Experience AI: Introducing AI/ML to Grade 6–8 students in the UK

Jane Waite[1][0000−0002−0270−2124], Ben Garside[1][0009−0003−0625−7341], Robert Whyte[2][0000−0002−1214−1957], Diana Kirby[2][0000−0003−3630−2548], and Sue Sentance[2][0000−0002−0259−7408]

[1] Raspberry Pi Foundation, Cambridge, UK
{jane.waite,ben.garside}@raspberrypi.org
[2] Raspberry Pi Computing Education Research Centre,
University of Cambridge, Cambridge, UK
{rw724,dgk21,ss2600}@cam.ac.uk

Abstract. The call for school students to learn about artificial intelligence and machine learning is increasing, yet what should be taught and how is still to be agreed upon. Despite this, resources are starting to be developed. Such material can be used to explore pedagogical content knowledge decisions and evaluate the impact of teaching activities. In this experience report, we detail the development and implementation of Experience AI, a free curriculum unit consisting of six one-hour lessons designed for classroom use by teachers of students aged 11 to 14 years old (Grade 6–8) in the United Kingdom. The lessons were designed with an underlying set of design principles developed in consultation with industry experts. The design principles we focus on in this report are (i) avoiding anthropomorphisation of language and images used in the resources, (ii) incorporating careers materials and activities, and (iii) increased teacher support for lesson delivery. The resources include teacher guides, classroom presentations, explanations of key terms, student activities, and assessment ideas. From an independent evaluation of the implementation of the lessons, initial survey results are reported (student conceptions of AI, student and teacher AI careers awareness, teacher self-efficacy when teaching about AI). Evidence from the evaluation has provided early yet encouraging evidence that teachers who used the curriculum unit may have improved their AI careers awareness and efficacy when teaching about AI. We suggest the design principles, lesson materials and evaluation instruments may be useful to other researchers working in this field.

Keywords: Artificial intelligence · machine learning · curriculum

1 Introduction

Even though artificial intelligence (AI) systems are becoming ubiquitous across society [9], AI technology is not widely understood by those affected [14]. Creating an AI-ready workforce is a significant focus for many governments [4]. In the United Kingdom (UK), AI policy has included the development of recommendations [18], policy options [21], and roadmaps for research and change [25]. AI education has also been called to be incorporated into the school and teacher preparation curricula [21, 25]. In the computing education research community, there are a growing number of initiatives to do this [24, 22]. However, there is limited empirical research to understand what and how AI should be taught, and what the impact of early initiatives might be [23, 17]. This presents a challenge for K–12 computing education researchers, education resource developers, and teachers alike to consider how AI and machine learning (ML) can be introduced.

In this experience report, we describe the design of a set of 6 lessons introducing AI/ML to students aged 11 to 14 (Grades 6 to 8) in the UK. The resources were developed with support from Google DeepMind through a volunteer industry expert working group. We also detail results from an independent evaluation of the school implementation of the resources.

2 Background and related work

What might be taught about AI has been defined in various but conflicting ways. For example, the AI4K12 working group suggested 'Five Big Ideas' for K–12 AI: perception, representation and reasoning, learning, natural interaction, and societal impact [24]. From an analysis of 30 K–12 instructional units on ML, Marques et al. [8] identified 12 ML topics (e.g. neural networks), 13 ML applications (e.g. sentiment analysis), and 7 ML processes (e.g. model evaluation). A computational thinking (CT) view of learning about AI has been suggested [23], whereby the difference between rule-driven and data-driven system development paradigms is emphasised, and a new CT2.0 is defined. Differences in the problem-solving workflow of CT2.0 (data-driven) to CT1.0 (rule-driven) are compared, including describing the job and collecting the data rather than formalising the problem; filtering, cleaning and labelling the data rather than designing an algorithmic solution; training a model rather than implementing the algorithm; and evaluating and using the model rather than compiling and executing the program [23]. Clearly, a consensus about what should be taught has yet to be reached. Olari and Romeike [12] argued that most AI literacy frameworks fail to capture data science (or "data literacy") concepts and skills, and Druga et al. [3] noted that a common language for AI and ML teaching resources had not been agreed upon.

As well as considering what should be taught, an open question is how AI and ML should be taught and which pedagogy should be used. A set of 15 design considerations have also been defined [7], including contextualising data, opportunities to program, and leveraging learners' interests. From a synthesis of AI teaching studies in K–12 and a supposition of what might work for teaching this age group and topic, a taxonomy of pedagogies for AI has been suggested, including active learning, personalised learning, participatory, problem-based, interactive, project, inquiry, and design-oriented learning [20]. Finally, a simple AI and ML learning framework, called SEAME has been proposed for use in reviewing AI teaching resources [26] and research activities [17]. The framework comprises four levels: Social and Ethical, Application, Model and Engine. In SEAME, the levels do not dictate the order in which learning must occur, and some activities will span more than one level. The levels provide an intuitive way for educators, resource developers, and researchers to frame the main aims of a learning experience. As students make progress, it is expected students can move between levels, and, at times, the boundaries between levels will blur.

3 Design principles

For the lesson design, a set of design principles were devised to guide the development of AI resources. These principles extended existing general resource design principles for computing resource development [16]. In this paper, we focus on three of the AI design principles. These principles we report on here were selected as we can comment on their enactment through evaluation data. The three design principles we focus on are:

- avoiding anthropomorphisation
- promoting awareness of AI/AI-related careers
- increased teacher support

Other design principles, but not reported on in detail here, include (i) using the SEAME framework [26] to help develop the learning objectives and learning progression [16], (ii) developing a set of working explanations for learning objective concepts and sub-concepts for the research team and educators [16], and (iii) using semantic profiling to align everyday contexts to abstract technical language [10] for explanations and lesson design.

Avoiding anthropomorphisation: A key consideration was the need to avoid anthropomorphisation in student- and teacher-facing materials. Anthropmorphisation is “the action or fact of attributing human characteristics, form, or personality to something non-human (in later use esp. an animal)”.[3]

The rationale for this choice was that attributing human characteristics to computers has led to programming misconceptions [13], and more specifically for AI, may lead to system developers (including novice programmers) developing incorrect mental models of how AI works, as the technology is humanised, black-boxed, and oversimplified [23]. Additionally, when using technology, young children have been found to view robots as peers rather than devices, seeing them as less smart ‘people’ or overestimating technology capabilities [2], or developing relationships with the devices [23], leading to high risks of either unintended influence, purposeful manipulation, or policing [28]. Compounding the issue of delegating the responsibility of the human system developer and human user to an imagined responsible AI agent [19], anthropomorphised AI agents have been predominately portrayed as white in colour and as such exacerbating racism in technology and society at large [1]. However, by engaging students to learn

[3] $https://www.oed.com/dictionary/anthropomorphization_n$

about AI and making the 'black box' transparent, students become more sceptical about the technology and recognise responsibility for AI design as associated with the humans who are 'in the loop' rather than the AI technology [2, 23]. For Experience AI, in practice, this principle was enacted by there being no illustrations that depicted devices with human-like faces and the replacement of vocabulary that was associated with the behaviour of people (see, look, recognise, create, make) with system-type words (detect, input, pattern match, generate, produce).

In line with suggestions by Druga et al. [3] to create a shared language among curriculum developers, a set of working explanations for learning objective concepts/sub-concepts was developed in partnership with an industry working group. For example, a working definition of AI literacy was defined as *"AI literacy is a set of competencies that enables people to use AI applications in everyday life creatively and ethically, to identify and evaluate AI technologies critically, and to have a basic knowledge and understanding of the key concepts and processes associated with AI applications, models, and engines."* [16].

Promoting awareness of AI/AI-related careers: A second design choice was to provide career examples in each lesson that shared social and ethical issues through relatable real-world examples of applications of ML models. The aims were to 1) engage students and 2) help them understand the relevance and impact AI has in their lives. The rationale for this choice was the need for students to understand the career and societal implications of AI developments. Prior work in computing education has highlighted the challenge in promoting the aspirations of young people, particularly female students, in pursuing careers in the field [5]. Studies have indicated multiple factors that impact young people's career aspirations in computing, including a lack of exposure to CS, the need for role models, and the influence of parents and self-efficacy beliefs. As such, researchers have underscored the importance of promoting awareness of AI career opportunities and the broader impact of AI across disciplines [29]. In our lesson materials, the careers principle was enacted by demonstrating the breadth of careers in AI/AI-related fields through real-world examples. Video interviews with researchers and scientists working at Google DeepMind were featured throughout the lessons to enrich classroom discussions on career goals.

Increased teacher support: Though teaching and learning of AI and ML in schools is an increasingly important topic being suggested for classroom teachers to consider [4], it is unlikely that teachers will have prior experience of AI/ML or appropriate pedagogical knowledge when working with school-aged learners [23, 30]. Therefore, in Experience AI, to increase teacher confidence and self-efficacy, teacher support was provided through the following means:

- *Student-facing concept videos* embedded in the lesson slides, including industry experts explaining key AI/ML concepts.
- *A free asynchronous online teacher professional development course* on introducing AI/ML, how to deliver the lessons, and understanding concepts.
- *Teacher support videos* to introduce lesson activities, including screen-casts demonstrating the steps students will need to follow for practical tasks.

– *Webinars* for teachers including in-depth discussions on AI/ML concepts with industry experts and guidance on how to deliver lesson activities.
– *Lesson plans* which include detailed guidance on how to deliver the lesson activities, including additional information about the concepts and suggestions on formative assessment questions to ask students.

Prior surveys of AI/ML teaching and learning resources indicate a lack of teacher documentation in most currently-available resources [26]. Our intention was to provide support materials to anticipate teacher needs and provide guidance on difficulties that could be faced in the classroom.

4 Lesson resources

The free unit of work consists of six lessons available for educators to download [16]. An overview of the lessons and links to the full resource set is provided in Appendix A. The lessons are intended to cover a six-week period. Each lesson is designed to be taught in a one-hour classroom-based lesson. The unit of work is aimed at school educators in the UK who teach students aged 11 to 14 years old, although they are available for use by any educator in other countries. There are no specific hardware requirements for the lessons, except for equipment to access the internet. A web browser is needed to download and view lesson material; no software needs to be installed locally. Machine Learning For Kids[4] is used as a web app and is accessible without the need for an account.

The ambition for the six lessons is to provide students with a foundational knowledge of AI concepts, contexts, and the careers involved in developing AI applications. Building on the four strands of SEAME [26], key concepts covered include: (i) rule-based vs data-driven approaches to programming; (ii) applications of AI; (iii) ML models; (iv) bias and ethics; (v) decision trees; (vi) the AI data life cycle; and (vii) careers in AI. A description of concepts covered in the lessons is provided in Appendix A.

For each lesson, there are lesson plans, teaching slides, student activity worksheets, teacher support videos, projects for students to select from depending on their personal interests, and student assessment activities (formative and summative). In addition, there are three overarching documents available to teachers: a unit overview, learning graphs (to demonstrate progression), and a set of explanations of key terms (e.g. concepts).

The learning objectives were reviewed to draw out candidate sub-concepts (key terms); for each sub-concept, an explanation was developed [16]. Example sub-concepts incorporated in the lessons include (this is not exhaustive) prediction, supervised learning, unsupervised learning, reinforcement learning, ML classification, ML class, ML label, ML decision trees node, data types, ML training, ML training data, ML test data, data bias, societal bias, ML explainability, ML accuracy, ML confidence, data cleaning, ML model card, computer vision, and generative AI. The degree to which knowledge is built for each of these

[4] https://machinelearningforkids.co.uk/

sub-concepts varies; some are introduced in passing to provide context for the exercise. Some sub-concepts demonstrate the breadth of ML to avoid introducing alternate conceptions, such as thinking that supervised learning is the only method of solving problems using ML. Other concepts are returned to multiple times, such as ML training.

5 Method

An independent evaluation was conducted for both student and teacher participants. The evaluation consisted of an online pre/post survey distributed to students and teachers via Qualtrics. The evaluation study was approved by the University of Cambridge Department of Computer Science and Technology Ethics Committee (#2023) and informed consent was obtained from all participants.

Instrument design: The student survey gathered quantitative data via five-point Likert-type questions relating to multiple constructs, including students' interest in AI and their awareness of AI/AI-related careers, and was adapted from existing validated instruments [29]. For interest in AI, example statements included *'I am interested in learning about AI'* and *'I want to learn more about AI outside of school'*. For AI careers awareness, example statements included *'I know about jobs that use AI'* and *'I am interested in jobs that use AI'*. Using data from 474 students, both scales obtained a Cronbach's alpha value of $\alpha = .939$ and $\alpha = .837$ respectively, indicating excellent internal consistency. Students were also asked to define AI via a free-text question (*'In the box below, write down what you think AI is'*).

The teacher survey focused on teachers' AI careers awareness and self-efficacy when teaching about AI via two 5-point Likert-type scales [27]. For AI careers awareness, example statements included *'I know about AI careers'* and *'I know where to find resources for teaching students about AI careers'*. For the *Personal AI Teaching Efficacy and Beliefs* scale, example statements included *'I am confident that I can explain to students how AI works'* and *'I understand AI concepts well enough to be effective in teaching about AI'*. Using data from 41 teachers, both scales obtained a Cronbach's alpha value of $\alpha = .921$ and $\alpha = .918$ respectively, indicating excellent internal consistency.

Data collection: Pre-lesson delivery data from 474 students and 41 teachers were collected through the survey. Correspondingly, post-lesson data was collected from 112 students and 6 teachers. Due to the disparities in sample sizes, any conclusions drawn should be treated with caution. Pre/post responses were matched where possible and appropriate statistical tests were employed. Further details are noted in the *Limitations* (Section 7.1). For qualitative data—such as student free-text questions—we decided to analyse all data currently collected as we intend to use a range of student responses to inform future work.

Data analysis: Due to the imbalance in sample sizes of the current data collected, we employed appropriate statistical techniques for analysing data. Student data that were normally distributed (assessed by Shapiro-Wilk's test, $p < .05$) were analysed using paired-samples t-tests, while non-normally distributed data were accordingly analysed using a series of Wilcoxon signed-rank tests. Similar tests were undertaken for teacher survey responses.

Principal component analysis (PCA) was used to assess inherent latent variables within the survey Likert-type items. The suitability of PCA was assessed prior to analysis (i.e. the assumptions of independent sampling, normality, linear relationships between pairs of variables, and the variables being correlated at a moderate level). For both surveys, PCA revealed components that were consistent with the underlying sub-scales, namely that there were strong loadings.

For the qualitative data of student definitions of AI (from the open-text question), responses were analysed by two researchers until agreement was reached. Using both pre- and post-data, we analysed responses based on whether their response ascribed anthropomorphic features when describing AI. We aimed to find evidence that students would attribute fewer human characteristics (e.g. emotions, consciousness, morality) after the lessons. Examples of anthropomorphic responses are shown in Table 1.

6 Results

Teacher survey: The teacher survey measured teachers' AI teaching self-efficacy and their AI careers awareness. Teachers' post-test *Personal AI Teaching Efficacy and Beliefs* scores were .70 higher than pre-test (95% CI, .22 to 1.18), a non-significant median increase, $z(3) = 1.46$, $p = .14$. Likewise, teachers' post-test *AI careers awareness* scores were .34 higher (95% CI, -.60 to 1.29) than pre-test scores, a non-significant mean increase, $t(3) = 1.15$, $p = .167$. For teachers' self-efficacy scores, the greatest mean differences were observed for multiple items relating to teacher confidence when teaching about AI such as *'I know the steps necessary to teach about AI effectively'* (+.86) and *'I am confident that I can answer students' questions about AI'* (+.73). Likewise, the greatest mean differences for AI careers awareness were observed across multiple items including *'I know about AI careers'* (+.80) and *'I know where to go to learn more about AI careers'* (+.80). These results suggest that some teachers who took part in this programme may have felt more aware of AI careers, and some may have felt more confident in their AI teaching skills after completing these programmes.

Student survey: The student survey measured students' attitudes and awareness of AI/AI-related careers. Only a small amount of paired data was available. Post-test *interest in AI* scores were .09 higher (95% CI, -.16 to .34) than pre-test scores, a non-significant increase ($t(30) = 1.659$, $p = .24$). However, post-test *AI careers awareness* scores were .73 (95% CI, .42 to 1.03) higher than pre-test, a statistically significant increase ($t(30) = 4.919$, $p = .05$). In particular, group differences in certain items (e.g. *'I know someone like me who works in an AI-related field'* (+0.55) and *'I plan to study AI after secondary school'* (+0.40)) suggest that student awareness and interest in AI careers may have been impacted by their participation in the lessons. Additional items that had large mean differences centred on student awareness of AI-related jobs (+0.42) and discussion of jobs with families and friends (+0.46).

Table 1. Anthropomorphism in student responses to "write what you think AI is"

AI description	Examples	Pre #	%	Post #	%
Anthropomorphic	*"AI is a robot which is starting to become more human like with feelings" / "A computer that has a mind of its own and can think its own thoughts."*	78	16.46%	12	10.71%
Non-anthropomorphic	*"A simulation of intelligence made by humans [...] It has no actual intelligence" / "The development of a concept or tool created to mimic the intelligence of humans"*	396	83.54%	100	89.29%

Student conceptions of AI: Students were asked to define AI before and after taking part in the lessons. We found evidence that most students demonstrated non-anthropomorphic descriptions (e.g. *"AI simulates human behaviour in machines which helps [in] tasks such as problem solving"*) in both pre- and post-responses (83.54% and 89.29%, respectively). Students gave proportionally fewer anthropomorphic answers (e.g. *"[AI is] someone who helps [you] online"*) in post-test data (16.46% vs 10.71%) (see Table 1).

7 Discussion

Previous work has detailed the extent of resources to support teaching/learning about AI/ML [26]. However, most resources were found to not include specific learning objectives, recommended age groups, or assessment materials. Likewise, a lack of common vocabulary underpinning learning material was noted [3]. Our curriculum, Experience AI, represents a significant attempt to provide educators (and researchers) with a research-informed set of teaching materials, including teacher guides, classroom presentations, explanations of key terms, student activities and assessment ideas. Evidence from the independent evaluation provided initial evidence that this approach could support teachers' AI career awareness and efficacy when teaching about AI. We discuss these results in relation to our original design principles and relate this to prior literature.

Avoiding anthropomorphisation: Design decisions were taken to avoid anthropomorphisation in language use, however, students were not explicitly taught about the topic. Results from the independent evaluation found proportionally fewer uses of anthropomorphic language, though this was not significant given the unbalanced sample. This result highlights a challenge in understanding student awareness from written responses and also that, without explicit teaching, the change in language and imagery and motivation for this was not transparent to students. Previous work in computing education has suggested that the use of natural language (e.g. metaphors) may lead to naive preconceptions [15]. It is possible that the use of anthropomorphic language (e.g. analogies to the human brain) may serve as a scaffolding measure to aid understanding. Follow-up work could focus on qualitative analysis of student perspectives of typical AI characteristics. This could include scenario-based tasks to discover whether more robust mental models [23] or fewer problematic associations of gender, race, over-reliance or an increased view of human responsibility for AI design [2, 1, 19, 28] are developed through avoiding—or discussing the limitations of—anthropomorphisation. While differences in student use of anthropomorphic language could not be identified, we were instead provided with important insights into students' emerging understanding of AI technologies.

Promoting AI/AI-related career awareness: The curriculum unit design featured a strong emphasis on the social and ethical dimensions of AI. As in the case of the DAILy curriculum [6], we sought to raise student awareness of careers and the extent to which AI features in their everyday lives. For instance, the role AI could play in students' careers both directly—working within the field—and indirectly—impacted by AI—was explored through real-world examples. Findings from the survey data collected were limited though suggested that both students and teachers improved their awareness of AI and AI-related careers. We were encouraged to observe mean differences for student responses of appropriate Likert items *'I know about jobs that use AI'* and *'I plan to study AI after secondary school'*, as these suggest that our AI career activities both illustrated a broad range of careers (research scientists, robotics engineers, ethics researchers employed at Google DeepMind) and promoted interest among participants. Students also compared different career pathways against multiple dimensions of AI (e.g. social and ethical, creating applications and tools, training models) to demonstrate the implications of AI in a broad range of fields. Embedding career pathways in lesson materials was also intended to support teachers. As gatekeepers to facilitate student awareness of AI careers [27], it was encouraging to see large mean differences in teachers' awareness of careers and career resources. However, further work is required to better understand which resources were impactful and how teachers' and students' perceptions of AI/AI-related careers were influenced.

Teacher support: Finally, teacher support was a focus for Experience AI. Following prior work that suggests a lack of suitable documentation (e.g. learning outcomes, differentiation, and assessment activities) in most currently-available resources [26], we focused on supporting teachers through teaching materials, online courses, webinars, and student-facing materials. Survey data from the independent evaluation found early evidence that this approach could support their efficacy when teaching about AI. In particular, teachers felt most confident knowing the steps to teach about AI and receiving and answering student questions. In future work, we could seek to better understand what specific measures could support teachers implementing the curriculum in schools.

7.1 Limitations

Our findings are limited to the extent that the independent evaluation was conducted separately to the lesson design and development. As such, only the impact of some of our design principles could be investigated in light of limited data collected (e.g. teacher self-efficacy). Nonetheless, findings from the independent evaluation were encouraging and provided some early indication of the impact of teachers' and students' participation in Experience AI. Survey data were collected anonymously meaning that no additional demographic data were collected beyond age, year group (if applicable), gender and ethnicity. This meant that follow-up work could not be organised, such as focusing on more in-depth qualitative analysis of the student and teacher experience when taking part in Experience AI. This would have provided extra insights into student reasoning around AI where survey data is limited. For example, one issue with the open-text question was that students might have interpreted it to mean what future capabilities *could* be achieved using AI technologies as opposed to its current functionalities and limitations. In-depth follow-up work with students is required to differentiate between these two perspectives. Finally, non-significant increases in student and teacher attitudes merely suggest effects and necessitate further data collection and analysis to rule out chance effects and provide confirmatory evidence.

8 Conclusion

This paper has described some of the key design principles of six lessons that can be used to introduce AI/ML to Grade 6–8 students in the UK. The design principles foregrounded in our lesson unit include avoiding anthropomorphisation, embedding careers and increased teacher support. Evidence from the independent evaluation indicates that student and teacher AI careers awareness may have been positively impacted. Evidence has also provided early yet encouraging evidence that teachers who taught the curriculum unit may have had higher self-efficacy when teaching about AI. By providing the full resource set (see Appendix A), we propose that the materials may be useful to educators new to AI/ML as well as other researchers.

Further work is needed to investigate whether the content covered and design decisions, such as teaching about the difference between rule-based and data-driven systems, are the most effective way to develop useful mental models. Similarly, whether starting with decision trees, rather than neural networks or other machine learning engines, is the most effective way to establish an effective progression of knowledge building. Also, whether the instructional approaches used sufficiently scaffold learning for all, or if they kept the nuances of ML too hidden to help students overcome either too much or too little trust in the predictions of ML models requires further research. We look forward to exploring these issues and hope that our resources will be useful to others to investigate.

Acknowledgments. We are grateful to Google DeepMind for funding the development and evaluation of this programme.

References

1. Cave, S., Dihal, K.: The Whiteness of AI. Philosophy & Technology **33**(4), 685–703 (Dec 2020). https://doi.org/10.1007/s13347-020-00415-6, https://link.springer.com/10.1007/s13347-020-00415-6
2. Druga, S., Ko, A.J.: How do children's perceptions of machine intelligence change when training and coding smart programs? In: Interaction Design and Children. pp. 49–61. ACM, Athens Greece (Jun 2021). https://doi.org/10.1145/3459990.3460712
3. Druga, S., Otero, N., Ko, A.J.: The landscape of teaching resources for AI education. In: Proceedings of the 27th ACM Conference on on Innovation and Technology in Computer Science Education Vol. 1. pp. 96–102. ITiCSE '22, Association for Computing Machinery, New York, NY, USA (Jul 2022)
4. Galindo, L., Perset, K., Sheeka, F.: An overview of national AI strategies and policies. OECD Going Digital Toolkit Notes **14** (2021). https://doi.org/https://doi.org/https://doi.org/10.1787/c05140d9-en
5. Hur, J.W., Andrzejewski, C.E., Marghitu, D.: Girls and computer science: experiences, perceptions, and career aspirations. Computer Science Education **3408**(October), 100–120 (2017)
6. Lee, I., Ali, S., Zhang, H., DiPaola, D., Breazeal, C.: Developing middle school students' AI literacy. In: Proceedings of the 52nd ACM Technical Symposium on Computer Science Education. pp. 191–197. SIGCSE '21, Association for Computing Machinery, New York, NY, USA (Mar 2021)
7. Long, D., Magerko, B.: What is AI Literacy? Competencies and Design Considerations. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. pp. 1–16. ACM, Honolulu HI USA (Apr 2020). https://doi.org/10.1145/3313831.3376727
8. Marques, L.S., Gresse Von Wangenheim, C., Hauck, J.C.R.: Teaching machine learning in school: A systematic mapping of the state of the art. Informatics in Education pp. 283–321 (Jun 2020)
9. Maslej, N., Fattorini, L., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Ngo, H., Niebles, J.C., Parli, V., Shoham, Y., Wald, R., Clark, J., Perrault, R.: The AI Index 2023 Annual Report (2023)
10. Maton, K.: Making semantic waves: A key to cumulative knowledge-building. Linguistics and Education **24**(1), 8–22 (Apr 2013)
11. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D., Gebru, T.: Model Cards for Model Reporting. In: Proceedings of the Conference on Fairness, Accountability, and Transparency. pp. 220–229 (Jan 2019). https://doi.org/10.1145/3287560.3287596, arXiv:1810.03993 [cs]
12. Olari, V., Romeike, R.: Addressing AI and Data Literacy in Teacher Education: A Review of Existing Educational Frameworks. In: The 16th Workshop in Primary and Secondary Computing Education. WiPSCE '21, ACM, New York, NY, USA (2021). https://doi.org/10.1145/3481312.3481351
13. Pea, R.D.: Language-independent conceptual "bugs? in novice programming. Journal of educational computing research **2**(1), 25–36 (1986)
14. Pedro, F., Subosa, M., Rivas, A., Valverde, P.: Artificial intelligence in education: Challenges and opportunities for sustainable development. Tech. rep., UNESCO (2019)
15. Qian, Y., Lehman, J.: Students' Misconceptions and Other Difficulties in Introductory Programming: A Literature Review. ACM Trans. Comput. Educ. **18**(1), 1:1–1:24 (2017)

16. Raspberry Pi Foundation: Experience AI design principles study (2024), rpf.io/ai-design-principles-study
17. Rizvi, S., Waite, J., Sentance, S.: Artificial Intelligence teaching and learning in K-12 from 2019 to 2022: A systematic literature review. Computers and Education: Artificial Intelligence **4**, 100145 (2023). https://doi.org/https://doi.org/10.1016/j.caeai.2023.100145
18. Royal Society: Machine learning: the power and promise of computers that learn by example. Royal Society (2017), https://royalsociety.org/-/media/policy/projects/machine-learning/publications/machine-learning-report.pdf
19. Salles, A., Evers, K., Farisco, M.: Anthropomorphism in AI. AJOB Neuroscience **11**(2), 88–95 (Apr 2020). https://doi.org/10.1080/21507740.2020.1740350, https://www.tandfonline.com/doi/full/10.1080/21507740.2020.1740350
20. Sanusi, I.T., Oyelere, S.S., Agbo, F.J., Suhonen, J.: Survey of Resources for Introducing Machine Learning in K-12 Context. In: 2021 IEEE Frontiers in Education Conference (FIE). pp. 1–9. IEEE, Lincoln, NE, USA (Oct 2021). https://doi.org/10.1109/FIE49875.2021.9637393, https://ieeexplore.ieee.org/document/9637393/
21. Select Committee on AI: AI in the UK: ready, willing and able? House of Lords (2018), https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf
22. Tedre, M., Denning, P., Toivonen, T.: CT 2.0. In: 21st Koli Calling International Conference on Computing Education Research. Koli Calling '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3488042.3488053
23. Tedre, M., Toivonen, T., Kahila, J., Vartiainen, H., Valtonen, T., Jormanainen, I., Pears, A.: Teaching Machine Learning in K-12 Computing Education: Potential and Pitfalls (2021). https://doi.org/10.48550/ARXIV.2106.11034, https://arxiv.org/abs/2106.11034
24. Touretzky, D., Gardner-McCune, C., Martin, F., Seehorn, D.: Envisioning AI for K-12: What Should Every Child Know about AI? Proceedings of the AAAI Conference on Artificial Intelligence **33**(01), 9795–9799 (Jul 2019). https://doi.org/10.1609/aaai.v33i01.33019795
25. UK AI Council: AI Roadmap. Nuffield Foundation (2019), https://assets.publishing.service.gov.uk/media/5ff3bc6e8fa8f53b76ccee23/AI_Council_AI_Roadmap.pdf
26. Waite, J., Tshukudu, E., Cucuiat, V., Whyte, R., Sentance, S.: Towards a Framework for Learning Content Analysis in K-12 AI/ML Education. In: 2023 IEEE Frontiers in Education Conference (FIE). pp. 1–5. IEEE Computer Society, Los Alamitos, CA, USA (Oct 2023). https://doi.org/10.1109/FIE58773.2023.10343368
27. Whyte, R., Kirby, D., Sentance, S.: Measuring Teacher Self-Efficacy and Careers Awareness in K-12 AI Education. In: Proceedings of the 2024 ITiCSE, July 8–10, 2024, Milan, Italy. ITiCSE '24, ACM, New York, NY, USA (Oct 2024). https://doi.org/10.1145/3649405.3659489
28. Williams, R., Ali, S., Devasia, N., DiPaola, D., Hong, J., Kaputsos, S.P., Jordan, B., Breazeal, C.: AI + ethics curricula for middle school youth: Lessons learned from three Project-Based curricula. Int J Artif Intell Educ pp. 1–59 (Aug 2022)
29. Zhang, H., Lee, I., Ali, S., DiPaola, D., Cheng, Y., Breazeal, C.: Integrating ethics and career futures with technical learning to promote AI literacy for middle school students: An exploratory study. Int J Artif Intell Educ pp. 1–35 (May 2022)
30. Zhou, X., Brummelen, J.V., Lin, P.: Designing AI Learning Experiences for K-12: Emerging Works, Future Opportunities and a Design Framework (2020)

A Appendix 1

All resources are free to use by anyone now and in perpetuity under a Creative Commons license (CC BY-NC-ND 4.0). The full resource set is available via the study website [16].

A.1 Lesson 1 - What is AI?

Students explore the current state of artificial intelligence (AI), and how it is used in the world. They explore the difference between rule-based systems and data-driven models (Figure 1) and consider the benefits and drawbacks of AI systems. The learning objectives covered are: (i) Describe the difference between ‘data-driven’ and ‘rule-based’ approaches to application development; (ii) Name examples of AI applications; (iii) Outline some benefits and issues of using AI applications.

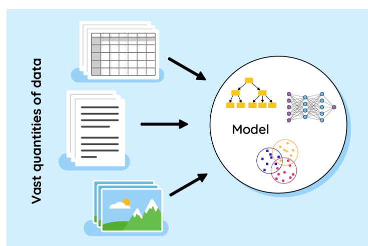

Fig. 1. Figure used to illustrate data-driven models.

A.2 Lesson 2 - How computers learn from data

The activities in this lesson help students think critically about which parts of a system use AI components and the role of ML models. Through a video, students hear from experts about the different types of ML and example problems solved. Students are introduced to a specific example of ML: classification. The learning objectives are: (iv) Define machine learning’s relationship to artificial intelligence; (v) Name the three common approaches to machine learning; (vi) Describe how classification can be solved using supervised learning.

A.3 Lesson 3 - Bias in, bias out

Students create an ML model using Machine Learning for Kids[5]. The model classifies images of apples and tomatoes (Figure 2), but students discover that

[5] https://machinelearningforkids.co.uk/

their model is flawed due to the limited training data set. Next, students explore training data bias and biased predictions. Learning objectives include: (vii) Describe the impact of data on the accuracy of a machine learning (ML) model; (viii) Explain the need for both training and test data; (ix) Explain how bias can influence the predictions generated by an ML model.

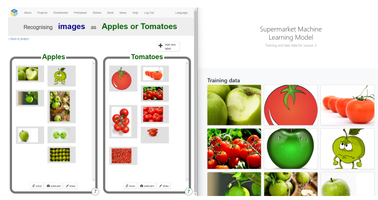

Fig. 2. Slides showing adding training data.

A.4 Lesson 4 - Decision trees

Students take their first in-depth look at a type of engine: decision trees. The activities build on students' learning from Lessons 1–3 about classification, training and test data, and the data-driven nature of models. The aim is for students to gain an understanding of the processes used to create ML models. Learning objectives: (x) Describe how decision trees are used to build a classification ML model; (xi) Describe how training data changes an ML model; (xii) Explain why ML is used to create decision trees.

A.5 Lesson 5 - Solving problems with ML models

Students are introduced to the AI project lifecycle. They follow the stages to create an ML model to solve a problem of their choice from example projects. They train the model and test it to determine its accuracy. (xiii) Describe the stages of the AI project lifecycle; (xiv) Use a machine learning tool to import data and train a model; (xv) Test and examine the accuracy of an ML model.

A.6 Lesson 6 - Model cards and careers

In this lesson, students complete the final stages of the AI project lifecycle: evaluating and explaining a model. To help them explain their model, students are introduced to model cards [11]. (xvi) Evaluate an ML model; (xvii) Produce a model card to explain an ML model; (xviii) Recognise the range of opportunities that exist in AI-related careers. The main instructional approaches used are active learning through discussion, project-based learning, and to provide student choice. Students conclude by exploring careers both in AI and fields in which it is used.