Introduction to Machine Learning with Robots and Playful Learning

Viktoriya Olari,[1] **Kostadin Cvejoski,**[2] **Øyvind Eide**[3]
[1,3] Department for Digital Humanities, University of Cologne, 50923 Cologne, Germany
[1,2] Fraunhofer IAIS, 53757 Sankt Augustin, Germany
[3] Center for Data and Simulation Science, University of Cologne, 50923 Cologne, Germany
viktoriya.olari@iais.fraunhofer.de, kostadin.cvejoski@iais.fraunhofer.de, oeide@uni-koeln.de

Abstract

Inspired by explanations of machine learning concepts in children's books, we developed an approach to introduce supervised, unsupervised, and reinforcement learning using a block-based programming language in combination with the benefits of educational robotics. Instead of using blocks as high-end APIs to access AI cloud services or to reproduce the machine learning algorithms, we use them as a means to put the student "in the algorithm's shoes." We adapt the training of neural networks, Q-learning, and *k*-means algorithms to a design and format suitable for children and equip the students with hands-on tools for playful experimentation. The children learn about direct supervision by modifying the weights in the neural networks and immediately observing the effects on the simulated robot. Following the ideas of constructionism, they experience how the algorithms and underlying machine learning concepts work in practice. We conducted and evaluated this approach with students in primary, middle, and high school. All the age groups perceived the topics to be very easy to moderately hard to grasp. Younger students experienced direct supervision as challenging, whereas they found Q-learning and *k*-means algorithms much more accessible. Most high-school students could cope with all the topics without particular difficulties.

Introduction

The guidelines and concrete proposals for AI curricula pay special attention to the technological aspects of AI, especially to machine learning (Sloman 2009; Blakeley and Breazeal 2019; Clarke 2019; Touretzky et al. 2019; Long and Magerko 2020; Wong et al. 2020). Children of all levels, from primary to high school, are expected to be able to cope with the central paradigms of machine learning: supervised, unsupervised, and reinforcement learning (Williams et al. 2019; Kahn et al. 2018; Jatzlau et al. 2019; Michaeli, Seegerer, and Romeike 2020). Lin et al. (2020) and Hitron et al. (2019) argue that understanding the concrete processes is particularly crucial in creating the proper mental models and avoiding misconceptions. However, few approaches focus on making the technical part of machine learning tangible for young learners (Williams, Park, and Breazeal 2019; Williams et al. 2019; Lin et al. 2020). Most current approaches either resemble a black box or are complicated and thus inaccessible to primary- or middle-school students (Jatzlau et al. 2019).

With our approach, we seek to fill this gap. Inspired by narrative techniques, designs of children's books, and the advantages of educational robotics and visual programming languages, we developed two extensions[1] for the open-source platform Open Roberta Lab[i]:

- The Neural Network Playground allows the user to experiment with simple neural networks. The student can train the network by modifying the weights and directly observing the effects on the simulated robot until it behaves as desired. In this way, the student grasps the concept of *direct supervision* – a process of adjusting the weights in the neural network until the output is satisfactory.
- With the Q-learning Playground, the student can tinker with the Q-learning algorithm by creating unique learning environments for the robot and playing with the parameters of the algorithm. Step by step, the student can debug the algorithm and explore how it is learning from the agent's perspective.

We also developed an unplugged activity to make unsupervised learning tangible by adapting the *k*-means algorithm. Our extensions are accompanied by a curriculum, which introduces young learners to the respective machine learning paradigms.

We tested our developments with 24 participants as representatives of three education levels: primary, middle, and high school. In the evaluation, we examined how children

[1] Codes, learning materials, and a video demonstrating the functionality of the extensions are available at github.com/vlebedynska/openroberta-lab/, retrieved: 17.12.2020.

of different ages perceived the topics and whether they had difficulties in understanding them.

In this paper, we first discuss the related work on the introduction of machine learning in schools as well as background studies on the use of robots, simulation, and playful learning in education. We then present design principles that guided us in developing the extensions and the curriculum. We continue with a presentation of the extensions and supplementary materials and describe our evaluation methods. In the end, we provide insights into the user study, summarize the children's feedback and discuss the results.

Background and Related Work

Machine Learning in Schools

Although there is a wide range of easy-to-use services introducing beginners to supervised machine learning, they usually use only a limited number of descriptive examples, such as image, text, sound classification (Teachable Machines, Machine Learning for Kids), or speech synthesis (Cognimates)[ii]. The main disadvantage of using such applications in education is that the mechanisms underlying training and classification remain hidden from the user (Hitron et al. 2019; Jatzlau et al. 2019). The children play with high-end systems and ready-trained models, with no opportunity to learn how the training is performed and how algorithms work behind the scenes. There have been increasing efforts to open the black box of supervised learning using visual programming languages. However, even then, these often either provide an interface to powerful high-end APIs (Druga 2018) or reproduce the underlying concepts without adapting them for the young learner. The proposal by Kahn et al. (2020) and Kahn et al. (2018) to teach deep learning, image classification, and speech synthesis with programming language Snap! is hardly suitable for children, due to its complicated technical terminology and implementation.

Introductory activities around reinforcement and unsupervised learning are rare. Michaeli, Seegerer, and Romeike (2020) conducted a case study introducing unsupervised machine learning with the learning vector quantization algorithm. However, the block-based approach used in the study is complex and targets high-school students. Some case studies have focused on teaching Q-learning as one of the reinforcement learning algorithms (Jatzlau et al. 2019; Toivonen, Jormanainen, and Tukiainen 2017). Again, these activities are aimed at undergraduate and high-school students and do not apply to younger children.

Blocks and Robots

Kahn and Winters (2017) and Jatzlau et al. (2019) argue that a block-based approach is child-friendly, intuitive, and straightforward. The use of blocks provides the user with easy access and low barriers to entry into programming applications (Druga 2018; Kahn et al. 2020). However, the use of blocks does not necessarily imply the simplification of the content itself. Although blocks may be used to introduce complex topics, the representation of the algorithms proposed by Kahn and Winters (2017) and Jatzlau et al. (2019) is not suitable for young children, due to its complicated vocabulary and numerous technical details.

There have been few attempts to implement robots in the classroom to teach young children about machine learning. Their success and effectiveness have been demonstrated in a small number of case studies with kindergarten and primary-school students (Druga et al. 2018; Williams, Park, and Breazeal 2019; Williams et al. 2019; Lin et al. 2020).

Simulation and Modeling

Robot simulators are as beneficial to the learning process as the operation of the real robot (Papert 1993). Simulators even have a decisive advantage – the time required for code-test-debug loops is considerably less than working with real robots (Dodds et al. 2006).

Simulation is an activity that actively engages with models, and it contains significant learning potential. In the current view of modeling as pragmatic processes with significant epistemological potential (Gelfert 2016; Ciula et al. 2018), active engagement with models, or "imaginary concreta" (Godfrey-Smith 2009: 108), has significant potential in research as well as in learning (Nersessian 2008). Indeed, simulation is a core concept with which computer games can be understood; it is creative confrontation with a narratological understanding (Aarseth 1998).

One finds experimental, playful, and practical problem solving in many computer games as well as in the long history of gaming and play (Salen et al. 2004; Flanagan 2009). It has significant untapped potential in a wide area of applications, from flight simulators to the recent growth in gamification. We believe that the bottom-up approach in our research, where a basic understanding of practical cybernetics is used to establish the basis for machine learning, provides a robust foundation for further development of teaching methodologies grounded in creative and playful modeling based on simulation systems.

Methodology and Curriculum Design

Considering the shortcomings of current approaches, we propose a curriculum for introducing machine learning with simulated robots and the visual programming language NEPO.

We build on the two "big ideas" of AI – perception and learning (Touretzky et al. 2019) – as practical guidance for designing AI curricula. Students are expected to create applications with simulated robots. This helps them to understand perception as a process in which sensors are used to extract data from the environment. At the same time, they immerse themselves in the challenges of supervised, unsupervised, and reinforcement learning by training neural networks and interacting with underlying algorithms.

We also took inspiration from the "Four P's of Creative Learning" (Resnick and Robinson 2017) as a modern framework that engages students in creative learning experiences. We strive to incorporate their ideas of playful learning (Papert 1993; Resnick and Robinson 2017; Resnick and Silverman 2005) into our extensions wherever possible. Although we designed some structured activities, such as Neural Network and Reinforcement Learning Cards, to help learners get started, our aim is that they serve as a stepping stone and not a final destination. We want to enable the participants to play with machine learning technologies and make something that interests them, following the ideas of constructionism (Queiroz et al. 2020; Michaeli, Seegerer, and Romeike 2020; Papert and Harel 1991).

We modified the design principal of embodied interaction (Long and Magerko 2020), by virtually putting the student in "the agent's shoes." Learners should immerse themselves in the behavior of the simulated robot. Doing so allows them to look behind the scenes and thus promotes transparency (Long and Magerko 2020) – another principle toward explaining AI.

We also based our approach on the principle of "low floors and wide walls" to accommodate children across various education and skill levels (Resnick and Silverman 2005). We have made it as easy as possible to get started with the extensions, while designing opportunities for students to dive deeper into the work on the topics. To this end, we took inspiration from the graphic design and storytelling of children's books. We followed the principles described by Castella (2018) in keeping our extensions and materials appealing and straightforward: Designs for children have to cater to everyone and leave room for exploration.

The entire curriculum consists of four thematic modules and lasts around 360 minutes, which the teacher can shorten or extend as needed.

Module 1, "How does your robot learn?," introduces children to the three paradigms of machine learning: supervised, unsupervised, and reinforcement learning (Russell and Norvig 2016). The children discuss two experiments proposed by Braitenberg, "Fear" and "Love" (Braitenberg 1986), which the facilitator performs with a robot in the front of the class. We chose these experiments because they are a simple entry point into the questions of what intelligence is and how it relates to learning. After the discussion, the facilitator holds a short input lecture that introduces machine learning.

In Module 2, "Teaching your robot," the children teach the robot various behaviors through direct supervision. They create simple neural networks by composing short programs in the Open Roberta Lab. When they start the program on the simulated robot, the program is compiled, and a neural network is created. They can now train the neural network by modifying the weights and observing the results directly from the behavior of the robot. The process of adjusting the weights until the robot behaves as desired is what we mean by *direct supervision* – the students are involved in the training process of the neural network and imitate it by manually adjusting the weights. As the children receive immediate feedback from the configuration of the network, they begin to understand how the robot learns. Immersing themselves in the training process allows them to focus on the underlying processes of supervised learning. At the same time, the children discover hands-on components of neural networks, such as nodes, layers, links, and weights. They can start with the Neural Network Cards, but they are then encouraged to explore and test the limits of what they can teach to the robot.

In Module 3, "Let your robot learn from experience," the children explore the Q-learning algorithm using the Open Roberta Lab and analyze how a robot learns through rewards. We chose Q-learning, because it is a simple model-free algorithm that has already been tested with children, and it is suitable for use in schools (Jatzlau et al. 2019). The children can create unique learning environments for the robot and experiment with the parameters of the algorithm. After they start the program on the robot, they observe and analyze the learning and reasoning process step by step with the Q-learning Playground. They may experience cases in which the robot fails to learn and finds no way out, and they can then correct the algorithm in the next iteration.

Finally, in Module 4, "Can robots learn autonomously?," the children are introduced to the *k*-means algorithm through unplugged activities. The facilitator leads a discussion on how the robot would group the objects on the table without any previous knowledge. He or she then sorts the items according to the *k*-means algorithm, without explaining what criteria was used for the grouping. The children are encouraged to make guesses. After discussing the grouping criteria and explaining the sorting principles, the students group the items themselves and let others guess their criteria. This activity aims to introduce them to cluster analysis and what it means to be in a group.

Extensions Design and Learning Materials

We designed our machine learning playgrounds based on the Open Roberta Lab, a visual block-based open-source programming platform. We chose this coding platform due to its focus on teaching programming with robots and its advanced ecosystem, including a robot simulation (Ketterl et al. 2015; Jost et al. 2014). We extended the platform with two playgrounds for the simulated LEGO EV3 robot.

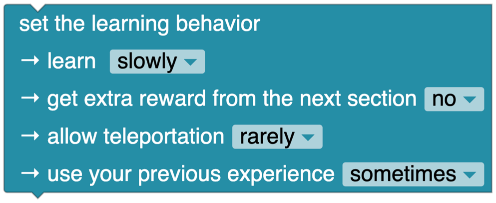

Figure 1. Block for setting up the learning behavior

First, we extended the block categories in the Open Roberta Lab with the new category of AI consisting of two subcategories: Neural Networks and Reinforcement Learning. For Neural Networks, we defined 10 new blocks:

- *Neural network* has two openings for the input and output layers. Technically, we included the case that the user can extend the block with hidden layers. However, the current block does not yet represent this graphically.
- *Neuron* is a block with a list in the backend where the user can plug in different types of input and output nodes. There is no limit to the number of neurons in a layer.
- *Input and output nodes*, enable the user to plug in the ultrasonic and the color sensors in light, RGB, and color modes as the neurons of the input layer. In the output layer, the user can plug in motor, LED, text, and sound.

For the subcategory of Reinforcement Learning, we implemented four new blocks, for which we adapted the appearance of the algorithm concerning technical vocabulary to make the algorithm tangible.

- *Map* allows users to set up the environment for the Q-learning algorithm.
- *Learning behavior* configures the parameters for the Q-learning algorithm. Figure 1 shows how users see the block. It is an example of how we have adapted the technical vocabulary of the Q-Learning algorithm for young students. The first parameter is Alpha, the second is Gamma, the third is Rho, and the fourth is Nu.
- *Gain experience* initializes the Q-learning Playground, where the children can specify the duration in seconds and the number of episodes.
- *Draw optimal path* draws the way for the robot to exit the labyrinth.

Second, we implemented two machine learning playgrounds. In the Neural Network Playground, the students experiment with direct supervision by adjusting the weights of the network. The Q-learning Playground promotes the transparency of the learning process and allows the users to start or stop the execution of the algorithm. They can also debug the algorithm step by step.

Figure 2 demonstrates the Neural Network Playground on the left and the simulation of the robot on the right. The input layer has two input neurons, which are ultrasonic sensors connected to ports 2 and 3. The output layer consists of two neurons (i.e., motors connected to ports b and c). If the user adjusts weights now, the robot immediately changes its behavior. In Figure 2, the weight between the first input and first output neuron is set to 1. It means that the value of ultrasonic sensor at port 2, which is currently 60, is completely transferred to motor at port b. Consequently, motor at port b rotates at a speed of 60. The same applies to the ultrasonic sensor at port 3 and the motor at port c. Such configuration of the neural network results in the following behavior of the robot: the closer the robot is to the object, the lower the value of the two ultrasonic sensors, and the slower the corresponding rotations of the motors.

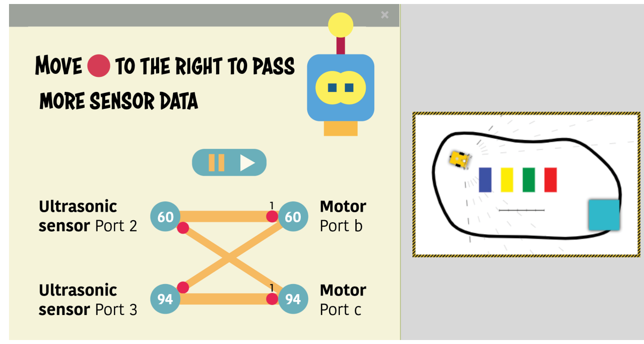

Figure 2. Neural Network Playground

We created eight Neural Network cards to introduce the user to direct supervision and neural networks. Figure 3 shows the learning card "Incognito."

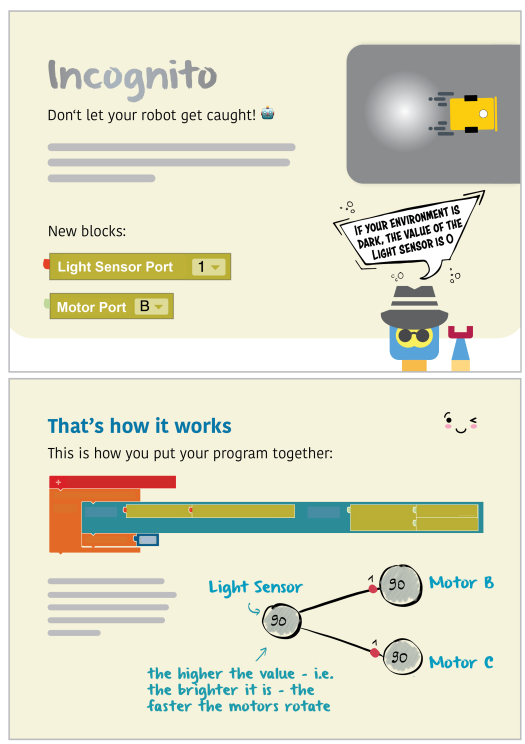

Figure 3. Neural Network Card "Incognito"

All the cards are similarly structured. On the front side, we describe the task and provide hints. Here, for instance, we show the new blocks that the learner needs and explain why the robot will attempt to drive onto the white area. On the back, we offer the solution for the program and our configuration of the neural network.

The vehicles proposed by Braitenberg (1986) were our inspiration for the cards. Hence, we adapted two scenarios: “Fear” teaches the robot to be “afraid” of obstacles, while “Friendship” instructs it to be friendly. We added six further examples: (1) Chameleon – Teach the robot to adapt itself to the environment, (2) Incognito – Teach the robot to avoid bright places, (3) Attention – Educate the robot on the traffic rules, (4) Loud Distance – Instruct the robot to measure the distance to the obstacle out loud, (5) Interest – Let the robot explore the landscape, and (6) Rally – Enable the robot to master the colored curves.

Figure 4 demonstrates the Q-Learning Playground. After the learner has created a program using the blocks and transferred it to the robot, the Q-Learning Playground is generated dynamically. The environment reflects the parameters that the user has set in the reinforcement learning blocks. On pressing the start button, the user starts the learning process, which can be observed in the navigation bar at the top and on the map. After the learning is finished, the optimal path out of the labyrinth is drawn, and the robot can now follow it. Figure 4 shows the last step, where the robot follows the optimal path.

We accompanied the Q-learning Playground with the materials, such as the Q-learning cards (on which the user can take notes), a Q&A, block descriptions, and a flow diagram on how Q-learning works. The materials developed aim to make the Q-learning algorithm tangible even for young learners.

For the unplugged activity that introduces the *k*-means algorithm, the facilitator only needs a set containing some objects and different colored Post-its. We assembled our collection from various drinking vessels and containers.

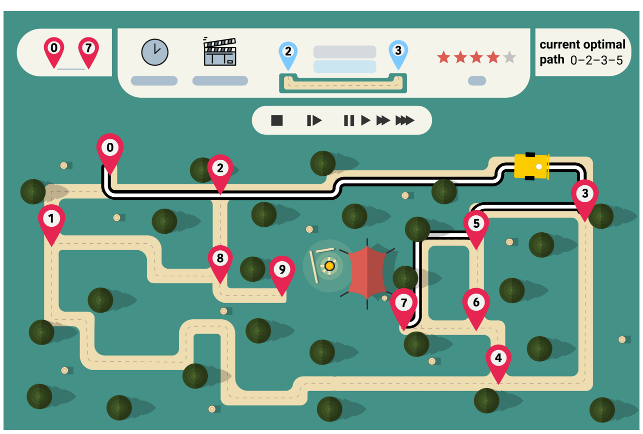

Figure 4. Q-Learning Playground

The facilitator sticks Post-its on a portion of the random items that serves as the cluster centers. The remaining objects in the set are then compared with each cluster center based on a criterion known only to the facilitator. After comparing each item with the cluster centers, the facilitator places the object behind the one cluster center that he or she thinks is the best fit. The version of the *k*-means algorithm that we performed with the children was simplified. It comprised only the first stage of the algorithm without post-clustering. However, it can be easily incorporated into the experiment.

User Study

We tested whether the learning experience with the extensions developed for machine learning promotes the children’s understanding of the underlying concepts of the subject. In particular, we evaluated how interesting and how difficult the students found the individual topics. We also asked the children what they thought AI and machine learning were, both at the beginning and the end of the session. In addition, we questioned their motivation to continue working on machine learning.

Method

Participants

Our aim in developing the extensions and teaching materials is to reach children of different ages and with no prior knowledge in machine learning. Therefore, we conducted a case study with students in various age groups. Twenty-four children participated in the study (Grades 3–4: four boys and five girls; 5–6: six boys and one girl; 7–9: seven boys only). We could not select the participants ourselves, since the sessions were organized as part of a summer-vacation program and had to be staffed on a “first come, first serve” basis. Some participants attended the session on their parents’ recommendation and were initially not enthusiastic about the workshops. The facilitator was informed that some of the students have special needs.

We held three sessions in total, one per day. Each session lasted six school hours (45 minutes) and was conducted in a block with short breaks. On the first day, we tested the extensions with the high-school students in Grades 7–9 (G1), on the second day with the primary-school students in Grades 3–4 grades (G2), and on the third day with the middle-school children in Grades 5–6 (G3). All the children had prior knowledge of working in the Open Roberta Lab with real LEGO EV3 robots, as they had participated in an introductory session on the previous day.

Procedure

In each session, we followed the modules as described in the “Methodology and Curriculum Design” section. We designed our presentation and instructions with a strong focus on the younger students and used them for all the age groups. First, we informally pre-assessed the knowledge of the children on machine learning and AI. We then completed the modules in order. At the end, the children filled out short questionnaires. On the second and third day, we changed the order of the modules, because the experiences from the first day indicated that reinforcement learning was too difficult to be tackled in the afternoon; the concentration of the students was lower in the afternoon than in the morning. We recorded the sessions, and an observer logged the activities for all three days.

Limitations

Since we could not influence the composition of the participant groups, we were not able to balance them by gender. We also could not collect detailed information about the background of each participant.[iii] Due to measures against coronavirus disease (COVID–19), we had two organizational restrictions: (1) Only a small number of children could participate in the sessions, and (2) the students were not allowed to work in groups. Therefore, we had to limit all activities to individual work.

In terms of content, we restricted the Q-learning environments to three maps and allowed the students to set as many obstacles as they wanted. This was necessary for obtaining comparable results at the end. However, our design does allow students to create and upload environments on their own under certain conditions.

We did not systematically examine whether our approach was effective in terms of measuring knowledge growth among students after each activity. Instead, we aimed to investigate how the students perceived the topics and whether they were able to cope with the complexity of the content. Future studies will focus on testing effectiveness, with more participants, a better gender balance, and diversity in terms of socioeconomic background.

Questionnaire

We were interested in the children's perception of the topics. Our goal was to understand how the children felt about the approaches and whether they had difficulty understanding them. On this basis, we developed a questionnaire with six items, based on a five-point semantic differential scale. We chose the semantic differential scale, because it enables quick measurement of attitudes and performs well with few items (Salkind 2006). Our items were: (1) How interesting did you find the topic "Supervised Learning and Neural Networks"? (2) How interesting did you find the topic "Unsupervised Learning"? (3) How interesting did you find the topic "Reinforcement Learning"? (4) Was the topic "Supervised Learning and Neural Networks" difficult to understand? (5) Was the topic "Unsupervised Learning" difficult to understand? (6) Was the topic "Reinforcement Learning" difficult to understand?

Interest	Score	Difficulty
very uninteresting	1	very difficult
uninteresting	2	difficult
neutral	3	neutral
interesting	4	easy
very interesting	5	very easy

Table 1. Distribution of scores for each response

To answer each question, the children could check a number on a scale from 1 to 5 between two pairs of adjectives: "very uninteresting" – "very interesting" for questions 1 to 3 and "very difficult" – "very easy" for questions 4 to 6. We then coded each response from 1 to 5, as shown in Table 1.

In order to obtain the overall attitude score $\overline{resp(I,G)}$ for each item I per group of participants G, we averaged the responses $resp(I,G)$ for each individual item:

$$\overline{resp(I,G)} = \frac{1}{|G|}\sum_{i\in G} resp(I,G)_i$$

We also asked the children about their general attitude toward further involvement with AI and machine learning. They could respond with "yes," "maybe," or "no." Furthermore, we invited them to provide written feedback (one sentence) about what they took with them from that day.

Results

Table 2 demonstrates the students' responses to the questionnaire items, and Figure 5 presents the responses graphically. The *x*-axis illustrates the absolute number of responses. The *y*-axis shows three topics divided by grade level. The first module is not considered, because it was only an introductory unit. Each bar of the diagram is aligned with the red dotted line that visually separates the responses with high scores (4–5) from the ones with lower scores (1–3).

Perception of Supervised Learning

The topic of supervised learning (Module 2) was the most difficult one for primary-school children, with an average score of 3.3. Children in middle school perceived it to be easy, with an average score of 4.0, as did the high-school students with 4.0. The middle-school students also found the topic of supervised learning to be the most interesting, compared to other groups. The average score here for fifth and fourth graders was 4.58, followed by third and fourth graders with 4.3 and sevens to ninth graders with 4.0.

The observations suggest that children of all grades were engaged and motivated by tinkering with neural networks and teaching the robot through direct supervision.

		Supervised Learning							Reinforcement Learning							Unsupervised Learning						
Score		1	2	3	4	5			1	2	3	4	5			1	2	3	4	5		
Interest	$\|G\|$						$\overline{resp(1,G)}$	SD						$\overline{resp(3,G)}$	SD						$\overline{resp(2,G)}$	SD
G1 Grades 7–9	7	1	0	1	1	4	4.0	1.52	1	0	0	2	4	4.14	1.67	1	0	2	1	3	3.71	1.14
G2 Grades 3–4	10	1	0	1	1	7	4.3	2.83	0	0	2	2	6	4.4	2.45	0	1	1	5	3	4.0	2.0
G3 Grades 5–6	7	0	1	0	0	6	4.58	2.61	0	0	1	2	4	4.42	1.67	0	0	2	2	3	4.14	1.34
Difficulty	$\|G\|$						$\overline{resp(4,G)}$	SD						$\overline{resp(6,G)}$	SD						$\overline{resp(5,G)}$	SD
G1 Grades 7–9	7	0	0	2	3	2	4.0	1.34	0	0	2	3	2	4.0	1.34	0	1	0	2	4	4.2	1.67
G2 Grades 3–4	10	1	2	2	3	2	3.3	0.71	0	1	1	3	5	4.2	2.0	1	0	0	2	7	4.4	2.92
G3 Grades 5–6	7	0	1	1	2	3	4.0	1.14	1	0	2	3	1	3.42	1.14	0	1	0	3	3	4.14	1.52

Table 2. Students' responses to questionnaire items 1–6

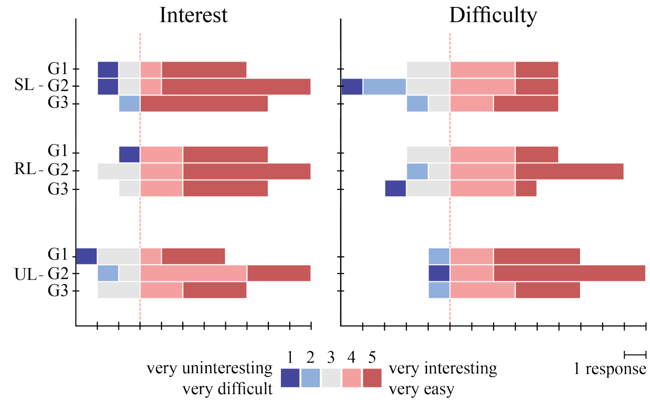

Figure 5. Attitudes of the participants in Grades 7–9 (G1), 3–4 (G2), and 5–6 (G3) toward the topics of supervised (SL), reinforcement (RL), and unsupervised learning (UL)

Most of the children completed only the task with the Neural Network cards. Few of them then had time to tinker with the applications based on their ideas. The feedback from the students in high and middle school was that the explanations were easy to follow. They also recommended improving some points of the user experience, such as the design of the playgrounds and the button locations.

Perception of Reinforcement Learning

Participants in all the age groups found the topic of reinforcement learning based on Q-learning (Module 3) neutral to very interesting. The average score of participants in Grades 3-4 was 4.4, and those of high-school students was 4.14. The middle-schoolers found the topic to be the most interesting, with an average score of 4.42. However, at the same time, they found reinforcement learning most challenging, with an average score of 3.42 for difficulty. High-school children perceived the topic to be more difficult than the primary-school students did (with 4.0 and 4.2 points, respectively).

The observer stated that each age group spent very different amounts of time creating the learning environments. Some children spent much time creating increasingly difficult environments, while others were more interested in testing. The older children had less motivation to carry out the experiments, and they were often more distracted than the middle- and primary-school students.

Perception of Unsupervised Learning

The middle-school students showed the greatest interest in the topic of unsupervised learning (Module 4) introduced by the unplugged activity, with an average score of 4.14. The lowest interest came from high-schoolers with 3.71, followed by primary-school children with 4.0. The average score for difficulty varied from easy to very easy in all three groups: 4.4 for primary-school, 4.14 for middle-school, and 4.28 for high-school students.

The observer noticed that the children were attentive when the facilitator conducted the experiment. They also actively participated in the discussion about the experiment afterward.

Student Motivation and Feedback

Of all the participants, 75% responded that they would continue to work on the topics, and 25% indicated that they might want to continue. One participant remarked: "I didn't like the topic with supervised learning so much, because I have a feeling that the tasks could also be solved with 'if-then' queries." Overall, however, we felt that the children had enriching sessions. One participant explained his experiences with reinforcement learning: "I found reinforcement learning very interesting, mainly because it improves by checking which way is the better one. […] AI is a bit more complicated than I thought, is really something that big ... can be tricky." Another participant reflected on his experiences with supervised and reinforcement learning and pointed out the moment when the robot could not find its way out despite its knowledge: "So, I take it from this day... I take all these ways with me […] I still can't describe […], but it's in any case, it's independence and that the [robot] can do something by himself without help, yes and also as an example he can say 'no,' which everybody is afraid of."

Discussion

By experimenting with the playgrounds and completing the four modules, children from primary to high school experienced the technical part of machine learning in practice. They taught the robot by training simple neural networks and explored how the robot can learn with rewards by experimenting with the Q-learning algorithm. They also familiarized themselves with unsupervised learning by exploring the *k*-means algorithm unplugged.

We demonstrated that our approach to introducing supervised, unsupervised, and reinforcement learning could raise the interest of students and be accessible even to young children. We include considerations for future approaches to teaching machine learning with robots and playfulness.

Introduce playfulness to machine learning extensions. We created engaging learning materials and neatly kept extensions, so that young students could immerse themselves in the machine learning topics. This approach, as described in "Methodology and Curriculum Design," was even well received by the young students. Although one of our objectives was to give the children more room for experimentation, this was only partly achieved. We observed that most students experimented with the underlying processes and algorithms based on the materials we provided. Only some students who were faster than others, and thus had time to work further on their projects, created more complicated learning environments or more complex network architectures. Therefore, tinkering with their projects should be expanded and emphasized in the future.

Let the robot learn and allow it to make mistakes. By teaching the simulated robot and experiencing the environment from the agent's perspective, the students gained insights into how the robot perceives the environment and how it learns. They also deepened their mental models of the capabilities and limitations of different machine learning approaches. The children were engaged in researching why the robot did not learn properly and why it could not find its way out from the labyrinth. They also wished to "train" the neural network so that the robot behaves correctly. We made a similar observation as Lin et al. (2020) that the robot's errors can be used to demonstrate that the agents are trainable and that they are not perfect.

Focus the design on the youngest students, and accommodate the older ones at the same time. Inspired by children's books on machine learning topics, we designed our materials and extensions with young students in mind. We adapted the topics in terms of technical vocabulary by reformulating the descriptions into the story-like narratives and revising the terminology. For visual communication, we used comics, hand drawings, and colorful illustrations. Both primary- and high-school students found the design of the extensions and materials appealing. This method of presenting complicated content can be used in the future.

Conclusion and Future Work

In this work, we presented our approach to the introduction of machine learning using robots, oriented towards playful learning and a child-centered design. The vast majority of children in all three age groups perceived the topics as exciting and easy to follow, and they expressed the intention to learn more about AI and machine learning. We promoted transparency in the underlying processes and algorithms by providing extensions of the Open Roberta Lab that reveal the machine learning algorithms. Although the extensions can be easily operated via the Open Roberta Lab interface, the critical complexity of the underlying processes is not lost. All three age groups could teach the robot different behaviors in the Neural Network Playground by exploring the basic principles of supervised learning. They also explored reinforcement learning step by step with the Q-learning Playground. For unsupervised learning, we adapted the *k*-means clustering algorithm, and the children explored it unplugged. Overall, modeling and practical simulation created more playfulness and fun in learning, without making the learning process less demanding and enriching.

We evaluated the children's perception of the proposed machine learning topics. The results indicate that the vast majority of the students found the topics engaging and easy to follow. We intend to test their understanding and the effectiveness of our extensions in the future.

We also aim to implement more open activities, so that children have more space to explore and experiment. In addition, we hope for more collaboration activities between the students, as this was not possible due to the limitations of COVID-19. We plan to integrate clustering into the robot simulation environment of the Open Roberta Lab and to migrate all the extensions from simulated to real robots.

Acknowledgments

We are grateful to the anonymous reviewers for their helpful comments. We would also like to thank the children who participated in this study as well as *Codingschule junior* for providing us with the opportunity to test the extensions and materials and for supporting us in organizing the sessions. Finally, we wish to thank the *Roberta Initiative* from Fraunhofer IAIS and Reinhard Budde, PhD, who supported us with their expertise on the Open Roberta Lab and other helpful advice.

This research is supported by the Competence Center for Machine Learning Rhine-Ruhr (ML2R), which is funded by the Federal Ministry of Education and Research of Germany (grant no. 01|S18038B).

References

Aarseth, E. J. 1998. *Cybertext: Perspectives on Ergodic Literature.* Johns Hopkins University Press. doi.org/10.2307/1513408

Blakeley, H. P., and C. Breazeal. 2019. *An Ethics of Artificial Intelligence: Curriculum for Middle School Students*. MIT Media Lab.

Braitenberg, V. 1986. *Vehicles: Experiments in Synthetic Psychology*. MIT Press. doi.org/10.1016/0004-3702(85)90057-8

Bredenfeld, A., and T. Leimbach. 2010. The Roberta® Initiative. In *2nd International Conference on SIMULATION, MODELING and PROGRAMMING for AUTONOMOUS ROBOTS (SIMPAR 2010) proceedings*, 558–67. Darmstadt (Germany).

Castella, K. 2018. *Designing for Kids: Creating for Playing, Learning, and Growing*. New York: Routledge. doi.org/10.4324/9781315266015-1

Ciula, A., Ø. Eide, C. Marras, and P. Sahle. 2018. Modelling, Thinking in Practice; An Introduction, In HISTORICAL SOCIAL RESEARCH, (31) 7-29. doi.org/10.12759/hsr.suppl.31.2018.7-29

Clarke, B. 2019. Artificial Intelligence - Alternate Curriculum Unit. In University of Oregon: Exploring Computer Science.

Deutscher Bundestag. 2020. Mehrheit der Fraktionen gegen den Begriff "Rasse" im Grundgesetz. In *19. Wahlperiode – 196. Sitzung: Tagesordnungspunkte 27 a bis 27 e*, 24766-80. Berlin.

Dodds, Z., L. G. Greenwald, A. Howard, S. Tejada, and J. B. Weinberg. 2006. Components, Curriculum, and Community: Robots and Robotics in Undergraduate AI Education, *AI Magazine*, 27: 11-22.

Druga, S. 2018. Growing Up with AI: Cognimates: from Coding to Teaching Machines. Master thesis, Program in Media Arts and Sciences, Massachusetts Institute of Technology, Cambridge, MA.

Druga, S., R. Williams, H. W. Park, and C. Breazeal. 2018. How smart are the smart toys? Children and parents' agent interaction and intelligence attribution. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*. doi.org/10.1145/3202185.3202741.

Flanagan, M. 2009. *Critical Play: Radical Game Design* (MIT Press). https://doi.org/10.7551/mitpress/7678.001.0001.

Gelfert, A. 2016. *How to Do Science with Models: A Philosophical Primer*. Springer International Publishing. doi.org/10.1007/978-3-319-27954-1.

Godfrey-Smith, P. 2009. Models and Fictions in Science, *Philosophical Studies*, 143. 10.1007/s11098-008-9313-2.

Hitron, T., Y. Orlev, I. Wald, A. Shamir, H. Erel, and O. Zuckerman. 2019. Can Children Understand Machine Learning Concepts? The Effect of Uncovering Black Boxes, In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. doi.org/10.1145/3290605.3300645.

Jatzlau, S., T. Michaeli, S. Seegerer, and R. Romeike. 2019. It's not Magic After All @ Machine Learning in Snap! using Reinforcement Learning, *2019 IEEE Blocks and Beyond Workshop (B&B)*: 37-41.

Jost, B., M. Ketterl, R. Budde, and T. Leimbach. 2014. Graphical Programming Environments for Educational Robots: Open Roberta - Yet Another One? In *2014 IEEE International Symposium on Multimedia*, 381-86. 10.1109/ISM.2014.24.

Kahn, K., Y. Lu, J. Zhang, N. Winters, and M. Gao. 2020. Deep Learning Programming by All, In *Constructionism 2020: Exploring, Testing and Extending our Understanding of Constructionism conference proceedings*, Dublin.

Kahn, K., R. Megasari, E. Piantari, and E. Junaeti. 2018. AI programming by children using Snap! block programming in a developing country. Vol 11082. Springer. doi.org/10.1007/978-3-319-98572-5.

Kahn, K., and N. Winters. 2017. Child-friendly programming interfaces to AI cloud services. In *Lavoué É., Drachsler H., Verbert K., Broisin J., Pérez-Sanagustín M. (eds) Data Driven Approaches in Digital Education. EC-TEL 2017. Lecture Notes in Computer Science*, vol 10474. Springer, Cham. doi.org/10.1007/978-3-319-66610-5_64.

Ketterl, M., B. Jost, T. Leimbach, and R. Budde. 2015. Open Roberta – a Web Based Approach Visually Program Real Educational Robots, *Læring & Medier (LOM)*, 8.

Lin, P., J. V. Brummelen, G. Lukin, R. Williams, and C. Breazeal. 2020. Zhorai: Designing a Conversational Agent for Children to Explore Machine Learning Concepts. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Association for the Advancement of Artificial Intelligence. doi.org/10.1609/aaai.v34i09.7061.

Long, D., and B. Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. doi.org/10.1145/3313831.3376727.

Michaeli, T., S. Seegerer, and R. Romeike. 2020. Looking Beyond Supervised Classification and Image Recognition @ Unsupervised Learning with Snap! In *Constructionism 2020: Exploring, Testing and Extending our Understanding of Constructionism conference proceedings*, Dublin.

Nersessian, N. J. 2008. *Creating Scientific Concepts* (MIT Press). doi.org/10.7551/mitpress/7967.001.0001.

Papert, S. 1993. *Mindstorms: Children, Computers, And Powerful Ideas*. Basic Books.

Papert, S., and I. Harel. 1991. *Constructionism: Research Reports and Essays, 1985-1990*. Ablex Publishing Corporation. doi.org/10.1037/031551.

Queiroz, R. L., F. b. F. Sampaio, C. Lima, and P. Lima. 2020. AI from concrete to abstract: demystifying artificial intelligence to the general public, *ArXiv* abs/2006.04013.

Resnick, M., and K. Robinson. 2017. *Lifelong Kindergarten: Cultivating Creativity Through Projects, Passion, Peers, and Play*. MIT Press. doi.org/10.7551/mitpress/11017.001.0001.

Resnick, M., and B. Silverman. 2005. Some reflections on designing construction kits for kids. In *Proceedings of the 2005 conference on Interaction design and children*: 117-22. doi.org/10.1145/1109540.1109556.

Russell, S., and P. Norvig. 2016. *Artificial Intelligence: A Modern Approach*. Pearson. doi.org/10.1016/j.artint.2011.01.005

Salen, K., K. S. Tekinbaş, E. Zimmerman, and Askews. 2004. *Rules of Play: Game Design Fundamentals* (MIT Press).

Salkind, D. N. J. J. 2006. *Encyclopedia of Measurement and Statistics*. SAGE Publications, Inc: Thousand Oaks. dx.doi.org/10.4135/9781412952644.n47.

Sloman, A. 2009. Teaching AI and Philosophy at School? *Newsletter on Philosophy and Computers*, 9(1), 42-48.

Toivonen, T., I. Jormanainen, and M. Tukiainen. 2017. An Open Robotics Environment Motivates Students to Learn the Key Concepts of Artificial Neural Networks and Reinforcement Learning. In *Robotics in Education*. Springer International Publishing. doi.org/10.1007/978-3-319-62875-2_29.

Touretzky, D., C. Gardner-McCune, F. Martin, and D. Seehorn. 2019. Envisioning AI for K-12: What Should Every Child Know about AI? In *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 9795-9799. Hilton Hawaiian Village, USA: AAAI Press. doi.org/10.1609/aaai.v33i01.33019795.

Williams, R., H. W. Park, and C. Breazeal. 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*. doi.org/10.1145/3290605.3300677.

Williams, R., H. W. Park, L. Oh, and C. Breazeal. 2019. PopBots: Designing an Artificial Intelligence Curriculum for Early Childhood Education, In *Proceedings of the 9th Symposium on Education Advances in Artificial Intelligence (EAAI '19)*. Menlo Park, CA, USA: AAAI Press. doi.org/10.1609/aaai.v33i01.33019729.

Wong, G. K. W., X. Ma, P. Dillenbourg, and J. Huan. 2020. Broadening artificial intelligence education in K-12, *Association for Computing Machinery (ACM) Inroads*, 11: 20-29.

[i] Open Roberta Lab: lab.open-roberta.org, retrieved: 16.12.2020.

[ii] Teachable Machines: teachablemachine.withgoogle.com, retrieved: 16.12.2020; machinelearningforkids.co.uk, retrieved: 16.12.2020; Cognimates: cognimates.me/home, retrieved: 16.12.2020.

[iii] The research took place in a German-language setting in Germany. Questions about “race” or affiliation to minorities are considered offensive and inappropriate in Germany (Deutscher Bundestag 2020).