

DESIGN PRINCIPLES AND RESOURCES FOR INTRODUCTORY AI LESSONS FOR 11- TO 14-YEAR-OLD LEARNERS

Jane Waite ¹ & Sue Sentance ²

¹ Raspberry Pi Computing Education Research Centre, University of Cambridge, UK,
jw2251@cam.ac.uk

² Raspberry Pi Computing Education Research Centre, University of Cambridge, UK

Focus Topics: Learning Materials

Introductory AI lessons - overview

Introduced in 2023, Experience AI is a free educational programme. The materials were developed by the Raspberry Pi Foundation in collaboration with Google DeepMind and aim to introduce AI and Machine Learning (ML) to 11- to 14-year-old students (<https://experience-ai.org/>). The programme of resources and support consists of six one-hour classroom lessons (a unit of work), additional one-off lessons (e.g., for teaching about AI in a biology lesson), a supplementary challenge, and online teacher professional development.

Although the resources were initially developed for teachers in England, they are available for any educator who teaches any age group in any country. For example, teachers may adapt resources for younger or older learners in their context. Also, for an increasing number of countries, there are translated and localised resources (e.g. for Romania <https://experience-ai.org/ro> and for Kenya <https://experience-ai.org/sw>) supported by in-country partners (<https://experience-ai.org/en/partners>).

There are no specific hardware or software requirements. The web app Machine Learning for Kids (Lane, 2021) is used for practical activities in lessons across the resources.

Lesson resources include lesson plans, teaching slides, student activity worksheets, student explanation videos, student projects linked to personal interests, and student assessment activities (formative and summative). There are also teacher support materials, including a set of learning graphs (that detail the order of learning objectives (LOs) in a diagram), a set of explanations of key and sub-concepts, a unit overview, a set of teacher support videos and webinars, and an online teacher professional development course. In each lesson, LOs are shared with students, an active learning approach is used with hands-on and collaborative activities to facilitate rich discussion, and case studies are provided to link learning to real contexts. Key concepts introduced in the lessons include rule-based vs data-driven approaches to programming, AI applications, ML models, bias and ethics, decision trees, the AI data life cycle, and careers in AI. Sub-concepts include (amongst others) prediction, supervised learning, unsupervised learning, reinforcement learning, ML classification, ML training, ML explainability, and ML model cards. As this is an introductory set of lessons for school-aged learners and is likely to be delivered by educators who are not AI specialists, the depth of knowledge building of concepts and sub-concepts varies. Some concepts/sub-concepts are introduced as core vocabulary to situate learning. In contrast, others demonstrate the breadth of ML or are introduced to avoid misconceptions forming. For example, the term ML model is key to situate learning about related sub-concepts such as types of ML learning and ML training. More technical concepts are addressed at a simple level; for example, only decision trees are introduced as a way for ML model implementation to start to be explained. What misconceptions might form as we start to teach about ML needs further research. Still, in anticipation, the idea of rule-based vs data-driven was introduced to address potential misconceptions around the difference between existing program design approaches and newer ML approaches (Section 2.5).

An independent evaluation of the Experience AI initiative is being conducted, with some early reporting indicating promising results (Waite et al., 2024)

Design principles

A set of design principles underpin the resources, including careful learning objective design, using the SEAME model (Waite et al., 2023), avoiding anthropomorphisation, developing a core vocabulary, incorporating Computational Thinking 2.0 (Tedre et al., 2021a), threading real-world examples throughout, using semantic profiling to link everyday to technical language and everyday

contexts to abstract understanding (Maton, 2013) for explanations and lesson design, and increased teacher support materials.

Learning objectives design

Which LOs were to be covered by the unit of work was informed by a review in 2021/2022 of AI and ML teaching and learning resources which found a lack of cohesive LOs, lack of progression of LOs, and a lack of assessment in AI learning resources (Waite et al., 2023). For the Experience AI lessons, LOs were derived from a synthesis in June 2022 of stated and implied learning objectives of AI teaching materials from CBSE, Code.org, MITRaise, ENARIS, Apps4Good, AI4K12, the Technovation project, Beverly Clarke resources, Machine Learning 4 Kids and Teaching London Computing (Raspberry Pi Foundation, 2023). It was clear after LOs synthesis that this initial set of LOs was not complete nor entirely cohesive, but for the intended age group of the unit of work, this set of common emerging LOs provided a starting point to develop progression. Bloom's taxonomy, the SEAME framework (Waite et al., 2023) and learning graphs were used to categorise, balance, and sequence the LOs (Raspberry Pi Foundation, 2023). Whether the LOs used are ideal requires further research, but they provided a good starting point.

SEAME framework

The SEAME framework was found to be a useful way for resource developers and researchers to categorise learning objectives and lesson activities. SEAME has been used to review AI teaching resources (Waite et al., 2023) and research activities (Rizvi et al., 2023) and comprises four levels: Social and Ethical, Application, Model and Engine. The levels do not dictate the order of learning, and some activities will span more than one level. It may be that as students make progress, they will become more able to move between levels, and, at times, the boundaries between levels will blur.

Avoiding anthropomorphisation

A firm and early design decision was to avoid anthropomorphisation in student- and teacher-facing material for language and images. For example, avoiding using illustrations that depicted devices with human-like faces and replacing vocabulary that was associated with the behaviour of people (see, look, recognise, create, make) with system-type words (detect, input, pattern match, generate, produce). The rationale for this choice was that attributing human characteristics to computers has led to programming misconceptions (e.g. Pea (1986)), 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 (Tedre et al., 2021b). 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 (e.g. Druga and Ko (2021), Williams et al.(2019)), or developing relationships with the devices (Druga et al., 2017; Tanaka et al., 2007), leading to high risks of either unintended influence, purposeful manipulation, or policing (Vollmer et al., 2018; Williams et al., 2022, 2018). Compounding the issue of delegating the responsibility of the human system developer and human user to an imagined responsible AI agent (Salles et al., 2020), anthropomorphised AI agents have been predominately portrayed as white in colour and, as such, exacerbating racism in technology and society at large (Cave and Dihal, 2020). However, by engaging students to learn about AI and making the box transparent, students become more sceptical about the technology and recognise responsibility for AI design and use as associated with the humans who are 'in the loop' rather than AI technology (Druga and Ko, 2021; Tedre et al., 2021b).

Core vocabulary and semantic waves

In line with calls for a common AI vocabulary (Druga et al., 2022), we developed long explanations of key concepts to provide a shared understanding of the topic by resource developers and teachers (Raspberry Pi Foundation, 2023). Each explanation was carefully constructed and included sources that informed the text, and justified the choices of specific terms, and followed a semantic wave to aid knowledge building (Maton, 2013).

Computational Thinking (CT) 2.0

The resources were initially developed with 11- to 14-year-old students in England in mind, who will have been taught to develop programs using an algorithmic/rule-based approach. To take this

prior learning into account, a CT1.0 vs CT2.0 view (Tedre et al., 2021a) was incorporated in the resources to start to address potential misconceptions of rule-based vs data-driven notional machines for ML systems. However, further research is needed to explore whether this approach has succeeded.

Relatable real-world case studies to exemplify social and ethical issues

A design choice was made to provide examples in each lesson that shared social and ethical issues through relatable real-world examples of applications of ML models and AI careers. The aim was to engage students and help them understand the relevance and impact AI has in their lives by providing a range of different contexts and opportunities to discuss the social and ethical implications.

Acknowledgements

We would like to thank Google DeepMind for funding the development of the Experience AI resources. Particular thanks are made to the Google DeepMind working group.

References

- Cave, S. and Dihal, K. (2020). The Whiteness of AI. *Philosophy & Technology*, 33(4):685–703.
- Druga, S. and Ko, A. J. (2021). How do children’s perceptions of machine intelligence change when training and coding smart programs? In *Interaction Design and Children*, pages 49–61, Athens Greece. ACM.
- Druga, S., Otero, N., and Ko, A. J. (2022). 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, ITiCSE ’22*, pages 96–102, New York, NY, USA. ACM.
- Druga, S., Williams, R., Breazeal, C., and Resnick, M. (2017). ”Hey Google is it OK if I eat you?”: Initial Explorations in Child-Agent Interaction. In *Proceedings of the 2017 Conference on Interaction Design and Children*, pages 595–600, Stanford California USA. ACM. 00000.
- Lane, D. (2021). *Machine Learning for Kids: A Project-Based Introduction to Artificial Intelligence*. No Starch Press, San Francisco, US.
- Maton, K. (2013). Making semantic waves: A key to cumulative knowledge-building. *Linguistics and Education*, 24(1):8–22.
- Pea, R. D. (1986). Language-independent conceptual “bugs?” in novice programming. *Journal of educational computing research*, 2(1):25–36.
- Raspberry Pi Foundation (2023). Experience AI research website. [/rpf.io/ai-design-principles-study](https://rpf.io/ai-design-principles-study).
- Rizvi, S., Waite, J., and Sentance, S. (2023). Artificial Intelligence teaching and learning in K-12 from 2019 to 2022: A systematic literature review. *Computers and Education: Artificial Intelligence*
- Salles, A., Evers, K., and Farisco, M. (2020). Anthropomorphism in AI. *AJOB Neurosc*, 11(2):88–95.
- Tanaka, F., Cicourel, A., and Movellan, J. R. (2007). Socialization between toddlers and robots at an early childhood education center. *Proceedings of the National Academy of Sciences*, 104(46):17954–17958.
- Tedre, M., Denning, P., and Toivonen, T. (2021a). CT 2.0. In *21st Koli Calling International Conference on Computing Education Research, Koli Calling ’21*, New York, NY, USA. ACM.
- Tedre, M., Toivonen, T., Kahila, J., Vartiainen, H., Valtonen, T., Jormanainen, I., and Pears, A. (2021b). Teaching Machine Learning in K-12 Computing Education: Potential and Pitfalls.
- Vollmer, A.-L., Read, R., Trippas, D., and Belpaeme, T. (2018). Children conform, adults resist: A robot group induced peer pressure on normative social conformity. *Science Robotics*, 3(21).
- Waite, J., Garside, B., Whyte, R., Kirby, D., and Sentance, S. (2024). Experience AI: Introducing AI/ML to Grade 6–8 students in the UK. In *Proceedings of the 2024 17th International Conference on Informatics in Schools, ISSEP ’24*. Alpen-Adria Universit t Klagenfurt.
- Waite, J., Tshukudu, E., Cucuiat, V., Whyte, R., and Sentance, S. (2023). Towards a framework for learning content analysis in K-12 AI/ML education. *Frontiers in Education*.
- Williams, R., Ali, S., Devasia, N., DiPaola, D., Hong, J., Kaputsos, S. P., Jordan, B., and Breazeal, C. (2022). AI + ethics curricula for middle school youth: Lessons learned from three Project-Based curricula. *Int J Artif Intell Educ*, pages 1–59.

- Williams, R., Machado, C. V., Druga, S., Breazeal, C., and Maes, P. (2018). "My doll says it's ok": a study of children's conformity to a talking doll. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*, pages 625–631, Trondheim Norway. ACM.
- Williams, R., Park, H. W., and Breazeal, C. (2019). A is for Artificial Intelligence: The Impact of Artificial Intelligence Activities on Young Children's Perceptions of Robots. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–11, Glasgow Scotland UK. ACM.