Sue Sentance and Jane Waite
University of Cambridge and Raspberry Pi Foundation

Perspectives on AI and data science education

Sentance, S. & Waite, J. (2022). Perspectives on AI and data science education. 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-sentance-waite

Perspectives on AI and data science education

Sue Sentance and Jane Waite
University of Cambridge and Raspberry Pi Foundation

Introduction

Artificial intelligence (AI) and data science education have a far-reaching and growing impact on our lives, and it is important for young people to understand them both from a technical and a societal perspective, and for educators to learn how to best support them to gain this understanding.

From September 2021 to March 2022, the Raspberry Pi Foundation ran a series of seven seminars on the topic of AI and data science education for young people. The objectives of the seminar series were:

1. To learn from experts in the field about their perspectives on the future of AI and data science education for young people;
2. To develop a community of interested researchers, teachers, and industry experts around this topic.

The invited speakers brought a range of different perspectives to the topic, in terms of their approaches to theory, resources, and their ambitions for AI and data science education. In this short article, we summarise our understanding of their presentations and how their work may contribute to a research agenda for this new and emerging field.

Five of these seminars have been supplemented with chapters in these proceedings. We also held a special panel session including young people and a UK Minister, which looked at policy and perspectives of AI and data science education as a school subject. You can view all the seminars at this link: rpf.io/ai-research-seminars

Why teach AI?

In the UK, like many countries, we have a very crowded school curriculum with many different subjects jostling for curriculum time. We have made progress globally in introducing computer science into some school-aged contexts, mostly at the secondary school level (Vegas et al., 2021). However, AI is a subject typically taught at the Master's level, although some undergraduate degrees in the topic have been available in recent years. Although AI is spoken about in many contexts, we may not even have a shared definition of what it really is or covers. So why would we even consider adding the teaching of AI either in school or in non-formal settings? We still have much to learn about how and what to teach in terms of AI, machine learning, and data science. During the seminar series, our speakers provided a range of different perspectives to the question of why we should teach AI. From these, it is possible to extract a number of different reasons for teaching AI.

Children are already growing up with AI: This is probably the most obvious reason that people cite when thinking about AI. Young children are already surrounded by devices and apps that use AI, so the argument goes that they should learn about how they work to become discerning consumers. Stefania Druga discussed her research on working with families who were developing an understanding of the potential of smart devices (Druga et al., 2021).

AI is impacting children's lives: AI may have far-reaching consequences in children's lives, where it's being used for decision-making around access to resources and support. From an ethical perspective, Mhairi Aitken holds the view that AI systems are already having a significant impact on young people's lives through systems

Table 1: Seminars hosted by the Raspberry Pi Foundation.

Title	Speaker(s)
AI ethics and engagement with children and young people	Dr Mhairi Aitken, The Alan Turing Institute
Exploring the data-driven world: Teaching AI and ML from a data-centric perspective	Professor Carsten Schulte, Yannik Fleischer, and Lukas Höper, University of Paderborn
ML education for K–12: Emerging trajectories	Professor Matti Tedre and Dr Henriikka Vartiainen, University of Eastern Finland
What is it about AI that makes it useful for teachers and learners?	Professor Rose Luckin, University College London
Teaching artificial intelligence in K–12	Professor Dave Touretzky, Carnegie Mellon University, and Professor Fred Martin, University of Massachusetts Lowell
Teaching youth to use AI to tackle the Sustainable Development Goals	Dr Tara Chklovski, Technovation
Democratising AI education with and for families	Stefania Druga, University of Washington

deployed in children's education, in apps that children use, and in children's lives as consumers. Children's data is being collected, and decisions are being made about them using AI; therefore, awareness of the impact of AI should be raised (Aitken & Briggs, 2022).

AI requires a new way of thinking: Two of our seminars covered the ways in which our understanding of computational thinking changes when we move away from traditional programming to more data-driven approaches. Matti Tedre and Henriikka Vartiainen propose a new version of computational thinking called CT 2.0. In contrast to CT 1.0, which is rule-driven, CT 2.0 is data-driven, so requires skills such as being able to experiment with data (Vartiainen et al., 2021). Dave Touretzky and Fred Martin proposed a broad version of AI thinking, which includes perception, reasoning, representation, machine learning, and language understanding (Touretzky et al., 2019).

People need to use AI safely and effectively: In order to build a citizenry of people who use AI safely and effectively, we need to educate them in the subject. Rose Luckin shared a very broad view of AI — in education, for education, and as part of our education. Luckin emphasised the importance of being able to customise AI tools to your context (Luckin et al., 2016).

We want to empower children to effect change: Understanding AI and data science will be very empowering in the years to come. It's important that AI education is inclusive and that opportunities to learn about AI are for everyone. To do this, we need to make AI education available and accessible.

Humans are starting to interact with machines in new ways: In our most theoretically-focused seminar, Carsten Schulte argued for a new discipline around machine behaviour and hybrid human interaction, focusing on the ways in which society and individuals interact with data-centric systems (Rohlfing et al., 2021).

We need a skilled AI workforce: This is another reason for teaching AI, one that was not named in any of the talks, but which is put forward by policymakers describing their country's development of AI: in order for a country to lead in developments in AI, it needs a trained workforce with the appropriate technical skills (Galindo et al., 2021). In the UK, we've seen the publication of the National AI Strategy[1] and the AI Roadmap[2], both highlighting the need for AI education, and this year's AI action plan[3] focuses on supporting the development of a diverse workforce in AI, and other countries in the world have similar policy drives.

Just as there are many views on why we should teach AI, experts and academics hold different views on what we should actually teach within the vast area of AI. It is clear that the motivation for teaching AI to young people impacts the actual

Table 2: Seven reasons for teaching AI in schools.

Why teach AI?	Implications for teaching content
Children are already growing up with AI	Young people need to learn that there are both drawbacks and advantages of innovative technologies, particularly where they use AI.
AI is impacting children's lives	Creators of systems that children will use should understand that AI may impact their privacy and that systems are being used to make decisions that affect them.
AI requires a new way of thinking	We should teach skills and knowledge around data-driven programming and how AI works in addition to traditional programming techniques.
People need to use it safely and effectively	Young people should be taught to use AI tools and applications.
We want to empower children to effect change	We should ensure that there are opportunities to learn all aspects of AI, both technical and socio-ethical, for all children at an early age.
Humans are starting to communicate with machines in new ways	We need to teach students about the ways that machines, including AI systems, impact individuals and society, and to be curious about the way machines behave.
We need a skilled AI workforce	We need to provide a progression of learning opportunities that lead towards highly technical courses in AI later on in school or ensure that facilitating subjects such as mathematics, physics, and computer science are taught effectively to all.

[1] https://www.gov.uk/government/publications/national-ai-strategy

[2] https://www.gov.uk/government/publications/ai-roadmap

[3] https://www.gov.uk/government/publications/national-ai-strategy-ai-action-plan/national-ai-strategy-ai-action-plan

content that might be taught, and at what age and stage it could be introduced. Table 2 summarises, at a very high level, the implications of certain motivations for teaching AI on the type of content that might be needed.

It is clear that we need to have specific goals in AI education, and that curriculum developers and educational resource developers may have different views on what we need to teach. In one of the seminars, we heard about the five big ideas of AI from Dave Touretzky and Fred Martin (Touretzky et al., 2019). The five big ideas from the AI4K12 project[4] are perception, representation and reasoning, learning, natural interaction, and societal impact. These have been really useful in both mapping to school standards in the US in computer science, and also in giving a framework for resource developers. In this way, the big ideas used by the AI4K12 framework help to show the breadth of AI content that we could cover. Our own research (in progress) has shown that many current resources focus on machine learning, so the AI4K12 framework highlights other areas of AI that could be studied.

However, there is another dimension: the degree to which we abstract from the technical aspects of AI. Do we teach children how to actually create AI, or do we teach them how it impacts them and how to be informed users of AI? And there is much in between those two aspects of the subject.

In Appendix 1, we have included a simple framework that we are using to categorise different levels of AI as SE (socio-ethical), A (applications), M (models), and E (the engine — or how AI works). This gives us a way of understanding different resources and their learning goals. It provides levels of abstraction for the subject, with the SE level most abstracted from the technical aspects. We are calling this the SEAME Framework.

Our seminar speakers had different perspectives on which of these elements were important to be understood. While Mhairi Aitken gave an excellent exposition of ethical issues and why we should engage children in them (focusing on the SE level), Dave Touretzky and Fred Martin talked about the fact that while young children might be using applications of AI (the A level), older children should engage more with the models of AI (the M level) and argued for transparent AI demonstrations that made the E level visible. Rose Luckin focused on teachers' knowledge of AI, which she argued could be developed by actually using data to create a model (the M level).

How and when would we teach it?

Computing is increasingly being introduced into curricula around the world, and in England has been mandatory for students aged 5–16 since 2014. Students in many countries can opt to take computer science as an elective in upper secondary or high school.

AI at the university level is likely to be included as part of a computer science department's course, but it may not necessarily follow that AI education will fit into the computing school curriculum. Some of the socio-ethical components could be addressed across a range of school subjects, for example.

Some of our seminar speakers gave examples of AI education in **non-formal** settings. For example, Stefania Druga shared the findings of studies carried out with families working in an informal way with their children (Druga et al., 2021), and Tara Chklovski gave examples of an annual challenge that children could sign up to as an extracurricular activity. Some of the work described in Matti Tedre and Henriikka Vartiainen's talk took place in the homes of young children aged from 3–9 (Vartiainen et al., 2021).

In terms of **formal** education, the Finnish speakers have also conducted some research with 11–12-year-olds in schools. The AI4K12 project is intended for formal education, through its mapping to the CSTA (Computer Science Teachers Association) standards (Touretzky

[4] https://ai4k12.org/

et al., 2019), although in reality many US states may not have the curriculum in place to deliver this content. ProDaBi[5], the data science and AI education project described by the researchers from Paderborn, is being designed for students in school at the lower secondary and upper secondary levels.

The **age range** of children in the studies we heard about in the seminars ranged from 3–18 years old, so it was clear that the discussion about AI education spans from kindergarten through primary and secondary education. One question that came up often from our seminar audiences was the extent to which teachers can be involved in research projects such as those described, and the level of training being developed for teachers to enable them to teach and understand AI. Some research projects we heard about were conducted in a participatory way, and certainly the AI4K12 project has developed sufficiently to have courses for teachers embedded into it. Apart from Rose Luckin, who described an adult-facing programme that could be used to support teachers wishing to understand AI, the seminar speakers did not particularly focus on the needs of the teacher in this context.

Our speakers had different views on **how** AI should be taught to young people. It was clear that AI is relevant to young people's lives, and Tara Chklovski highlighted how young people could be engaged in building solutions to problems that they could see in their own lives by accessing technology (Chklovski et al., 2021). Other speakers discussed how AI might require a change to the way we think. For example, Matti Tedre, of the University of Eastern Finland, proposed CT 2.0, which he's written about elsewhere, explaining that a data-driven approach to solving a problem is fundamentally different to writing an algorithm to solve it (Tedre et al., 2021). Carsten Schulte and colleagues, of the University of Paderborn in Germany, also highlighted issues around the role of code and the approach to accuracy, and how these are both different in machine learning in comparison to traditional programming (Rohlfing et al., 2021). Both research groups are developing resources that reflect these differences, and studying the way that learners interact with them. This is interesting work, and we will be following the updates of these two research groups with much interest.

Where is research needed?

What our speakers said they wanted to do next gives us an interesting range of ideas for future research. Tara Chklovski wants to continue to broaden participation by ensuring that more girls and underrepresented groups in computing can access the opportunities to develop AI skills through team challenges. Stefania Druga calls on us to consider family life as a third space for AI learning and suggests there is much more research to do in this area. The AI4K12 project is concerned with reaching more US states and also proposed further work on tools development around teaching AI thinking. Rose Luckin's work at UCL is much broader than our specific context and extends to the use of AI in education, where there is much to do to ensure that this is implemented ethically. Linked to this, Mhairi Aitken's future work will involve actually engaging with children to support ethical practices in AI: this is crucial, as we so often ignore the young person's voice. Matti Tedre and Henriikka Vartiainen left us with many questions and challenges regarding how we can make the shift from CT 1.0 to CT 2.0. Aligned to this, Carsten Schulte's summary included a call to action for us to conduct research that helps us to understand the data-driven and emergent ecosystem, and to investigate how that might impact a paradigm change in teaching.

Drawing together these ideas for future work, we have suggested four areas which we believe should be included within a research agenda for AI education:

[5] https://www.prodabi.de/

1. **Teaching and learning.** It's clear that our traditional approach to teaching programming, which involves writing an algorithm that can be implemented, will need to change as we introduce young people to more data-driven approaches to solving problems. What does CT 2.0 look like in a teaching and learning environment?

2. **Learner voice.** We need to engage learners in research around their perception of issues that affect them. There are links to culturally responsive computing research as well as opportunities to develop learners' thinking around social justice and equity in our teaching about AI.

3. **Teachers/educators.** At the Raspberry Pi Foundation, we've recently conducted a literature review of empirical research in the area, which demonstrates that educators are not often included in studies or not seen as a stakeholder in AI education work. There is much to do here.

4. **Tools and resources.** We've conducted a mapping exercise looking at AI resources written for children, which surfaces a complex picture, beyond simply what and who is taught. For example, the choice of software can limit the transparency of what is being taught. A framework for resource development would be a useful addition to the field.

We're going to start working on some of the four areas listed and would encourage others to do the same. There is much research to be done!

References

Aitken, M. & Briggs, M. (2022). In AI, data science, and young people. Understanding computing education (Vol 3). Proceedings of the Raspberry Pi Foundation Research Seminars. rpf.io/seminar-proceedings-vol-3-aitken-briggs

Chklovski, T., Jung, R., Anderson, R., & Young, K. (2021). Comparing 2 Years of Empowering Families to Solve Real-World Problems with AI. *KI-Künstliche Intelligenz, 35*(2), 207–219. https://doi.org/10.1007/s13218-021-00738-2

Druga, S., Yip, J., Preston, M., & Dillon, D. (2021). The 4As: Ask, Adapt, Author, Analyze - AI Literacy Framework for Families. In Algorithmic Rights and Protections for Children. https://doi.org/10.1162/ba67f642.646d0673

Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). *Intelligence unleashed: An argument for AI in education*. UCL Knowledge Lab. http://oro.open.ac.uk/50104/

Galindo, L., Perset, K., & Sheeka, F. (2021). An overview of national AI strategies and policies. *Going Digital Toolkit Note*, 14. https://doi.org/10.1787/c05140d9-en

Rohlfing, K. J., Cimiano, P., Scharlau, I., Matzner, T., Buhl, H. M., Buschmeier, H., Esposito, E., Grimminger, A., Hammer, B., Hab-Umbach, R., Horwath, I., Hullermeier, E., Kern, F., Kopp, S., Thommes, K., Ngonga Ngomo, A.-C., Schulte, C., Wachsmuth, H., Wagner, P., & Wrede, B. (2021). Explanation as a Social Practice: Toward a Conceptual Framework for the Social Design of AI Systems. *IEEE Transactions on Cognitive and Developmental Systems, 13*(3), 717–728. https://doi.org/10.1109/TCDS.2020.3044366

Tedre, M., Denning, P., & Toivonen, T. (2021). CT 2.0. *21st Koli Calling International Conference on Computing Education Research*, Article 3. https://doi.org/10.1145/3488042.3488053

Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019). Envisioning AI for K–12: What Should Every Child Know about AI? *Proceedings of the AAAI Conference on Artificial Intelligence, 33*, 9795–9799. https://doi.org/10.1609/aaai.v33i01.33019795

Vartiainen, H., Toivonen, T., Jormanainen, I., Kahila, J., Tedre, M., & Valtonen, T. (2021). Machine learning for middle schoolers: Learning through data-driven design. *International Journal of Child–Computer Interaction, 29*, Article 100281. https://doi.org/10.1016/j.ijcci.2021.100281

Vegas, E., Hansen, M., & Fowler, B. (2021). *Building skills for life: How to expand and improve computer science education around the world*. The Brookings Institution. https://www.brookings.edu/wp-content/uploads/2021/10/Building_skills_for_life.pdf

Appendix 1: The SEAME framework

A simple learning levels framework can be used to categorise research and resources. This is derived from a framework developed at Queen Mary University of London by Jane Waite and Paul Curzon[6].

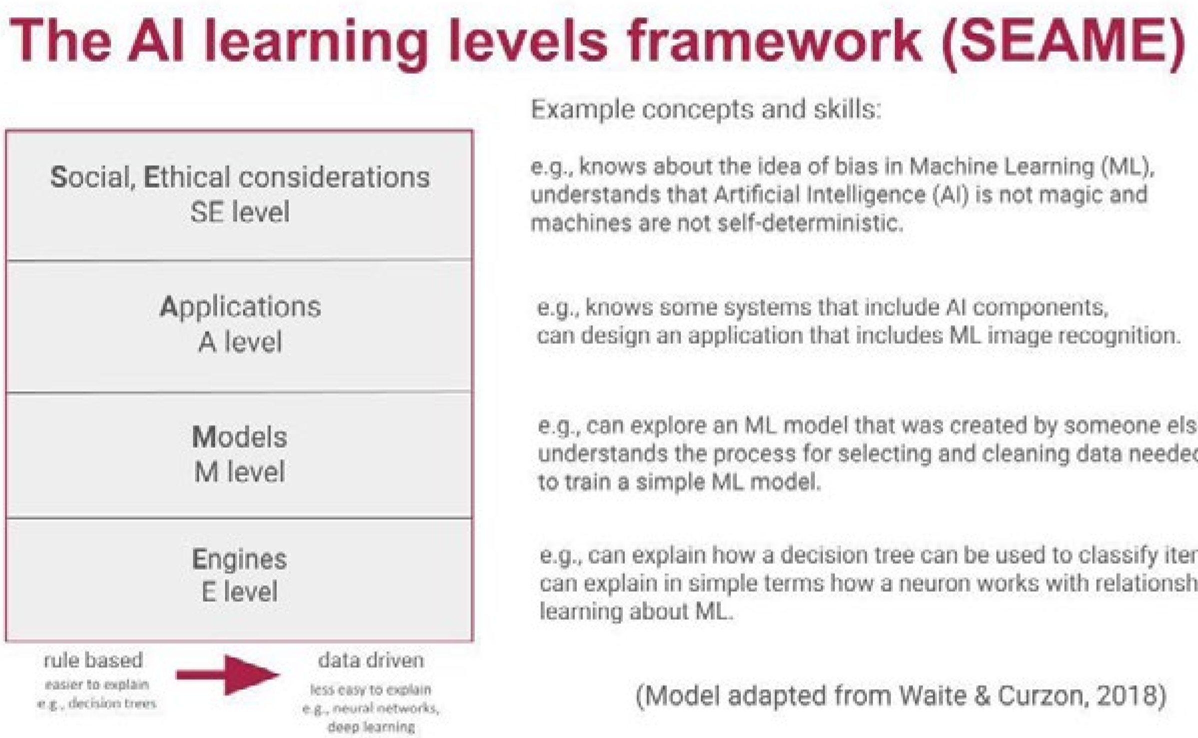

Figure 1. The AI learning levels framework (a revised version of Waite & Curzon, 2018[6]).

The SEAME AI learning levels framework used in our studies is shown in Figure 1. This framework has four levels and provides a simple way to reflect upon the content included in AI resources and activities.

SE: This is the level of **social and ethical** considerations.

A: This is the **applications** level, where we might use, modify, or create applications that have some AI or ML component.

M: This is the **models** level, where we train the model with data. Models output recommendations and predictions for use in applications.

E: This is the **engines** level, including neural networks, generative algorithms, data structures, etc. This is the most hidden level, which we are not aware of when we use an application with an ML component.

[6] https://teachinglondoncomputing.org/machine-learning/

The framework will require full evaluation, but is currently providing a valuable way for the Raspberry Pi Foundation to review and reflect on available research and resources. It is not intended to cover data science resources and research, as there are aspects of data science that are more statistics related, but it covers aspects of early data literacy.

How to use the framework

The framework can be used to categorise resources developed using AI. Some examples are given below.

The ethical dilemmas of self-driving vehicles as discussed with students can be described as level SE, the level relating to ethical and societal considerations.

Some activities might span two levels. For example, an activity where students use an existing ‘rock-paper-scissors’ application that uses an ML model to recognise hand shapes works at the Applications level. If students then move on to train the model to improve accuracy by adding more image data, they work at the Model level.

Other resources drill down through the layers for a single concept. For example, if studying bias, an activity might start with an example of the societal impact of bias. Students might then discuss the applications they use personally to reflect on bias, and the activity might finish with students exploring data in a simple ML model. This involves students working through layers SE, A, and M.

Another approach to using the framework is to see whether some age groups might have more learning activities available at one level than another and whether this changes over time. For example, younger learners might work mostly at levels SE and A, and older learners might move between the levels with increasing clarity as they develop their knowledge.

Sue Sentance (University of Cambridge and Raspberry Pi Foundation, UK)

Dr Sue Sentance is Chief Learning Officer at the Raspberry Pi Foundation and Director of the Raspberry Pi Computing Education Research Centre. She researches the teaching of programming in schools, teacher professional development, and physical computing. Her academic background is in computer science, artificial intelligence, and education, and she is a qualified teacher and teacher educator. She has created and researched the PRIMM methodology for structuring programming lessons in school.

Jane Waite (University of Cambridge and Raspberry Pi Foundation, UK)

Jane Waite is a computing education researcher who has worked in industry and as a classroom teacher for many years. She currently works as the senior research scientist at the Raspberry Pi Foundation in their research team. The research team is part of the Raspberry Pi Computing Education Research Centre, a joint initiative with the University of Cambridge. Jane is currently working on a wide range of research projects, including investigating culturally relevant pedagogy for teaching computing and looking at the underpinning concepts for teaching and learning AI in schools. She has published on a wide range of topics, such as on computational thinking with Tim Bell and Paul Curzon, pedagogy for teaching programming for the Royal Society, and primary program design, which is her main interest and passion.

Raspberry Pi Foundation

www.raspberrypi.org

Raspberry Pi Foundation

@RaspberryPi_org

Raspberry Pi Foundation