created at Buckingham University.[37]

Thus, since digital machines and networks are now part of our lives, it is important to regulate their action. The issues of training people and regulating organisations are not opposed but complementary. Maybe the machines’ capacities for reading, writing and processing should be limited?

The risk behind uncontrolled implementations based on a lot of AI is to confine students to highly stressful working conditions, as the example of the use of the *Summit* programme in Kansas shows (Boninger, Molnar and Saldaña, 2020), constant competition, even at an early age, leading to depersonalised education.

Some in the USA, like the Brookings Institution,[38] see it as an unavoidable direction of education, referring to the country’s keeping its supremacy in the race to AI technology.[39] However, they admit that such a system may leave a lot of students ‘behind’, which raises deep moral issues and political leaders need to understand the underlying dynamics of the technologies in question. Others think that public interest should be protected by putting in place surveillance and responsibility mechanisms linked to digital platforms and personalised learning programmes (Boninger, Molnar and Saldaña, 2020).

(2) What AI in education?

When looking at the analysis of the developments of AI in China and what may be observed in the big digital companies, some trends appear. To conquer a big enough market, it is necessary to rely on a platform to attract customers, win their trust, create their loyalty, provide them with free resources and services, favour the constitution of communities (students, teachers, parents) and collect quantities of data. In China, the Wechat application (which is considered as a mode of living) encourages the link between communities. Then there is a double market, that of control (of schools by the State) and that of individual achievement (families). The latter is favoured by the competition among children and high-stake tests (importance of homework in China and Gaokao). Families and persons are ready to pay for coaching and personalised help, the efficiency of which partly depends on the processing of the large amount of collected data.

Focus: Some control measures of online education have been announced by the Chinese government (May 2021).

It is difficult to attribute the recent regulations to respect for ethical questions, but the Chinese government is intensifying its control measures and, sometimes, repression against its online education industry, forcing start-ups, which are often supported by big highly profitable groups like Tencent and Alibaba (though they may not always be very efficient from a teaching point of view) to put their multi-billion-dollar projects of quotation on

36 International Journal of Artificial Intelligence in Education (2000), 11, 163-176.
37 https://www.buckingham.ac.uk/research-the-institute-for-ethical-ai-in-education/
38 https://www.brookings.edu/about-us/
39 https://www.brookings.edu/research/why-we-need-to-rethink-education-in-the-artificial-intelligence-age/

hold. Officially, it is done to reduce the huge pressure lying on children through private lessons – 'mind-numbing online courses with uncertain benefits.'[40] For that, some measures should be taken – imposing time restrictions after school, forbidding online courses for under-6 children, limiting school fees, improving course quality (qualification of teachers) and controlling aggressive and dishonest advertisement.[41] Thus, several companies have supposedly hired the same actress to act as an English and mathematics teacher on their platforms. In one of those online videos, the actress sings the praises of a 33-hour long contract of direct lessons which only costs 8 dollars. She warns that the absence of lessons has consequences and states that 'parents themselves may be ruining their own children.'[42] That calls to mind the regulations of 10 years ago in South Korea (see part I). There are also concerns about the widening gap between the poor and the rich, i.e., those who can afford additional courses. The Chinese EdTech have an interest in education in rural areas through philanthropist projects. That allows them to test new technologies, which schools could not finance, and build trusting relations with future consumers and local governments.[43] According to a European report on education technology in China (Feijoo et al., 2021), the will of the latter to reduce the educational gap between the areas without infrastructures and urban zones could create economy of scale for affordable and technologically advanced innovations from which Europe could draw inspiration. While it is difficult to precisely predict what is going to develop in China and be exported, some dynamics have undoubtedly been created and will make offers for education. A foresight article on the use of platforms in Swedish schools (Hillman et al., 2020) debates the replacement of an American platform (which they call Poodle) by a Chinese one. The latter is supposedly well accepted, especially since it is well-adapted to support group work. Its economic model is not linked to advertisement, but to better course choices and better preparation of young people to being hired in partner companies. The risk is to replace a curriculum which is supported by democratically treated and negotiated values by a learning course created by commercial interests and algorithms.

In France, a quick assessment shows little interest in AI for school so far, or few data available (before the possible implementation of a 'data education hub'), some closure with the GDPR, and no training of teachers. Educational publishers seem not to be prepared and little interested, and private lessons, with a few exceptions like *Educlever*, need few data. It remains to be seen what effective ambition the AI strategy in education, which has been announced by the State within the strategy of accelerating 'education and the digital', will have.

To conclude, what challenges linked to AI in education may be retained? First, make good use of it for what it can offer to analyse the main trends and spot recurrent features behind the large amount of data, contributing to direct general policies. Second, help education actors improve the conditions of good learning respectful of persons.

To do so, avoid the temptation of control, but be aware of the ambiguous nature of the collection and processing of students' data – if there is little change for students, because the educational system does

40 https://fortune.com/2021/05/31/china-edtech-private-tutoring-ipo-delay-crackdown-student-overwork/
41 https://kr-asia.com/what-the-world-could-learn-from-chinas-edtech-crackdown
42 https://fortune.com/2021/05/31/china-edtech-private-tutoring-ipo-delay-crackdown-student-overwork/
43 https://thediplomat.com/2020/12/edtech-in-rural-china/

not manage to value them, this will only allow to control teachers. Avoid believing in the superpower of machines and their capacity to *understand* the needs of human beings. Thus, some mathematicians have obtained a strange result on automated learning – the capacity to be learnt may be unsure[44] in the sense that it cannot be proven or refuted by using the standard mathematical axioms, which gives a scientific limit to what a machine can do.

Following Paul Emerich France, who has worked for several years in the USA in a school using personalised teaching and has turned away from it, in order to reestablish equity and humanity in classes, would it be possible to develop humanised personalisation and not the dehumanised personalisation[45] that technologies ensure? Maybe one should not have to choose between Dewey and Thorndike, but better articulate them, build an open school which encourages risk taking and accepts mistakes, admits imperfections, but can offer activities instrumented by the digital technologies, and limited but efficient for training, acquiring necessary skills, etc.

44 https://www-nature-com.lama.univ-amu.fr/articles/s42256-018-0002-3

45 https://paulemerich.com/2020/07/06/three-tips-for-personalizing-in-a-pandemic/