Rose Luckin (University College London)

How can we make AI education a priority without using scare tactics?

Luckin, R. (2022). How can we make AI education a priority without using scare tactics? In AI, data science, and young people. Understanding computing education (Vol 3). Proceedings of the Raspberry Pi Foundation Research Seminars.

Available at: rpf.io/seminar-proceedings-vol-3-luckin

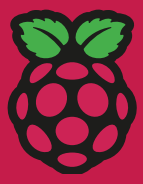

How can we make AI education a priority without using scare tactics?

Rose Luckin (University College London)

Introduction

I was listening to a Financial Times Tech Tonic podcast (Kynge, 2022) the other day, entitled *US–China Tech Race: brave new world* and it reminded me of an experience that I had during one of the last face-to-face events that I attended before the pandemic.

I was at AI Everything in Dubai: a conference and an exhibition. I had only just arrived and decided to start in the main conference hall to try and get a sense of the event. I listened to some of the keynote talks and found myself mesmerised with horror as I listened to the head of security for Huawei technology. He was telling a captivated audience how his company had developed amazing face recognition technology that was able to take the ageing process and its impact on the face into account. This meant that even if the AI technology had been trained with a picture of my 30-year-old face, when I hit my 50th birthday and beyond, the technology would still be able to analyse my face and recognise my identity. I had to admit that this was impressive technology. However, I have developed the habit of finding highly impressive technology developments that use AI disconcerting as well as exciting.

On this occasion, my unsettled feelings grew. The Huawei executive couched his narrative about the development of this technology in a story about the heartbreak of missing children and how this face recognition technology had been developed with the purpose of reuniting parents with children who had gone missing. The talk was accompanied by high-definition video footage of tearful parents being reunited with children whom they had not seen for many years. Their reunion was all thanks to Huawei's amazing face recognition technology. The technology had brought together some 169 missing children with their parents and was going on to help reunite more and more families. The speaker himself was dewy-eyed, as was the video footage. I was open mouthed, my jaw having dropped in disbelief that such a flimsy fig leaf of social responsibility could possibly be used to make a potentially deeply worrying technology seem like the best thing to have been invented this decade.

I turned around to look at the rest of the audience and their faces, assuming that they, like me, would be 'gobsmacked', but people just seemed to be taking it in, believing it, and seeing the good. Had I just become a cynical academic all too ready to criticise? I really didn't think I had, but I was left with a bad feeling.

So why was I reminded of this incident from 2019 while I was listening to the Tech Tonic podcast in 2022? And what does this have to do with the talk that I gave as part of the Raspberry Pi Foundation seminar series?

Let me answer the first question to start with. The Tech Tonic podcast was recounting a sad tale that has come to be known as The Countryman Case for reasons that will become apparent in just a moment. To cut a long (and fascinating) story short: in 2014, a young man, called Luka, was killed in a hit and run accident

on Serbia's Branko Bridge in Belgrade in the middle of the night. There were no witnesses. The police were not actively trying to solve the case, but Luka's father would not give up; he spent days standing on the bridge protesting and eventually the police started to investigate. Some grainy CCTV footage of the accident showed that the car that had killed Luka was a Mini Countryman, hence the name of the case. The police were unable to find the car, but they did manage to identify the driver. The driver could not be found anywhere in Serbia and they circulated his photograph to other countries and cities, including Beijing in China. Amazingly, less than three days later, the Chinese found the driver and the Chinese authorities immediately deported him back to Serbia.

The police in Beijing were able to find the criminal driver so quickly because of the advanced facial recognition technology used across the city. For several years now, China has been building a huge surveillance system with webs of cameras across the country, all of which have facial recognition technology, and that is why they could find the criminal driver in the Countryman Case so quickly.

Today, Branko Bridge in Belgrade is also monitored by cutting-edge Chinese surveillance cameras purchased from that same technology company whose representative I heard speaking at the AI Everything conference in 2019: Huawei. Of course, Huawei are not the only company who are making advanced facial recognition technologies, but they are one of the leading players in this field. It is also true that Serbia is not the only country that has made large purchases of these technologies from Chinese companies: according to the Tech Tonic podcast, 64 countries, from Africa to the Middle East and Europe have made such investments.

When countries invest in this facial recognition surveillance technology, they are enabling their police forces to stand a much better chance of capturing criminals, such as the man who killed Luka on the Branko Bridge, and they are probably enabling the speedier recovery of missing children and their abductors if they are still in the same country or in another country that uses this high-spec facial recognition technology. However, these countries are also equipping themselves either intentionally or unintentionally with the tools to enable them to become Big Brother style state surveillance entities.

Returning to the situation in Serbia for just a moment, it is good to hear that whilst the majority of the population believe that China is a beneficial trading partner for their country, there have been significant protests against the use of the facial recognition capacities of the cameras that have been purchased. As a result, the cameras installed across the country, including those overlooking the Branko Bridge, have the capacity to conduct facial recognition, but that capacity has not been turned on. One of the important learnings that we can take away from this story is that when people understand the implications of AI, both good and bad, then they can make more informed decisions about how they want it to be used. As AI becomes increasingly ubiquitous, the need for an educated population becomes even more important if democracy is to be upheld.

People and AI

The reason I started this chapter with these two recollections is that I am increasingly worried by the number of people who tell me that it is not important for teachers, parents, students, and the public in general to understand AI. I am told: "They just need to know how to use it". This is a dangerous situation, and it requires urgent attention. Firstly, I don't believe you can really know how to use AI without understanding something about what AI is and how it operates.

I do not mean that everybody needs to know how to programme and build an AI system, or that they need to understand the complex mathematics within a neural network. What I mean is that people need to understand what AI can do and what AI cannot do, along with the basics about how traditional AI, often referred to as Good Old-Fashioned AI (GOFAI), and modern machine learning AI operate in a non-technical way, as well as how they are different.

The case of the facial recognition technology being portrayed as purely beneficial in the stories at the start of this chapter offers an example of why people need to understand more about AI if they are to protect themselves and to appreciate the genuine risks, so that they are not driven away from beneficial AI by people who are scaremongering. People need to be able to tell truth from fiction and to make informed decisions. To make informed decisions, they need to be informed, or to put it in another way, they need to understand. Teachers need to understand because they can then help their students to understand. In the same way that teachers understand how to teach people to read and to write, they need to be able to teach people to be AI literate. And yet, it is very difficult to engage teachers in learning about AI. It is hard to persuade people who already have 101 things to do that this additional thing should be prioritised, particularly when it is not part of the curriculum or assessment framework.

I firmly believe that AI has a great deal of beneficial potential, way beyond the AI systems currently in use within the classrooms of the western world. I also believe that the vast majority of those companies who are selling AI technology into educational institutions are not posing risks like the facial recognition software I discussed at the start of the chapter. However, there are instances where systems that describe themselves as using AI provide scant information about the AI they use and how it delivers benefits. There are even examples where systems and companies that are described as using AI do not actually use any AI (for a range of views, see: Narayanan, 2019; Marr, 2018; Hao, 2019; Ram, 2019).

Education, educators, and AI

There are also a significant number of examples of AI being used in invasive and worrying ways in education. For example, students being monitored every minute of their day in China (Xie, 2019), CCTV being used to track down students not wearing masks in the USA (Keierleber, 2022), and classrooms in China using brain-wave trackers to check if a student is concentrating (Wall Street Journal, 2019).

Did the decision makers who brought these AI systems into their establishments know what the consequences, both positive and negative, would be? Were the educators fully involved in the decision? Is this just happening outside Europe, or is it the tip of the iceberg of a worrying global trend?

For those who doubt the sophistication of the AI being developed for education, it is useful to look at research labs to see what studies are being conducted and what ideas are being pursued. Such work will certainly inform the future and will likely foreshadow what will become commercially scaled. For example, the increase in availability and affordability of wearable and remote sensing technologies enables the study of groups of students working together by enabling the capture and analysis of voice, facial expression, speech, and bodily movements. Voice and facial expression data has been used for the analysis and categorisation of discussions using AI in the form of Natural Language Processing (NLP) and machine learning (Stewart et al., 2021). Posture detection and facial recognition has been used to classify the participation states of learners using Bayesian modelling (Kasparova et al., 2020). When AI is combined with other innovative science and technologies, the possibilities

grow profoundly. For example, functional Near-Infrared Spectroscopy (fNIRS) can detect neural signals that indicate when a person is engaging in reflective thinking. Heart rate variability, blood pressure, temperature, and electrodermal activity levels, can all be sensed, collected, and analysed to look for signals of frustration and stress. Gestures and movement patterns can be analysed by using 3D and 2D video to assess student engagement in collaboration; gaze patterns can illustrate student attention, eye-tracking data can be analysed to reveal students' emotions, cognitive load, and focus. The possibilities and potential are unbounded. Sadly, so too are the risks.

When the Institute for Ethical AI was created in 2018, it was because we were fearful that the benefits of this technology would be lost because an extremely negative and unethical event would occur due to the use of AI, and this would close down opportunities that could have positively transformed the lives of many people, particularly those who are disadvantaged. We produced a framework to help educational procurement ask the questions of themselves and of AI sellers that would help to ensure that the AI was beneficial. This is useful, but it is just one step in a much greater process through which educators must understand enough to make wise decisions about buying and using AI. There is also an important role for regulation, which must not be overlooked. But regulation will never be able to keep up with the developments in AI, and therefore *education is essential, urgent, and important*. We must educate the educators.

Engaging the educators

The challenge of engaging teachers in learning about AI is a tough one. I have recently published a book called *AI for School Teachers* (Luckin et al., 2022), which I wrote with an academic colleague and a headteacher. We wanted to write with a headteacher in order to try and ensure that we wrote in a manner that was relevant to teachers. In addition, our headteacher co-author was not someone who was 'a techie.' She is someone who values technology and believes it is important, but also someone who is not particularly proficient with using technology, nor did she initially understand a great deal about AI. We believed that if we could help her to understand AI, then she would be able to help us to understand how to write about it for other teachers. We all enjoyed the writing process and went to great lengths to find convincing, authentic educational examples for all the different ways in which we discussed and explained AI.

For example, we used a very common activity: planning a school trip to explain how Good Old-Fashioned AI (GOFAI) could be used to develop an AI application to help with stepping teachers through all the processes and decision points involved. We included checklists. For example, we suggested that teachers thought about the educational challenges they were facing, because these should lead decisions about how AI should be adopted. We provided a set of questions to prompt teachers about the challenges they might be facing that might be addressed by AI:

EXHIBIT X: QUESTIONS TO HELP IDENTIFY CHALLENGES THAT A SCHOOL TEACHER MIGHT BE FACING

- Are the children in my class attending regularly?
- Have the children developed good relationships amongst their peers?
- Are the children in my class well nourished?
- Have I planned lessons that meet the expected learning outcomes?
- Does my marking and feedback effectively engage students such that they act to improve their learning?
- How do I evaluate parental perceptions of me as a teacher?
- How do I ensure that soft skills are developed by my students?
- How do I know that we are meeting the expected standards for teaching and learning?
- What should I do to close the gender gap in maths?

Figure 1. Questions to help teachers identify challenges (Luckin et al., 2022, p.19).

And similarly, questions for headteachers and school leaders:

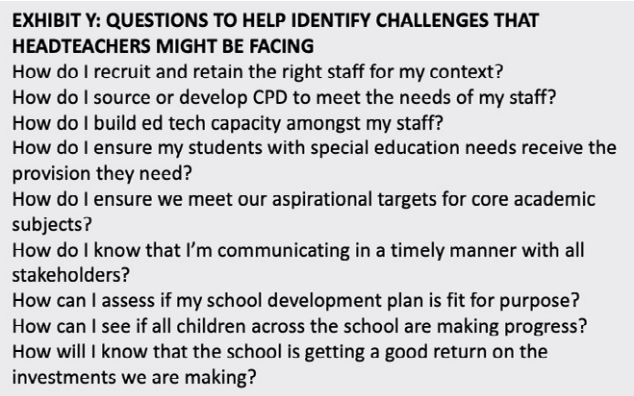

EXHIBIT Y: QUESTIONS TO HELP IDENTIFY CHALLENGES THAT HEADTEACHERS MIGHT BE FACING

How do I recruit and retain the right staff for my context?
How do I source or develop CPD to meet the needs of my staff?
How do I build ed tech capacity amongst my staff?
How do I ensure my students with special education needs receive the provision they need?
How do I ensure we meet our aspirational targets for core academic subjects?
How do I know that I'm communicating in a timely manner with all stakeholders?
How can I assess if my school development plan is fit for purpose?
How can I see if all children across the school are making progress?
How will I know that the school is getting a good return on the investments we are making?

Figure 2. Questions to help school leaders identify challenges (Luckin et al., 2022, p. 20).

We also set out a self-questioning process to help teachers and their leaders to decide which of their set of challenges is the best one to focus on first for thinking about AI (Luckin et al., 2022, pp. 20–23):

1. "Ask these of yourself, your colleagues, team, peers, managers or stakeholders, and use the answers to narrow your pool. What do you already know about this challenge? [Score 3 if you know a great deal, Score 2 if a modest amount, Score 1 if you don't know much, and Score 0 if you know nothing about this at all].

2. What kind of information is it possible for you to know that you don't know now? For example, if you are wanting to know more about the attainment gaps between different pupil groups, think hard about exactly what you could know about the pupils, their friends, family, context, etc. Or, perhaps you are concerned about bullying - there are different types of bullying in different degrees. For example, cyber, physical, name-calling, etc. It would be possible for you to explore the environmental conditions in the school that allow for these incidents to happen. [Score 3 if you are confident that you could know a great deal more, Score 2 if you believe that you could know a modest amount more, Score 1 if you are not sure that there is a great deal more that you could know, and Score 0 if you believe there is nothing more that you could know].

3. To what extent is the challenge you are facing controllable, and by whom? Are all systems and procedures understood clearly by all staff teaching and support? Are they audited, reported, and monitored? [Score 3 if the challenge is (a) controllable, (b) by someone at the school or within the school group, and (c) you do have all the systems in place to control the challenge; Score 2 if any two of (a), (b), and (c) are true; Score 1 if any one of (a), (b), and (c) is true; and Score 0 if none of (a), (b), or (c) is true]. For example, recruiting, training, and maintaining the best staff team. Any organisation only has limited control over the recruitment challenge, because whilst it can optimise all elements of the recruitment process that it adopts it cannot control how many people apply. Hopefully you have confidence that there are systems in place to help you optimise the elements of the recruitment process that are within your control, and AI can certainly help with that. However, the organisation cannot alter the number of people who are looking for the sort of employment that is on offer. Similarly, the school cannot control the pool of applicants that have the appropriate qualifications, skills, and expertise for the roles that need to be filled. [In this example, the score would be 2, because the whole of the recruitment process is not under your control, but you do have the systems in place to maximise the aspects of the process that are within your control].

4. What level of uncertainty is there? There may well be a level of uncertainty with a challenge based on incomplete reporting procedures

by staff and children, for example, or due to a challenge being surfaced through anecdotal evidence. [Score 3 if the level of uncertainty is negligible, Score 2 if there is a modest amount of uncertainty, Score 1 if there is a great deal of uncertainty, and Score 0 if there is no certainty at all].

5. Do you already have any data to help you understand this challenge or can you access data about this? [Score 3 if you have or can access a large amount of data from different sources, Score 2 if you have or can access a modest amount of data from different sources, Score 1 if you have or can access a very little data from any source, and Score 0 if you neither have, not have access to any data]. For example, you might have data derived from existing surveys, parent comments or complaints, behaviour logs and risk assessments.

6. Can you collect more data if you don't have enough data to help you understand this challenge and work out how best to tackle it? There are always opportunities to collect more data from students. [Score 3 if you can collect a large amount of relevant data, Score 2 if you can collect a modest amount of relevant data, Score 1 if you can only collect a small amount of relevant data, and Score 0 if you are unable to collect any new data at all].

7. How accurate can you be in your assessment of the challenge and your prediction about the best way to tackle it? [Score 3 if you can be very accurate, Score 2 if you can be quite accurate, and Score 1 if you can only be imprecise and therefore not very accurate at all, and Score 0 if you cannot be accurate at all]. For example, cyber bullying is a challenge that can be difficult to assess accurately, because it can occur outside of school grounds and systems.

8. Do you or your organisation have the appetite and capability to change to address this challenge? [Score 6 if the answer is "yes" and Score 0 if the answer is "no"]. If the answer is no, for whatever reason, it may not be a good investment of your time to be looking at ways AI can help you tackle the challenges in new ways.

9. Is the challenge AI compatible? [Score 3 if it is very AI compatible, Score 2 if it is modestly compatible, Score 1 if it is not very compatible, and Score 0 if it is completely incompatible]. This may be a difficult question for you to answer at the moment, but the section of this chapter entitled "Who has got the power, Artificial or Human Intelligence?" will help, as we hope will the rest of the book.

10. Finally, and most importantly, how important is solving this challenge to you or to your organisation? [Score 6 if it is crucial to solve this challenge, Score 4 if it is important to solve this challenge, Score 2 if it is quite important, and Score 0 if it is not important at all]."

And we found novel ways to communicate the intricacies of a machine learning algorithm, for example, through cookery, in this extract from the book (Luckin et al., 2022, pp. 72–73):

"I find it helpful to think about this situation as being a bit like cooking. There are lots of many types of cooking. We can bake, we can fry, we can broil, boil, or braise. We can grill, we can poach, we can smoke, sear, or sous vide. Just for a moment imagine that you are part of one of those TV shows where you are presented with a set of ingredients that are placed on a table and hidden under a cloth. You pull back the cloth to reveal the ingredients from which you must make something wonderful. Your instructions state that you are to make a dessert and that you must use all the ingredients in making your dessert.

Think about the ingredients as being a little bit like the data to which we want to apply AI. Back to the table. You have just whipped off the ingredients in our imaginary cooking show. You have eggs, and you have raspberries, plus there is cream and sugar.

What type of cooking method would be best applied to these ingredients? Frying is not an option, because whilst we could fry an egg, we have to use all the ingredients not just eggs. Grilling also looks unlikely to be suitable. Similarly, broiling or boiling is not really appropriate, but maybe baking could work for this.

The same type of situation exists when it comes to applying AI to our data "ingredients"; we need to decide what sort of AI could and should be applied. This decision is largely driven by the ingredients available and the challenge that we need to address, just as it is with cooking. There may be several options available to us for the same set of ingredients. Experience will help us to know which option to try first. Fortunately, unlike food, data can be subjected to multiple AI techniques that are appropriate to the type of data and the challenge being addressed.

Back to the cooking ingredients. We know we are required to solve the challenge of creating a dessert from the ingredients available. We also know that baking is likely to be the most appropriate cooking method to apply. The options available are now constrained by these parameters, but there are still options. Should we make raspberry pavlova or should we make raspberry souffle? Which of these is going to best meet the requirements of the cooking show, the challenge? We decide on raspberry souffle, because we have more experience of making this and therefore believe that a good result is more likely than with pavlova. Now this choice has been made, we know the method that we need to complete in order to produce the solution to our challenge: a dessert using the ingredients available to us.

The situation with data and applying AI is not so dissimilar. We have looked at our data (the ingredients) and the challenge (exploring the quality of teaching and learning when moving some provision online). We decide that the most suitable type of AI (cf. type of cooking) to be applied to these ingredients and this challenge is machine learning (cf. baking). Finally, we make a choice about the type of result we want to produce: finding patterns in the data for the teaching and learning interactions that have happened online and face to face (cf. pavlova or souffle). We can therefore now also choose the method of machine learning that we are going to apply; we choose unsupervised machine learning.

Returning for a moment to the raspberry souffle situation. We are now faced with needing to go through a set of preparations to be able to apply the baking process to the ingredients and produce a souffle. First, we have to wash the raspberries. Then we have to crack the eggs and whisk them. And then we have to add the sugar into the whisked eggs. We also have to whip up the cream and add that to the beaten eggs and the sugar. Finally, we add the washed raspberries. Then we need to mix it all together in the bowl. We now have the souffle mix and just need to put the mix into a dish, or a set of individual portion dishes, and we will be ready to apply the cooking method of baking to the prepared ingredients. As you can see, there is a lot of preparation. In fact, it may take longer to do all that preparation than it does to bake the souffle, which is really quick to bake.

For our education data and educational challenge situation, we want to explore the extent to which we have maintained the quality of teaching and learning, as things have moved online. It is important to note that there are many ways in which we could analyse our data, many of which have nothing to do with AI, but the point here is to see what extra insights and understanding the use of AI techniques can bring to the kind of analysis that is normally done with educational data. In our example here, we also want to understand more about AI."

Next steps

However, in conversations with teachers as they read the book, I still feel that we have missed an important connection, a connection that would really motivate and mean that understanding AI would become a priority. But what is that missing connection? Should we try to scare teachers with more lurid examples than those with which I started this chapter? Examples of ways in which the suits of AI already being used in education could be misused, abused, and cause harm? Should we gaze at the next generation of AI that is likely to be appearing in the classroom and illustrate the benefits and the risks in full and scary detail? I certainly believe we need something dramatic, but I don't feel comfortable with using scare tactics. Nor do I feel comfortable with the thought of teachers, students, parents, and the public being hoodwinked by smart salespeople who know how to tell a positive story and avoid catalysing any concern within their audience.

I am still experimenting to find other ways that will capture their attention, to make people sit up and get them to believe that understanding AI is vital. I have no smart answer to conclude the chapter, just a plea for more attention to be paid to finding effective ways to motivate educators to want to understand AI and the right tools for helping them to succeed.

References

Hao, K. (2019). China has started a grand experiment in AI education. It could reshape how the world learns. *MIT Technology Review, 123*(1). https://fully-human.org/wp-content/uploads/2019/09/China-AI-experiment-education.pdf

Kasparova, A., Celiktutan, O., & Cukurova, M. (2020, October). Inferring student engagement in collaborative problem solving from visual cues. In *Companion Publication of the 2020 International Conference on Multimodal Interaction* (pp.177–181). https://doi.org/10.1145/3395035.3425961

Keierleber, M. (2022, March 30). 'Really alarming': the rise of smart cameras used to catch maskless students in US schools. *The Guardian*. https://www.theguardian.com/world/2022/mar/30/smart-cameras-us-schools-artificial-intelligence

Kynge, J. (Host). (2022, April 25). US-China Tech Race: brave new world [Audio podcast episode]. In *Tech Tonic*. Financial Times. https://www.ft.com/content/33ef2363-6f97-43dd-8b59-7fca025684cc

Luckin, R., George, K., & Cukurova, M. (2022). *AI for School Teachers*. CRC Press. https://doi.org/10.1201/9781003193173

Marr, B. (2018, July 25). How is AI used in education. Real world examples of today and a peek into the future. *Forbes*. https://www.forbes.com/sites/bernardmarr/2018/07/25/how-is-ai-used-in-education-real-world-examples-of-today-and-a-peek-into-the-future/?sh=70581967586e

Narayanan, A. (2019). *How to recognize AI snake oil* [PowerPoint slides]. Princeton University. https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf

Ram, A. (2019, March 5). Europe's AI start-ups often do not use AI, study finds. *Financial Times*. https://www.ft.com/content/21b19010-3e9f-11e9-b896-fe36ec32aece

Stewart, A.E.B., Keirn, Z., & D'Mello, S.K. (2021). Multimodal modeling of collaborative problem-solving facets in triads. *User Modeling and User-Adapted Interaction*, *31*(4), 713–751. https://doi.org/10.1007/s11257-021-09290-y

Wall Street Journal. (2019, October 1). *How China is using Artificial Intelligence in classrooms* [Video]. YouTube. https://www.youtube.com/watch?v=JMLsHI8aV0g

Xie, E. (2019, September 16). Artificial intelligence is watching China's students but how well can it really see? *South China Morning Post*. https://www.scmp.com/news/china/politics/article/3027349/artificial-intelligence-watching-chinas-students-how-well-can

Rose Luckin (University College London, UK)

Rosemary (Rose) Luckin is Professor of Learner Centred Design at UCL Knowledge Lab. She was named one of the 20 most influential people in education in the Seldon List, 2017. Rose is Founder of EDUCATE Ventures Research Ltd., a London hub for start-ups, reseachers, and educators. She is treasurer of the International Society for AI in Education and co-founder of the Institute for Ethical AI in Education. Prior to joining Knowledge Lab in 2006, Rose was Pro-Vice Chancellor for Teaching and Learning at the University of Sussex.

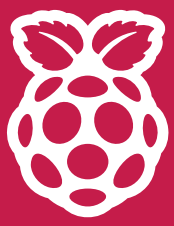

Raspberry Pi
Foundation

www.raspberrypi.org

Raspberry Pi Foundation

@RaspberryPi_org

Raspberry Pi Foundation