

Learn to Machine Learn: Designing a Game Based Approach for Teaching Machine Learning to Primary and Secondary Education Students

Iro Voulgari
Institute of Digital Games
University of Malta
Msida, Malta
iro.voulgari@um.edu.mt

Marvin Zammit
Institute of Digital Games
University of Malta
Msida, Malta
marvin.zammit@um.edu.mt

Elias Stouraitis
Palladio School
Athens, Greece
ilias.stouraitis@palladio.edu.gr

Antonios Liapis
Institute of Digital Games
University of Malta
Msida, Malta
antonios.liapis@um.edu.mt

Georgios N. Yannakakis
Institute of Digital Games
University of Malta
Msida, Malta
georgios.yannakakis@um.edu.mt

ABSTRACT

With the ubiquitous role of Artificial Intelligence (AI) in everyday applications such as smartphones and social media, children need digital literacy skills to navigate the digital world, critically view, and reflect on the social and ethical implications of the design and architecture of AI systems. To address this increasing need for AI literacy skills, particularly for younger students, this paper presents the rationale of the LearnML project which aims to develop a framework and game-based educational material for promoting AI literacy among primary and secondary education students. We also describe the design and initial assessment of the game “ArtBot”, developed as part of the LearnML project. We review existing literature, discuss the educational game design and development of “ArtBot” and describe the initial feedback of students and teachers. Our goal is to provide insights and suggest guidelines for the implementation of game-based learning environments for supporting AI literacy skills of students.

CCS CONCEPTS

• **Applied computing** → **Computer-assisted instruction**; • **Social and professional topics** → **Computational thinking**; **K-12 education**; **Computing literacy**; **Information technology education**; **Computer science education**; **Children**; **Adolescents**.

KEYWORDS

Game-based learning, AI literacy, AI education, machine learning, digital literacy skills, primary education, secondary education, educational game design

ACM Reference Format:

Iro Voulgari, Marvin Zammit, Elias Stouraitis, Antonios Liapis, and Georgios N. Yannakakis. 2021. Learn to Machine Learn: Designing a Game Based Approach for Teaching Machine Learning to Primary and Secondary Education Students. In *Interaction Design and Children (IDC '21)*, June 24–30, 2021, Athens, Greece. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3459990.3465176>

1 INTRODUCTION

Artificial Intelligence (AI) and Machine Learning (ML) are already ubiquitous in everyday life in applications such as speech and image recognition, smartphones, smart toys, recommendation systems, personalised information on social media and search engines, and autonomous vehicles. The term AI describes the general field of algorithms and processes by which an application performs tasks that seem to mimic human behaviour or intelligence. ML is a subset of these algorithms through which a computing process progressively adapts and improves its performance in a specific task or set of tasks [20]. Besides the value and benefits of AI, concerns have been raised on the potential pitfalls and challenges regarding the decisions such systems might make, the autonomy of the AI agents, and the values and bias embedded in their design [11]. In this context, children and youth have to develop digital literacy skills to navigate the digital space and cope with this major shift in the way information is accessed, processed, and distributed. Students need to be able to understand how computational systems work and their value for societal issues [1] and also develop the skills to access, critically analyse and interpret data and information, recognise emerging cultural and social biases embedded in the design of computer systems, and the ethical and political implications [6, 7, 18]. Such skills and competences can empower students to adopt a more critical and inquisitive approach towards existing technological environments and to engage more actively in the design of new technology systems [15].

In this context, an emphasis is placed over the past few years on AI and ML literacy, addressing not only the technical aspects of AI and ML but also on the ethical and cultural implications. AI literacy addresses not only Computer Science graduate students, but also

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
IDC '21, June 24–30, 2021, Athens, Greece
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8452-0/21/06.
<https://doi.org/10.1145/3459990.3465176>

primary and secondary education students with varied degrees of expertise in AI, ML, and computer programming. Despite the inherent challenges of teaching such abstract and complex concepts to children [16], a number of successful approaches have been proposed for early AI literacy. An example is the Artificial Intelligence Curriculum [21, 22] designed to empower early childhood children to understand how AI devices from their everyday experience work, main AI concepts such as labelling and classification, and discuss issues such as the meaning of intelligence. Another AI curriculum for middle school students [8] aimed at raising awareness on technical concepts of AI and their ethical implications, such as algorithmic bias. As another example, card-based design workshops for students aged 16-20 [2] were carried out to support exploration and reflection on the moral decisions ML systems are designed to make. Such workshops and approaches encourage insightful discussions about ethics and ML technology, rather than focus on the technical aspects and the design of a functional prototype [2].

Situated in this context, the “Learn to Machine Learn” (LearnML) project aims to develop an AI education framework and toolbox for supporting primary and secondary education students and teachers, including a game-based learning toolbox and relevant teacher training, teaching and learning material. Aiming to promote a culture of openness and participatory design, the project involved educators and students from the early stages of the design of the learning goals and the educational material, scenarios, and games, through focus groups and workshops. Our main goal in this paper is to examine and provide insights on the game design and development considerations for effective implementation of games for AI education. We present and review the design and initial assessment by students and teachers of the digital game “ArtBot”, developed as part of the LearnML education toolbox.

2 GAMES AND APPLICATIONS FOR AI AND ML EDUCATION

Digital games seems to be an appropriate tool for approaching the complex concepts and processes involved in AI systems [3]. A good number and a wide range of games, applications, and platforms seem to be available for children as young as 4 years old, addressing different aspects of AI and ML, such as technical concepts, and ethical and social implications [5]. In a study by Zimmermann-Niefeld et al. [23], middle and high school students developed ML models based on data from their athletic activities, linking their existing interests with their ML model. Students aged 12-13 developed their own models using Google Teachable Machine (GTM) [14] demonstrating that GTM was a feasible tool for young students with little or no experience in programming. Diverse audiences (high school students, university students, law enforcement officers and park rangers) were introduced to ML concepts through an online and a board role playing game, situated in a real-world problem (the protection of wildlife). Participants reported positive attitudes toward the games and increased interest in ML [13]. Similarly, Vartiainen et al. implemented design-based workshops, grounded on Papert’s constructionist approach where 12-13 year old students designed their own ML applications. The workshops were grounded on real-world problems close to the students’ experience and interests,

using Google Teachable Machine and the researchers’ own educational application for object recognition [17]. The results were promising; although the students’ conceptions of ML seemed to be closely linked to their own applications and the tools used, the workshops were a good entry point for exploring ML concepts, exhibit empathy about other people’s needs, and engage in inductive reasoning about the quality of their datasets and accuracy of their models. Druga et al. [4] focused on children aged 7-12 and concluded that they were able to engage in meaningful and creative ways with AI agents, such as the Alexa home assistant and the Cozmo robot.

AI and ML education tools tend to follow two directions: a) open-ended tools and platforms where students can design and create their own models and b) digital games scaffolding students through gameplay. These two directions may complement each other linking scaffolding with design-based pedagogy [5]. Existing studies suggest that such tools, when used in an appropriate pedagogical context and learning activities, may help students understand principles and processes of AI and ML and engage in meaningful reflection and discussions. Another critical aspect that emerged is the importance of engaging students through meaningful activities, close to their experience and interests. When students are more engaged, they explore the role and implications of AI inductively and make inferences about its potential value and limitations.

In this context, we aimed to integrate elements of these two directions in our game design, by guiding the players and also providing opportunities for exploration and experimentation. A meaningful story and narrative that would trigger the interest of the students was also important for our design, so as to motivate and engage them with the game.

3 A METHODOLOGY FOR DESIGNING A GAME-BASED APPROACH TO TEACH MACHINE LEARNING

In this section we discuss the design, development, and initial feedback on the “ArtBot” game from students and teachers, and draw insights from this experience. Prior to the design of the game-based educational material, discussions and focus groups with teachers and students were conducted, where we examined their perceptions and potential misconceptions about AI, their ideas about how AI could be introduced into formal education, and potential educational scenarios and activities. Some of the main points that emerged from these focus groups were:

- (1) Games selected and developed should be easy to install and use even in low-end systems. Some of the teachers reported very limited hardware and software resources in schools.
- (2) AI behaviour and processes should be made as clear as possible to the students and teachers; they should be able to understand how an AI system works.
- (3) Ethical issues and concerns about the role and implications of AI in everyday life and society should be addressed.
- (4) The educational material (educational scenarios, resources, games) should be close to the students’ experience.
- (5) Educational material should allow for interdisciplinary teaching approaches, and be applicable to different fields such as computing, history, arts, language, and literature.

- (6) Games and activities would be easier to implement in the classroom if they correspond to existing learning objectives of the curriculum.

We tried to address these points, as will be further discussed in the next section, by designing, for example, a game with low technical requirements, by defining appropriate learning goals regarding the technical and ethical aspects of AI, and by implementing a meaningful narrative that would allow for interdisciplinary teaching approaches.

3.1 ArtBot: Educational Game Design

“ArtBot” was designed by a team of educators, game developers, and AI experts with the aim to support AI literacy of primary and secondary education students. Our first step was to identify the learning objectives and core ML principles to be introduced through the game. The players should be able to:

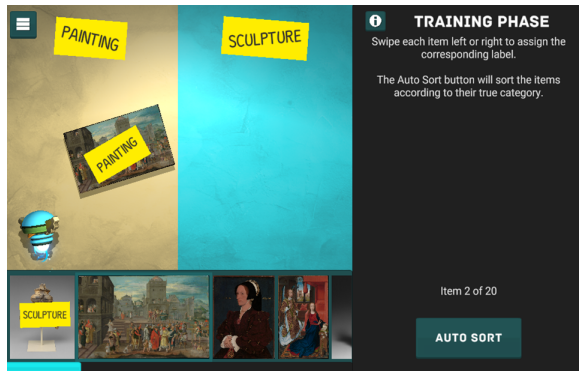
- (1) Understand the process of *Supervised Learning*: students are introduced to the concepts of training datasets, testing datasets, classification, labelling, image recognition, decision trees, and prediction accuracy, and understand the role of these elements in the architecture and behaviour of an AI system.
- (2) Understand the process of *Reinforcement Learning*: students are introduced to the concepts or rewards and penalties, learning duration, learning rate, exploration, exploitation, and pathfinding.
- (3) Understand the impact of the decisions made during the design and programming of an AI system on the behaviour and output of the system, as well as reflect on the concept of algorithmic bias and understand how human values and bias may define the behaviour of AI applications.
- (4) Discuss about the impact of the behaviour of an AI system on everyday life situations e.g. face recognition, autonomous vehicles.

The game was developed in the Unity game engine to facilitate deployment to a range of platforms. The game has been initially launched on Windows (32 and 64-bit) operating system and deployed on the Learn ML website (<http://learnml.eu/games.php>) using WebGL (see Figures 1 and 2 for screenshots of the game).

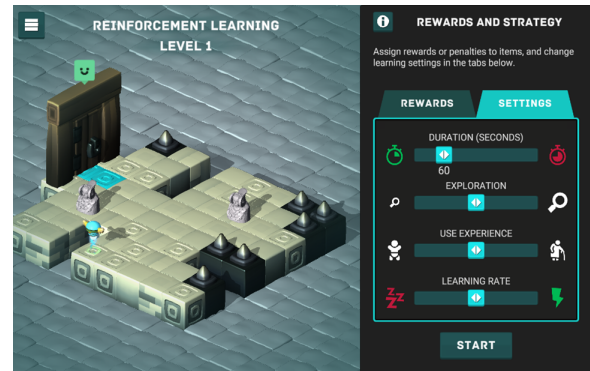
We aimed to introduce the students to elements and processes of supervised and reinforcement learning by combining a meaningful narrative and simulation mechanics. Players have the quest to find and retrieve valuable art objects that have been stolen and hidden. Through the first part of the game, the process of supervised learning is introduced; players train their AI helper to recognise specific art objects (paintings and sculptures). They classify a set of training data, experiment with different parameters, and then see how well the helper was trained by observing how it classifies a set of testing data. More specifically, a set of 40 photographs of paintings and sculptures available in the public domain were downloaded from the online open collection of The Metropolitan Museum of Art [9]. This superset of images was chosen to contain equal amounts of paintings and sculptures, and every time that this mini-game is played, a subset of 20 images is chosen at random as the training set. The players must label the images as either a painting or a sculpture (see Figure 1a). This design choice highlights both the essential

part of AI, which is data labelling, and shows the importance of classification at this stage, since the algorithm will then process the images according to the assigned labels. The disadvantage to this approach was that the size of the training set had to be kept as small as possible to avoid making the labelling task tedious. A small training set inevitably means an unreliable output of the learning algorithm, resulting in a bias towards the features prominent in the data set. However, this was determined to be less important than illustrating the concept of data labelling, and the results obtained were good enough for the scope of the game. The algorithm used to classify the images consisted of a decomposition of the palette of the images in the training set to a number of primary colours. The players are given a choice as to how many primary colours they want the algorithm to use, and the hue spectrum is subdivided accordingly. From this grouping a percentage of image pixels in each group is obtained and used to train a decision tree using a C4.5 algorithm [12]. The player may also determine the maximum depth of the decision tree, which was capped at 3 for the purpose of readability in the game (see Figure 1b). After the training is complete, the labels are inferred both for the training image set, as well as for the remaining images in the data superset. For the latter, the label is compared to the ground truth since no user labelling is available. An accuracy score is calculated for both the training and the testing set and shown to the player.

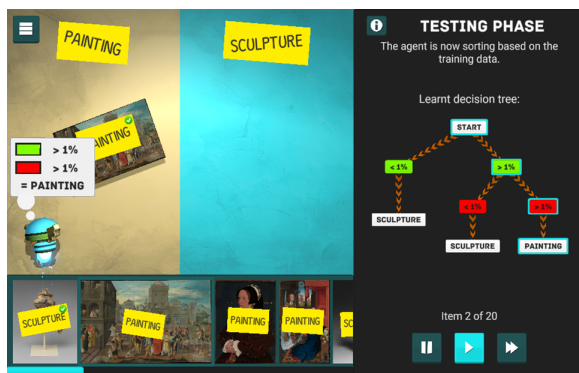
During the second part of the game, the players and their AI helper need to navigate through a series of dungeons and collect the stolen art objects. The players are introduced to the processes of reinforcement learning; they guide their helper by indicating what type of objects to look for and which ones to avoid (e.g. traps), by assigning rewards to the right objects. The AI helper tries to find its path based on the parameters set by the players, such as the exploration and exploitation rates. The players watch the process, they can pause or accelerate it (see Figure 1b), and think what settings would help the AI to find as many objects as possible. This part was implemented using a standard Q-learning algorithm [19]. The agent has to navigate a level that contains an exit, hazards (spikes), collectables (statues) and non-navigable spaces (holes), by performing a move in one of the four cardinal directions. The movements of the agent are discrete, as is the space of the level. For faster evaluation of the learning algorithm, the size of each level was limited to an 8×8 grid. This part allows the players to modify two sets of parameters: (1) The rewards or penalties assigned to hazards, rewards and the exit spaces, and (2) the parameters of the learning algorithm. The rewards or penalties are set by players using simple sliders ranging from a penalty (shown with a red icon) to a reward (shown with a green icon) for the hazards and the collectables. For the exit, the reward slider only allowed for a small or a large positive reward, not a negative one. This prevents the agent from roaming aimlessly the level if a negative reward is assigned to the exit, since in that case the game would only end with the destruction of the agent on a hazard space. These settings allow the players to observe how the type and magnitude of rewards (positive or negative) affect the behaviour learned by the agent. The second part of the settings control the main parameters of the algorithm (see Figure 2a), namely: (1) The total time for which the algorithm is allowed to run (learning duration), (2) bias of algorithm to exploration of unvisited locations (i.e. taking random actions),



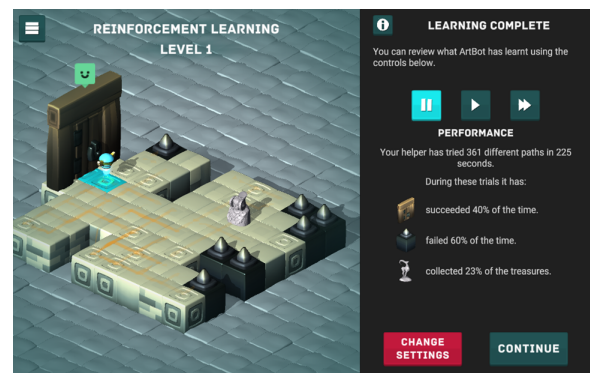
(a) Player setting up the supervised learning task



(a) Player setting up the reinforcement learning task



(b) AI agent simulating the supervised learning task



(b) AI agent simulating the reinforcement learning task

Figure 1: Screenshots of the first part of the “ArtBot” game: (a) The player labels the artefacts, (b) The agent classifies the training and the testing data sets.

Figure 2: Screenshots of the second part of the “ArtBot” game: (a) The player defines the variables of the agent’s reinforcement learning, (b) A final report is generated on the outcomes of the agent’s pathfinding.

versus using the rewards learned up to that point to determine the next best move (exploitation), (3) the discount factor, which controls the weight given to knowledge obtained in previous cycles in comparison to a new reward when it is being calculated (use experience), (4) the learning rate, which controls how much an obtained reward will influence the movement policy (learning rate).

The game guides the player through these sets of actions, but also provides opportunities for exploration, experimentation, and reflection. Students are introduced to the learning content through a more open, constructivist approach. They are encouraged to explore, experiment, and construct their knowledge by observing the outcomes of their actions, evaluate the results, make and test their hypotheses. Through the design of the game we tried to avoid common stereotypes and address students’ misconceptions of AI, such as the anthropomorphic nature of AI systems. We selected an unidentified object to represent the AI helper, so as to challenge the idea that AI only applies to robots, but that it may be applied to a variety of technology and computer systems. Players do have the option to choose and modify the appearance of their AI helper, including different models and colour variants; however, all models are fairly abstract and non-anthropomorphic. By setting the game in the context of cultural heritage (art objects) our aim was to address

the multidisciplinary application of AI systems, beyond computing and programming, such as archaeology, art, and transportation.

3.2 Game Assessment

Initial feedback on the game was collected through two online surveys for teachers and students. The students’ survey included 19 closed and open-ended questions, while the teachers’ survey included 20 closed and open-ended questions (2 of them on demographic characteristics). The survey was completed by 130 students from Greece, aged 12-17 years, and 17 teachers of different disciplines (physics, mathematics, computing, literature, language, and primary education teachers). 68 (52%) of the students identified as male, 58 (44%) identified as female, and 5 preferred not to say. 5 (29%) of the teachers identified as male while 12 (71%) identified as female.

The positive comments of the students mainly referred to the learning aspect of the game; the students enjoyed learning while playing and teaching the AI agent based on criteria they could choose, the settings they had to think about. They enjoyed that they had to think of the appropriate strategies and “use their brain”. They also reported enjoying the graphics, and the fact that they could play

with their friends. One of the students reported that “It’s different from the usual games”. Although most of the students reported that they wouldn’t change anything in the game, in some responses the students referred to the game as either too monotonous, boring, easy, and slow, or too complicated and time consuming at times.

The positive aspects of the game as reported by the teachers were the ease-of-use, the graphics, the interesting topic, the additional information included, and also the possibility of the players to select the appearance of their AI helper. Based on the teachers’ comments on the negative aspects, however, it seems that additional material would be necessary for explaining the different variables (e.g. the discount factor which was labelled as “use experience”) and how the settings and variables affect the behaviour of the AI helper. For some of the teachers it seemed difficult to understand what were the optimal settings for succeeding in the quests. The challenge of introducing advanced concepts of AI to novice students, as well as the critical role of supporting material and the guidance of a teacher were discussed by Parker et al. [10] in their review of the game “ViPER” which also aimed to teach AI and ML concepts to young students.

The closed-ended questions of the survey assessed the ease of use of the game, perceived enjoyment, whether respondents would like to use the game in the classroom, and their previous game experience. We report positive responses as those with agreement or strong agreement in the 5-item Likert scale. Indicatively, 88% of students and 77% of teachers responded that it was easy for them to understand how to play the game, 59% of students and 88% of teachers responded that they would recommend the game to friends and colleagues respectively, 73% of students and 82% of teachers responded that the game helped them understand what machine learning is. 82% of teachers responded that the game would help students understand the concept of AI.

4 CONCLUSIONS AND FUTURE WORK

The “ArtBot” game seemed to be well received by students and teachers; it seemed to be enjoyable and facilitated understanding of main ML concepts. Involving teachers and students from the early stages of design was valuable as it allowed us to consider their interests, previous experience with AI, potential technical limitations, and concerns about the role of AI. Further study is certainly needed for examining in depth the interaction of the students and teachers with the game in classroom settings, potential misconceptions that may emerge, learning outcomes, opportunities for critical thinking and reflection, and possible difficulties. One of the main points for future actions that emerged from the first assessment with students and teachers is the need to develop supporting material for elaborating on the main concepts presented, and also educational scenarios and activities for formal and non-formal education. Specifically, it would be important to provide topics for further discussion on the implications of the quality of the datasets and the training of the algorithms, resources for further reading, examples of real world implementations and implications.

Our goal in the LearnML project and through the “ArtBot” game was to develop supporting AI and ML education material and to trigger the critical thinking of students on the aspects, factors and bias that may shape AI agents and systems. We aim to provide

teachers, students, and the public with not only awareness and knowledge of AI and ML implications, but also practical tools and ways to envision and bring about positive change, using AI and ML as drivers for systemic change. Building upon the experiences and insights of existing applications, practices and tools for AI and ML education, it is important that curricula in primary and secondary education are re-designed so that the students can develop the literacies needed for critically engaging with AI applications and becoming more informed and conscious citizens [20]. Issues of inclusion and democratisation of access to AI literacy need to also be addressed, considering the differences in the attitudes towards AI, and experience with programming and AI technologies in children of different social, economical and cultural backgrounds [4].

ACKNOWLEDGMENTS

This work is supported by the “Learn to Machine Learn” (LearnML) project, under the Erasmus+ Strategic Partnership program (Project Number: 2019-1-MT01-KA201-051220).

REFERENCES

- [1] Tilde Bekker, Saskia Bakker, Iris Douma, Janneke van der Poel, and Koen Scheltenaar. 2015. Teaching children digital literacy through design-based learning with digital toolkits in schools. *International Journal of Child-Computer Interaction* 5 (2015), 29–38. <https://doi.org/10.1016/j.ijcci.2015.12.001>
- [2] Karl-Emil Kjær Bilstrup, Magnus H. Kaspersen, and Marianne Graves Petersen. 2020. Staging Reflections on Ethical Dilemmas in Machine Learning: A Card-Based Design Workshop for High School Students. In *Proceedings of the Designing Interactive Systems Conference*. 1211–1222. <https://doi.org/10.1145/3357236.3395558>
- [3] Douglas Clark, Brian Nelson, Pratim Sengupta, and Cynthia D’Angelo. 2009. Rethinking Science Learning Through Digital Games and Simulations: Genres, Examples, and Evidence. In *Learning science: Computer games, simulations, and education workshop*.
- [4] Stefania Druga, Sarah T. Vu, Eesh Likhith, and Tammy Qiu. 2019. Inclusive AI Literacy for Kids around the World. In *Proceedings of FabLearn*. 104–111. <https://doi.org/10.1145/3311890.3311904>
- [5] Michail Giannakos, Iro Voulgari, Sofia Papavaslopoulou, Zacharoula Papamitsiou, and Georgios Yannakakis. 2020. *Games for Artificial Intelligence and Machine Learning Education: Review and Perspectives*. 117–133. https://doi.org/10.1007/978-981-15-6747-6_7
- [6] Ting-Chia Hsu, Shao-Chen Chang, and Yu-Ting Hung. 2018. How to learn and how to teach computational thinking: Suggestions based on a review of the literature. *Computers & Education* 126 (2018), 296–310. <https://doi.org/10.1016/j.compedu.2018.07.004>
- [7] Tibor Koltay. 2011. The media and the literacies: media literacy, information literacy, digital literacy. *Media, Culture & Society* 33, 2 (2011), 211–221. <https://doi.org/10.1177/0163443710393382>
- [8] MIT Media Lab. 2021. *AI + Ethics Curriculum for Middle School Project*. MIT Media Lab, Personal Robots group. <https://www.media.mit.edu/projects/ai-ethics-for-middle-school/overview/> Retrieved April 13, 2021.
- [9] MET Museum of Art. 2020. *The MET Open Collection*. The Metropolitan Museum of Art, New York. <https://www.metmuseum.org> Retrieved November 12, 2020.
- [10] James Parker and Katrin Becker. 2014. ViPER : Game That Teaches Machine Learning Concepts - A Postmortem. In *Proceedings of the IEEE Games and Entertainment Media Conference*.
- [11] Stuart Russell, Daniel Dewey, and Max Tegmark. 2015. Research Priorities for Robust and Beneficial Artificial Intelligence. *AI Magazine* 36, 4 (2015), 105–114. <https://doi.org/10.1609/aimag.v36i4.2577>
- [12] Steven L. Salzberg. 1994. C4.5: Programs for Machine Learning by J. Ross Quinlan. *Machine Learning* 16, 3 (1994), 235–240. <https://doi.org/10.1007/BF00993309>
- [13] Nicole Sintov, Debarun Kar, Thanh Nguyen, Fei Fang, Kevin Hoffman, Arnaud Lyet, and Milind Tambe. 2016. From the Lab to the Classroom and beyond: Extending a Game-Based Research Platform for Teaching AI to Diverse Audiences. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 4107–4112.
- [14] Tapani Toivonen, Ilkka Jormanainen, Juho Kahila, Matti Tedre, Teemu Valtonen, and Henriikka Vartiainen. 2020. Co-Designing Machine Learning Apps in K–12 With Primary School Children. In *Proceedings of the IEEE International Conference on Advanced Learning Technologies (ICALT)*. 308–310. <https://doi.org/10.1109/ICALT49669.2020.00099>

- [15] Tommaso Turchi, Daniela Fogli, and Alessio Malizia. 2019. Fostering computational thinking through collaborative game-based learning. *Multimedia Tools and Applications* 78, 10 (2019), 13649–13673. <https://doi.org/10.1007/s11042-019-7229-9>
- [16] Henriikka Vartiainen, Matti Tedre, and Teemu Valtonen. 2020. Learning machine learning with very young children: Who is teaching whom? *International Journal of Child-Computer Interaction* 25 (2020). <https://doi.org/10.1016/j.ijcci.2020.100182>
- [17] Henriikka Vartiainen, Tapani Toivonen, Ilkka Jormanainen, Juho Kahila, Matti Tedre, and Teemu Valtonen. 2021. Machine learning for middle schoolers: Learning through data-driven design. *International Journal of Child-Computer Interaction* 29 (2021). <https://doi.org/10.1016/j.ijcci.2021.100281>
- [18] Steve Vosloo. 2018. *Guidelines for digital inclusion*. United Nations Educational, Scientific and Cultural Organization, Paris. <https://en.unesco.org/themes/literacy-all/pearson-initiative/guidelines> Retrieved 7 May 2021.
- [19] Christopher J. C. H. Watkins and Peter Dayan. 1992. Q-learning. *Machine Learning* 8, 3 (1992), 279–292. <https://doi.org/10.1007/BF00992698>
- [20] Mary E. Webb, Andrew Fluck, Johannes Magenheimer, Joyce Malyn-Smith, Juliet Waters, Michelle Deschênes, and Jason Zagami. 2020. Machine learning for human learners: opportunities, issues, tensions and threats. *Educational Technology Research and Development* (2020). <https://doi.org/10.1007/s11423-020-09858-2>
- [21] Randi Williams, Hae Won Park, and Cynthia Breazeal. 2019. A is for Artificial Intelligence: The Impact of Artificial Intelligence Activities on Young Children's Perceptions of Robots. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3290605.3300677>
- [22] Randi Williams, Hae Won Park, Lauren Oh, and Cynthia Breazeal. 2019. PopBots: Designing an Artificial Intelligence Curriculum for Early Childhood Education. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 9729–9736. <https://doi.org/10.1609/aaai.v33i01.33019729>
- [23] Abigail Zimmermann-Niefield, Makenna Turner, Bridget Murphy, Shaun K. Kane, and R. Benjamin Shapiro. 2019. Youth Learning Machine Learning through Building Models of Athletic Moves. In *Proceedings of the ACM International Conference on Interaction Design and Children*. 121–132. <https://doi.org/10.1145/3311927.3323139>