The Road to AI Literacy Education: From Pedagogical Needs to Tangible Game Design

Marvin Zammit, Iro Voulgari, Antonios Liapis, Georgios Yannakakis
Institute of Digital Games, University of Malta, Malta
marvin.zammit@um.edu.mt
iro.voulgari@um.edu.mt
antonios.liapis@um.edu.mt
georgios.yannakakis@um.edu.mt

Abstract: The advancement of artificial intelligence (AI) and its increased use in everyday life have exacerbated the need to understand its underlying processes, and to raise awareness about its shortcomings and faults. Consequently there has been an increased effort to teach basic concepts of AI and machine learning (ML) from an early age. Digital games have been shown to be effective as teaching and learning tools, and there are ongoing efforts to increase AI literacy through educational digital games. To this effect, this work describes the process followed to extrapolate the design of such an educational game directly from the pedagogical needs emerging from stakeholders. Seven focus groups and workshops were conducted in Greece, Malta and Norway, with 55 teachers, researchers, practitioners and policy-makers, and 22 primary and secondary education students in total. These workshops identified seven goals which informed the design of a game for AI literacy called ArtBot. The game design process and the final interface and gameplay loop of ArtBot are explained in relation to these goals. The game was subsequently deployed across a variety of platforms, to enable dissemination to a broad audience in classrooms and elsewhere. The game was part of the tools developed in the framework of the LearnML Erasmus+ project. A preliminary online survey was used to gauge how well the game was received by teachers and students, with an overall positive result. A longer-term data collection of the game usage statistics has been initiated and will be analysed over the course of the game dissemination.

Keywords: ArtBot, Artificial Intelligence, Education, Game Design, Game-Based Learning, LearnML, Machine Learning

1. Introduction

In this paper, we describe the design and implementation of a digital game for Artificial Intelligence (AI) education and more specifically on Machine Learning (ML). AI generally refers to algorithms and processes through which computer systems may perform tasks that seem to mimic human behaviour or intelligence (Webb et al., 2020). ML is a subset of AI which involves the processes and models used for an AI system to learn, adapt and accomplish a task or goal more efficiently and effectively (Webb et al., 2020). ML is a key component of everyday applications that mediate our social, cultural and political interactions, such as recommendation systems, voice and image recognition applications, search engines, personalised information in newsfeeds and social media, and self-driving cars (Rahwan et al., 2019; Webb et al., 2020).

Latest advancements in AI and ML provide tremendous benefits and opportunities in sectors such as healthcare, education, and sustainable development (Division for Policies and Lifelong Learning Systems, UNESCO's Education, 2019; Vinuesa et al., 2020). Concerns have also been raised on the limitations, the role and challenges of AI specifically in areas involving ethical decisions, the propagation of information, and the autonomous behaviour of AI systems (e.g. vehicles or weapons) (Future of Life Institute, 2015; Russell et al., 2015). AI education is critical for societal understanding of its role and implications (Webb et al., 2020). Particularly, children and young people need to understand how AI and ML systems work, and develop critical thinking that enables them to identify benefits and challenges, to navigate and take an active role in shaping this digital environment. AI literacy skills would allow students to identify and interpret elements that define ML processes, such as the algorithms, systems, design decisions implemented, datasets used, and learning parameters (Rahwan et al., 2019). In this context, our goal was to design and develop an AI literacy game for primary and secondary education students, by addressing not only the technical aspects of AI and ML, but also encourage critical thinking on the ethical, societal and cultural implications of such systems.

Previous studies suggest that digital games are a valuable tool for AI education; through the simulation of complex systems, games can support systems thinking - the skills of “identifying and understanding systems, predicting their behaviors, and devising modifications to them” (Arnold and Wade, 2015; Waddington and Fennewald, 2018). Over the past few years, a number of games have been developed for supporting Computational Thinking and AI education specifically, addressing varied age groups and the technical, ethical, and social aspects and implications of AI (Clark et al., 2009; Giannakos et al., 2020). Turchi et al. (2019), for instance, introduced AI concepts to varied audiences of students and professionals through an online and a board role playing game addressing a real-world problem (the protection of wildlife). Clarke & Noriega (2003) developed a war simulation game for teaching AI, where students had to design the behaviour of AI agents. In “While True: Learn()” (Nival Interactive, 2018) the players assume the role of a ML specialist and complete tasks using visual programming. In “Human Resource Machine'' (Gabler and Gray, 2015) the players are introduced to concepts such as automation and optimisation by programming the employees in an office environment. In “Minecraft. Hour of Code: AI for Good'' (Microsoft Corporation, 2019) players are introduced to basic concepts of AI and its potential for protecting the environment, by programming a robot to predict forest fires. Parker et al. discuss the design, evaluation, and challenges of the game “ViPER'' which aimed to teach ML concepts to middle school students (Parker and Becker, 2014). Building upon these studies and applications, we describe the design, development, and initial reception of the game “ArtBot”, which is part of a game-based toolbox for AI education developed in the framework of the “Learn to Machine Learn” (LearnML) Erasmus+ project. The objective of this 3-year project is to improve digital literacy for children and teenagers by implementing a game-based learning toolbox that realises ML training scenarios in primary and secondary classroom settings. Our goal in this paper is to share our insights and design principles for the development of games for AI education and specifically understanding of ML.

2. The ArtBot Game: Design and Implementation

2.1 Pedagogical Design

Before designing the game, 7 focus groups and workshops were conducted, in Greece, Malta and Norway, with 55 teachers, researchers, practitioners, policymakers and 22 primary and secondary education students in total. The workshops examined participants’ perceptions about AI, introduced them to main AI concepts, and recorded their suggestions on the implementation of AI education to formal education (Giannakos and Papavlasopoulou, 2020). Some of the main points that emerged from these discussions include:

- fear of the changes to everyday life and habits that AI systems might bring;
- AI ethics;
- misconception that AI is mainly linked to robots;
- importance of students’ critical thinking skills and critical presentation of AI;
- lack of discussions at school about the social impact of AI;
- the need to familiarise students with how AI works, to help them attain a functional understanding of ML and the benefits and challenges of AI.

Certain challenges have also emerged through these discussions, such as:

- the limited time available at school settings;
- the need to adopt more student-centred approaches where students would be able to experiment and construct their knowledge;
- the lack of programming knowledge by teachers and, in many cases, students;
- limited resources available (such as computer access);
- danger of overwhelming teachers and students with too much information;
- the benefits and challenges of a multidisciplinary approach through various school subjects.

During the design of the game, we attempted to address these points and integrate them as a set of design principles and learning objectives presented below.

The game should be close to the students’ experience: Improving student engagement through meaningful activities, close to their experience and interests, has been raised through our workshops and also in existing studies on AI education (Toivonen et al., 2020; Zimmermann-Niefield et al., 2019). In this way, students explore the role and implications of AI inductively, and make inferences about the potential value and limitations of AI. For engaging the students we used a familiar game format and mechanics and tried to combine them with a meaningful, adventure narrative.

Allow for experimentation and knowledge construction: Our aim was to scaffold the learning path and the introduction of complex ML concepts to students, but also allow them to construct their knowledge through exploration following exploratory learning principles and constructivist approaches to learning (De Freitas, 2006; Egenfeldt-Nielsen, 2007; Marone, 2016). The game guides the player through a set of actions, and by integrating simulation mechanics, it provides opportunities for exploration, experimentation, and reflection. Students are introduced to the learning content through a more open, constructivist approach and are encouraged to explore, experiment, and construct their knowledge by manipulating variables, observing the outcomes of their actions, evaluating results, making and testing hypotheses. Thereby players understand the impact of the decisions made during the development of an AI system on its behaviour and output.

AI and ML concepts: The learning objectives regarding the technical aspects of ML were for the students to understand the process of (a) supervised learning and concepts such as training datasets, testing datasets, classification, and decision trees and (b) reinforcement learning concepts including rewards and penalties, learning duration, learning rate, exploration, exploitation, and pathfinding.

Ethical issues and concerns: Although ethical issues are not directly addressed in the game, the players are encouraged to discuss the role and impact of an AI system in situations such as face recognition, autonomous vehicles through open questions and prompts. Additional accompanying material will be designed to address these issues and concerns more explicitly, using the game as a background and motivation for further activities and discussions.

It should allow for an interdisciplinary teaching approach and be applicable to different fields: We decided to set the game outside of a technological setting, such as a computer lab, but instead in the context of cultural heritage (art objects). This addresses the multidisciplinary application of AI systems, beyond computing and programming, in such fields as archaeology, art, and transportation.

Address misconceptions about the systems that may integrate AI: We tried to avoid common stereotypes and address students' misconceptions of AI, such as the anthropomorphic nature of AI systems. The AI helper is an unidentified artefact, inspired by ancient artefacts. While players can customize the appearance of the AI helper avatar, none of the available options are anthropomorphic.

The game should be easy to install and use even in low-end systems: We addressed this point by deploying the finished game to as many digital platforms as possible. We targeted the Windows operating system for the personal computer platform primarily, but also made the game available for Android mobile devices. In addition, the game was deployed as a web browser game so that it is accessible also by other devices and operating systems.

2.2 Implementation

ArtBot is an educational game intended to improve AI literacy among primary and secondary education students through a series of tasks requiring the application of supervised and reinforcement learning. The premise of the game is that valuable statues have been stolen and an AI helper, ArtBot, has been provided to the player to help find and retrieve them. The game comprises 2 mini-games, each intended to illustrate one of these ML processes. In the first part, the player must teach ArtBot to recognise the difference between paintings and statues. In the second mini-game, the player has to guide it through a set of dungeon-like rooms by collecting statues and finding the exit.

ArtBot was developed for multiple platforms, initially for Windows operating systems, supporting both 32-bit and 64-bit systems, since some schools may still operate with older hardware. Subsequently it was deployed as a browser game using WebGL technology[1], and on Android devices via the Google Play store.

[1] accessible at http://learnml.eu/artbot.php

The story is presented through a combination of text and animated graphics (Figure 1) where players may choose their preferred avatar and its colours from a set of predetermined options, deliberately chosen not to be anthropomorphic to veer away from the notion that AI is only related to robots.

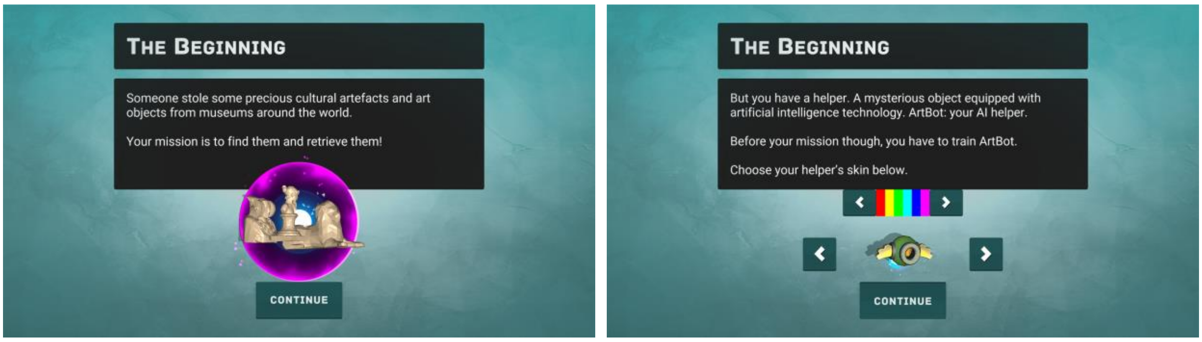

Figure 1: The story of the game is introduced via a set of short animations and accompanying text. Players are given a choice from a set of 5 avatars and a selection of colour schemes.

The minimalistic interface of the game has few on-screen instructions to avoid visual clutter and perceived complexity. However, throughout the game, players are provided with information buttons giving an optional verbose description of the interface, as well as more information about underlying AI processes.

2.2.1 Part 1: The supervised learning mini-game

The first part of the game introduces basic concepts of supervised learning. Players train their AI helper to categorise art objects as paintings or sculptures by classifying a set of training data. They experiment with different parameters and see how well the helper was trained by observing how it classifies a set of testing data.

For this part, a set of 40 photographs available in the public domain were obtained from the open collection of The Metropolitan Museum of Art. At every play-through, a subset of 20 images is chosen at random as the training set. The players label each image by swiping left or right, mimicking the essential and time-consuming task of data labelling in real-world AI. Players are intentionally allowed to label images incorrectly, and the algorithm will then process the images according to their choices, thereby highlighting the importance of initial data processing.

The size of the training set was kept small to avoid making the labelling task boring. This inevitably yields a less generalisable output of the learning algorithm, resulting in overfitting to features specific to that dataset. However, it was deemed more important to illustrate the concept of data labelling, and the classification results obtained in the game were reliable enough for its purpose.

Players can control the number of colour bands that each image is processed into, based on a closest colour quantization similar to Liapis (2018). The player may also determine the maximum depth of the decision tree. For large training sets it is customary to limit this depth to prevent the algorithm from overfitting to the data. In this game, the maximum depth was capped at 3, to show how the learning algorithm is affected in the case of a large decision tree, and also makes the resulting visualisation less daunting to a younger audience. These settings, shown in Figure 2, allow the players to explore how parameters impact the behaviour and accuracy of the AI system.

All pixels in each image are then analysed and grouped into the closest colour band. Players can see the original images recoloured with pixels in each colour group, to observe how the AI helper 'sees' the images. From this quantization, the ratio of image pixels in each group is obtained for supervised learning by a decision tree using a C4.5 algorithm implemented from the Accord.NET framework (Souza, 2014).

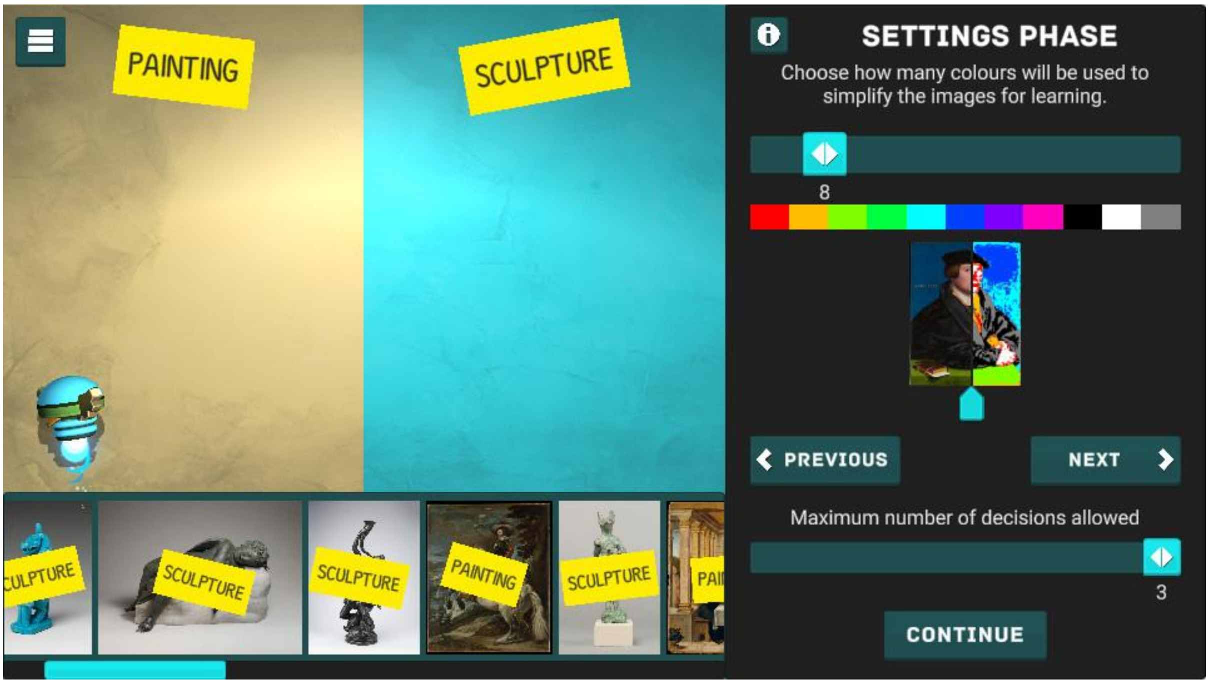

Figure 2: Settings used by the classification algorithm. The players are able to select a number of primary colours to subdivide the image and they can also see the effects of the subdivision on the images themselves. The maximum number of decisions that the tree is allowed (i.e. its depth) is also variable.

After training is complete, the players are introduced to the testing phase. The labels are inferred both for the training images, as well as for those remaining in the data superset (the testing set). For the latter, the label is compared to the ground truth since no user labelling is available. Students can view the trained algorithm "at work", with a graphical representation of the decision tree. Each node represents the basis of each decision taken. For example, in Figure 3 a yellow node labelled as "<2%" means that this branch will be followed if less than 2% of the pixels in the image are in the yellow group. Accuracy scores for the training and testing sets are finally shown to the player. The accuracy score refers to how close the labelling conducted by the AI helper was to the initial labelling done by the players (or to the ground truth, in the case of the testing set). The labelled images and whether they are accurate or not are shown in a carousel at the bottom of the screen, separated into two tabs; one for each data set (Figure 3).

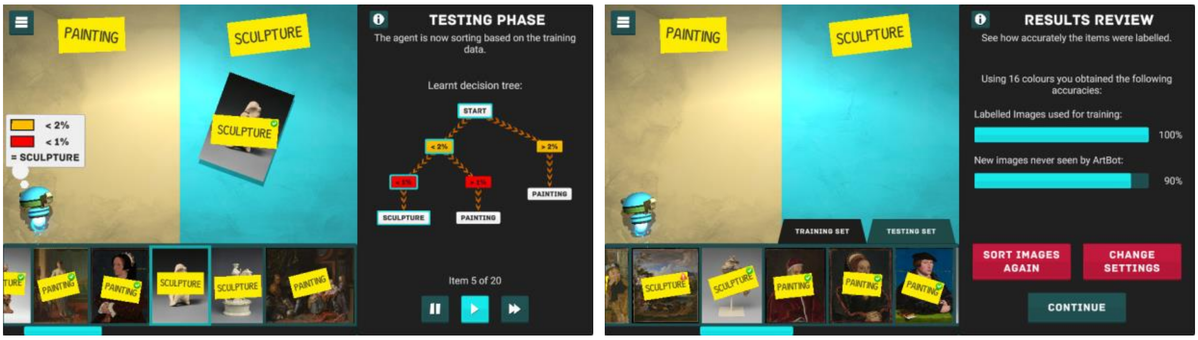

Figure 3: (left) The criteria of the classification are visualised in a decision tree, and in a speech bubble above the AI helper. (right) Players can view the classified training and testing sets, as well as accuracy scores for each set.

2.2.2 Part 2: The reinforcement learning mini-game

In the second part of the game, the players help ArtBot to navigate through 10 dungeons and collect the stolen art objects. 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 positive and negative rewards to each category. The helper searches for a path based on the parameters set by the players, such as the *exploration* versus *exploitation* rate. The players watch the process, they can pause or accelerate it, and reflect on which settings would be optimal for helping the AI agent find as many objects as possible. The agent must navigate a level that contains an exit, hazards (spikes), collectables (statues) and non-navigable spaces (holes), by performing moves in the four cardinal directions. The movements of the agent are discrete, as is the space of the level. For faster evaluation of the algorithm, the size of each level was limited to an 8 x 8 grid. The rules of the game are intuitive; if the helper moves into a hazard tile it is destroyed and the episode ends. If the helper lands on a tile occupied by a statue, the statue is collected (removed from the level). The episode also terminates when the AI helper reaches the exit tile.

The agent learns by repeatedly taking moves/actions and observing the rewards obtained. The reward for each state and action combination is recorded in a table (the Q-table) and the reward is evaluated by using a standard Q-learning algorithm (Watkins and Dayan, 1992). The state is recorded as a simple ASCII matrix, and the actions are recorded as up, down, left or right. This naive representation of state and actions means that the agent will search for a policy that will give the best move for a specific level, but it will not be valid for a different level layout since the state definition would be completely different.

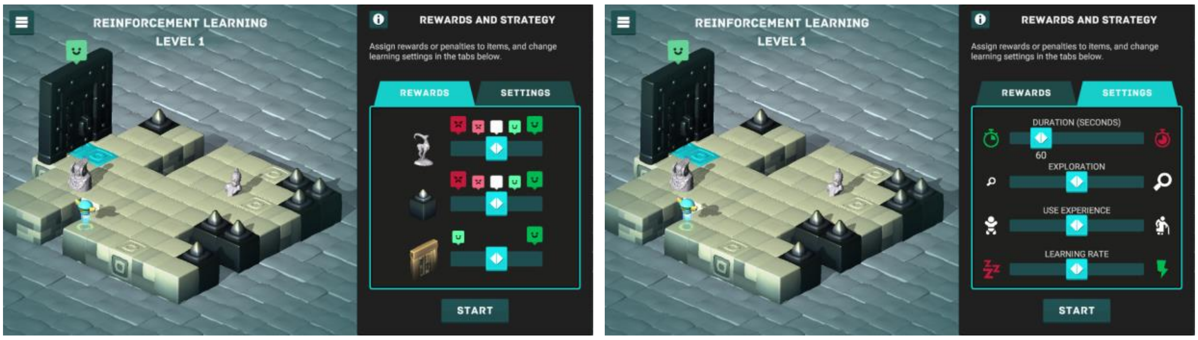

Figure 5: View of a user setting up the parameters for a level in the reinforcement learning mini-game, showing the controls that allow the players to set positive and negative rewards (left) and those for the parameters of the Q-learning algorithm (right).

In this mini-game the players are able to modify the rewards or penalties assigned to the different level tiles, and the parameters of the learning algorithm (Figure 5). The rewards are set using simple sliders ranging from a penalty (showing a red icon) to a reward (showing a green icon) for the hazards and the collectables. If a negative reward were to be assigned to the exit, the game would only end with the destruction of the agent on a hazard space. Therefore, the reward slider for the exit only allows a positive reward, preventing the agent from roaming aimlessly. These settings allow the players to observe how the type and magnitude of each reward (positive or negative) affect the learned behaviour.

The second part of the settings control the main parameters of the algorithm, namely:

1. The total time for which the algorithm is allowed to run ("Learning duration" slider).
2. The bias of the algorithm towards exploration of unvisited locations, i.e. taking random actions, versus using the rewards learned up to that point to determine the next best move ("Exploration" slider).
3. The discount factor, which controls the weight given to knowledge obtained in previous cycles in comparison to a newly obtained reward ("Use experience" slider).
4. The learning rate, which controls how much an obtained reward will influence the movement policy ("Learning rate" slider).

Once the settings are confirmed and the learning algorithm begins, the agent starts taking random moves, updating the Q-table with the rewards obtained according to the settings. As the algorithm progresses and the agent records more information, the actions selected will be less random, replaced by the best moves for the given state. The rate at which this randomness decreases is governed by the exploration slider.

The player is shown the AI helper avatar taking these actions in the level. Since the algorithm requires several hundreds or thousands of episodes for a reliable policy to be found, the game runs the algorithm in the background at a much faster rate than it would be possible to visualise. Consequently, players only observe one episode for every hundred that the algorithm runs. Whilst these episodes are being played out, a trail marks the paths taken by the avatar, adding to its intensity over the tiles which are visited more frequently (Figure 6).

When training is complete, players are given the option to observe (at normal or fast speed) the optimal path(s) found by the algorithm. Statistics about the learning process are also displayed: the duration of the training, the total number of episodes performed, the ratio of simulations that an exit was found, the ratio of simulations that hazards were encountered, and the percentage of statues that were collected (Figure 6).

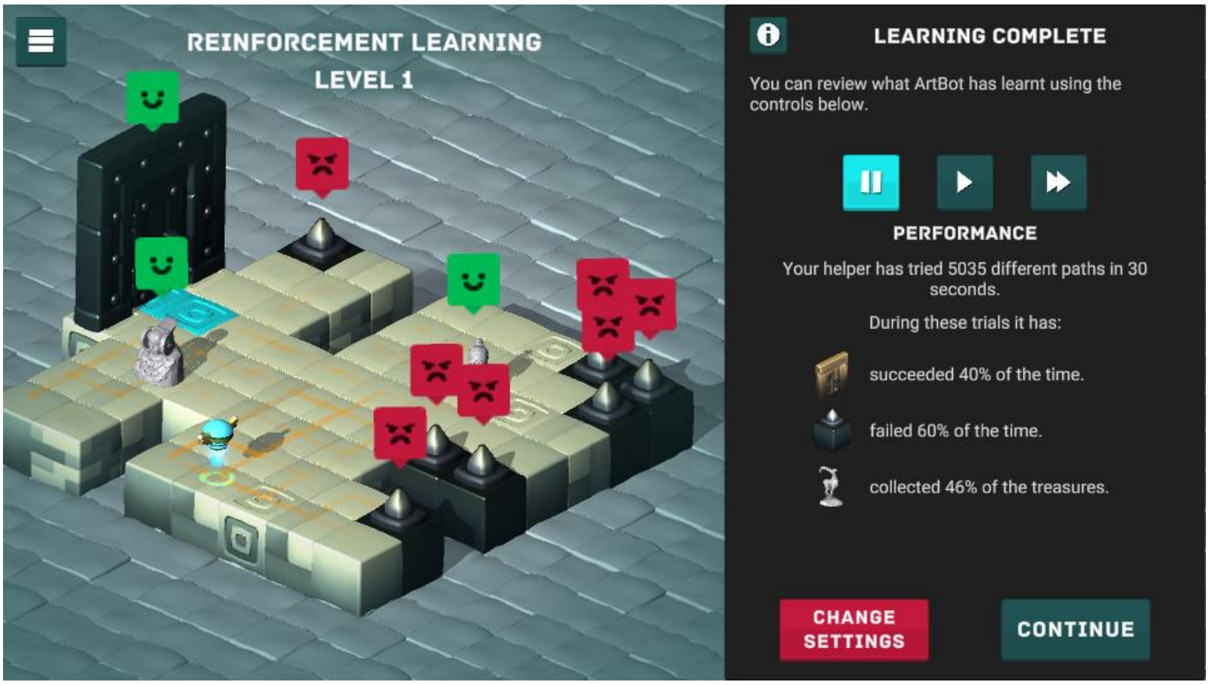

Figure 6: The summary screen of the reinforcement learning mini-game, showing the statistics collected during the learning process.

The player can progress to the next level if the optimal path found includes the exit and at least one statue. The player may also retry the level with different settings, giving them the opportunity to witness how different parameters affect the optimal path found.

A total of ten levels of this mini-game are available, each varying the challenge to the learning process by adjusting the proximity of statues to hazards and exits, the directness of the path towards the exit, etc.

3. Reflections

One of the initial challenges for our game design was the identification of learning objectives. AI and ML entail abstract and complex concepts which might be difficult for young students to understand. We tried to ensure that the concepts explored in this game were scientifically sound and grounded in real-world machine learning processes, and also simplify them by identifying the main components and removing elements that might distract from the learning goals.

Early involvement of the community (educators, researchers, students, policy makers, AI experts) informed not only the design of ArtBot but even, a step before that, the decision to develop such a game. It was important for us to address interests, needs, and varied backgrounds of students as well as of teachers. In line with participatory design principles and the empowerment of children and educators (Druin, 1999; Könings et al., 2014), involving the community and providing the space to actively participate in the design decisions was valuable. It allowed us to design a game that could address the real educational requirements of students and teachers, and also consider situational factors of formal education settings and the classroom, such as the available infrastructure and the promotion of a culture of interdisciplinary collaboration among teachers of different subjects.

Although the game includes additional information about the implemented algorithms in the pop-up information panels, the scope of the game was to give an overview of ML processes and how these are affected by parameters controlled by a human operator. Our objective was not to teach how the algorithms operate step-by-step, since we observed in the initial workshops that not all educators were familiar with them. Given that LearnML aims to increase AI literacy across education levels and curricula, ArtBot is intended more as a playground for exploring concepts, parameters and biases of ML algorithms.

Initial reception of the game, as measured with online surveys with open-ended and closed-ended questions, was generally positive (Voulgari et al., 2021). Some of the positive aspects reported were the game-based learning approach, the thinking and strategy required, the graphics, and the information included. Negative aspects reported such as boring, monotonous, or time-consuming gameplay and confusing instruction will be considered in future iterations of the game. Moreover, more teacher training will be carried out, and educational scenarios for classroom use will be prepared to further align gameplay with learning goals via e.g. debriefing (Plumettaz-Sieber et al., 2019). Further assessment of the game is certainly needed, with more students and teachers across more countries, in classroom settings, combined with qualitative approaches (e.g. observation, interviews) and game analytics. Anonymous usage data collection has been included in the game, which will be analysed once the game has been deployed to a broad set of classrooms.

4. Conclusion

The objective of this study was to dissect the design of a game aimed to improve AI literacy. The design process was initiated by soliciting direct feedback from the primary stakeholders, using the needs of educators and students to inform the design choices. The resulting ArtBot game teaches supervised learning and reinforcement learning under a consistent visual style and narrative, and allows for a constructivist playground for students to explore the impact of parameters, rewards, and training data on the learned behaviour of AI agents. The project will continue to explore the impact of ArtBot in diverse classroom settings, ages, and school subjects. Moreover, additional teaching material will be developed to accompany the game in classrooms and provide additional information about the ML algorithms in use. This will facilitate discussions of topics such as AI bias and AI ethics, which are necessary to complete the educational experience, and to maximize the impact of gameplay on reflection and positive change.

5. Acknowledgements

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).

6. References

Arnold, R.D. and Wade, J.P. (2015) "A Definition of Systems Thinking: A Systems Approach", *Procedia Computer Science, 2015 Conference on Systems Engineering Research* Vol 44, pp 669–678.

Clark, D., Nelson, B., Sengupta, P. and D'Angelo, C. (2009) "Rethinking Science Learning Through Digital Games and Simulations: Genres, Examples, and Evidence", *Learning Science: Computer Games, Simulations, and Education Workshop Sponsored by the National Academy of Sciences, Washington, DC*.

Clarke, D. and Noriega, L. (2003) "Games Design For the Teaching of Artificial Intelligence", *Interactive Convergence: Research in Multimedia*. Prague, Czech Republic.

De Freitas, S. (2006) "Using games and simulations for supporting learning", *Learning, Media & Technology,* Vol 31, pp 343–358.

Division for Policies and Lifelong Learning Systems, UNESCO's Education (2019), "Artificial Intelligence in Education: Challenges and Opportunities for Sustainable Development (Working Paper No. 7)", *Working Papers on Education Policy*. UNESCO, Paris, France.

Druin, A. (1999) "Cooperative Inquiry: Developing New Technologies for Children with Children", *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM. pp 592–599.

Egenfeldt-Nielsen, S. (2007), “Third Generation Educational Use of Computer Games”, *Journal of Educational Multimedia and Hypermedia,* Vol 16, pp 263–281.

Future of Life Institute (2015) “Autonomous Weapons: An Open Letter from AI & Robotics Researchers”.

Gabler K. and Gray K. (2015) “Human Resource Machine”, [online], Tomorrow Corporation, http://tomorrowcorporation.com/humanresourcemachine

Giannakos, M. and Papavlasopoulou, S. (2020) “IO1: A LearnML Pedagogical Framework Development (Intellectual Output 1)”, *Technical Report, Learn to Machine Learn (LearnML)*, Erasmus+ project.

Giannakos, M., Voulgari, I., Papavlasopoulou, S., Papamitsiou, Z. and Yannakakis, G. (2020) “Games for Artificial Intelligence and Machine Learning Education: Review and Perspectives”, in: Giannakos, M. (Ed.), *Non-Formal and Informal Science Learning in the ICT Era, Lecture Notes in Educational Technology*. Springer. pp 117–133.

Könings, K.D., Seidel, T. and van Merriënboer, J.J.G. (2014) “Participatory design of learning environments: integrating perspectives of students, teachers, and designers”, *Instructional Science,* Vol 42, pp 1–9.

Liapis, A. (2018) “Recomposing the Pokémon Color Palette”, *Proceedings of the Applications of Evolutionary Computation Conference*, Springer, pp 308-324.

Microsoft Corporation (2019) “Minecraft. Hour of Code: AI for Good”, [online], https://education.minecraft.net/hour-of-code

Marone, V. (2016) “Playful Constructivism: Making Sense of Digital Games for Learning and Creativity Through Play, Design, and Participation”, *Journal For Virtual Worlds Research,* Vol 9.

Nival Interactive (2018), “while True: learn()”, [online], https://luden.io/wtl/

Parker, J.R. and Becker, Katrin (2014) “ViPER : Game That Teaches Machine Learning Concepts - A Postmortem”, *Proceedings of the IEEE Games and Entertainment Media Conference,* IEEE.

Plumettaz-Sieber M., Bonnat C. and Sanchez E. (2019) “Debriefing and Knowledge Processing an Empirical Study About Game-Based Learning for Computer Education”, In: Liapis A., Yannakakis G., Gentile M. and Ninaus M. (eds), *Proceedings of the Games and Learning Alliance conference. Lecture Notes in Computer Science*. Springer.

Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J.W., Christakis, N.A., Couzin, I.D., Jackson, M.O., Jennings, N.R., Kamar, E., Kloumann, I.M., Larochelle, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D.C., Pentland, A. ‘Sandy,’ Roberts, M.E., Shariff, A., Tenenbaum, J.B. and Wellman, M. (2019) “Machine behaviour”, *Nature,* Vol 568, pp 477–486.

Russell, S., Dewey, D. and Tegmark, M. (2015) “Research Priorities for Robust and Beneficial Artificial Intelligence”, *AI Magazine,* Vol 36, No. 4: Winter 2015.

Souza, C.R. (2014) "The Accord.NET Framework", [online], São Carlos, Brazil, http://accord-framework.net.

The Metropolitan Museum of Art, “The MET Open Collection”, [online], New York, https://www.metmuseum.org/art/collection/search, accessed 12th November 2020.

Toivonen, T., Jormanainen, I., Kahila, J., Tedre, M., Valtonen, T. and Vartiainen, H. (2020) “Co-Designing Machine Learning Apps in K–12 With Primary School Children”, *Proceedings of the IEEE International Conference on Advanced Learning Technologies (ICALT)*. pp 308–310.

Turchi, T., Fogli, D. and Malizia, A. (2019) “Fostering computational thinking through collaborative game-based learning”, *Multimedia Tools and Applications,* Vol 78, pp 13649–13673.

Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S.D., Tegmark, M. and Fuso Nerini, F. (2020) "The role of artificial intelligence in achieving the Sustainable Development Goals", *Nature Communications,* Vol 11, Article No. 233.

Voulgari, I., Zammit, M., Stouraitis, E., Liapis, A. and Yannakakis, G.N. (2021) "Learn to Machine Learn: Designing a Game Based Approach for Teaching Machine Learning to Primary and Secondary Education Students", *Proceedings of the ACM International Conference on Interaction Design and Children*. ACM.

Waddington, D.I. and Fennewald, T. (2018) "Grim FATE: Learning About Systems Thinking in an In-Depth Climate Change Simulation", *Simulation & Gaming,* Vol 49, pp 168–194.

Watkins, C.J.C.H. and Dayan, P. (1992) "Q-learning", *Machine Learning,* Vol 8, No. 3, pp 279–292.

Webb, M.E., Fluck, A., Magenheim, J., Malyn-Smith, J., Waters, J., Deschênes, M. and Zagami, J. (2020) "Machine learning for human learners: opportunities, issues, tensions and threats", *Education Technology Research and Development.*

Zimmermann-Niefield, A., Turner, M., Murphy, B., Kane, S.K. and Shapiro, R.B. (2019) "Youth Learning Machine Learning through Building Models of Athletic Moves", *Proceedings of the ACM International Conference on Interaction Design and Children*. ACM, pp 121–132.