

Minecraft as a Platform for Project-Based Learning in AI

Sameer Singh

Department of Computer Science
University of California, Irvine
sameer@uci.edu

Abstract

Undergraduate courses that focus on open-ended, project-based learning teach students how to define concrete goals, transfer conceptual understanding of algorithms to code, and evaluate/analyze/present their solution. However, AI, along with machine learning, is getting increasingly varied in terms of both the approaches and applications, making it challenging to design project courses that span a sufficiently wide spectrum of AI. For these reasons, existing AI project courses are restricted to a narrow set of approaches (e.g. only reinforcement learning) or applications (e.g. only computer vision).

In this paper, we propose to use Minecraft as the platform for teaching AI via project-based learning. Minecraft is an open-world sandbox game with elements of exploration, resource gathering, crafting, construction, and combat, and is supported by the Malmo library that provides a programmatic interface to the player observations and actions at various levels of granularity. In Minecraft, students can design projects to use approaches like search-based AI, reinforcement learning, supervised learning, and constraint satisfaction, on data types like text, audio, images, and tabular data. We describe our experience with an open-ended, undergraduate AI projects course using Minecraft that includes 82 different projects, covering themes that ranged from navigation, instruction following, object detection, combat, and music/image generation.

Introduction

Project-based learning (Markham 2011; McManus and Costello 2019) is an effective pedagogical tool to train students via open-ended, hands-on team projects. These projects not only hone technical skills, by training the students to identify, design, and implement appropriate solutions, but also encourage self-confidence and creativity (they need to define their own project), and provide invaluable leadership and social/team-working experience. However, designing a project-based course for artificial intelligence is challenging. AI is increasingly spanning a large number of different technical approaches, such as supervised learning, reinforcement learning, search-based techniques, and Bayesian methods, and application domains, such as text, images, simulated environments, and time series. Since it is difficult to design

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

projects that cover a reasonable subset of approaches and domains, existing courses instead focus on a single environment or domain, such as robotics, games, computer vision, text, etc. Although these projects are beneficial for technical training in problem-solving, they ignore the open-ended and exploratory component of project-based learning, a crucial part of applying AI to the real-world: which part of a given scenario would benefit from machine learning or AI?

To support a wide variety of open-ended projects in AI and machine learning, we propose to use Minecraft for project-based learning of AI. The sandbox nature of Minecraft provides a dynamic, open-ended environment, additionally supported by the Malmo library that allows programmatic sensing and control at multiple levels of granularity. The combination of Minecraft and Malmo provides the opportunity for the students to design problems that cover a wide spectrum of AI and machine learning techniques, at an appropriate level of challenge. We will describe our design of an undergraduate AI projects course to use Malmo/Minecraft, and briefly mention some of the projects the students proposed, covering themes that ranged from navigation, instruction following, object detection, combat, and music/image generation.

Minecraft, and Malmo

Minecraft is an open-world sandbox game with elements of exploration, resource gathering, crafting, construction, and combat. We select it as a platform for AI projects primarily because of how open it is; we want the students to imagine the many possibilities of AI and machine learning in environments that attempt to mimic our own. Even though it is a game, the environment itself is quite complex and dynamic, with lots of possible actions and inputs, thus making it quite complex as a platform for AI. Using Minecraft also makes it much more attractive and accessible to a diverse group of students, especially since it transcends beyond the violence and aggression associated with traditional computer games.

Malmo (Johnson et al. 2016) is an AI experimentation platform built on top of Minecraft, providing programmatic access to the in-game player via a loop that supports observations and actions. In particular, Malmo provides the player agent observations that include: (1) grid view of the environment, (2) raw pixels of the agent's view, (3) objects and

distances of the visible entities, and (4) world state, like inventory, weapon status, health, etc. For control, the player can perform the following actions: (1) create the world, (2) grid-aligned movement, (3) continuous movement (e.g. speed and direction), and (4) commands for inventory, crafting, combat, and building structures. This variety of observations and actions, such as support for both grid-world and continuous state/action spaces, provide the flexibility for students to decide what challenge they want to focus on. There are a number of existing competitions that use Malmo (Guss et al. 2019; Perez-Liebana et al. 2019), however they are quite targeted, focusing less on the open nature of Minecraft but instead are more similar to other reinforcement learning environments with well-defined goals, such as OpenAI Gym (Brockman et al. 2016) and UnityAI (Juliani et al. 2018).

Course Description

We developed a project-based undergraduate AI course to explore the use of Minecraft platform for this purpose. The course is open-ended: the students have to define their own projects in the Minecraft/Malmo environment. As long as they are using a machine learning or an AI algorithm in a non-trivial manner, the student could choose any aspect of the game, whether it was based on structure building, navigation, inventory management, or combat. The goal of the course is for them to motivate their project goal, use appropriate approaches as the solution, evaluate/analyze their ideas, and present their implementation to the class.

The primary requirement for the course was to carry out a project in Minecraft as part of the team, with a proposal at the beginning, status report, and a final report. Since we wanted to impart aspects of real-world software engineering and communication, the *reports* were in form of a public webpage of their project they had to maintain, along with a public repository of their code (on Github) and summary of their ideas as videos (uploaded on YouTube).

This 10-week course has been offered twice at UC Irvine, in 2017 and 2019, with 108 and 130 undergraduate students respectively. Most of the students were juniors or seniors majoring in CS, and had taken undergraduate courses in AI and machine learning (prerequisites for our course). Each project contained maximum 3 students, and we totaled 82 projects over the two offerings. A third offering of the course is currently underway at the time of writing (Fall 2019).

Example Projects

We observed a wide variety of AI projects as a result of these offerings; we briefly highlight a few interesting ones here.

- Implemented Mario *inside* Minecraft, and used Deep RL for automated side-scrolling platforming,
- Speech interface for the player (e.g. “chop that tree”),
- Automated generation of in-game artifacts, such as using GANs for realistic trees, and RNNs for music generation,
- Archery by controlling angle and force to hit moving targets at different speeds and distances,
- Identifying plan for most effective planting of seeds and harvesting of crops for maximizing the crop yield,

- Implementing more complex versions of existing AI assignments, such as Wumpus world, the cliff walking problem in RL, and PacMan (DeNero and Klein 2010), and
- Real-time object detection and classification from pixels. The experience affirms the need for flexible, open-ended projects; although the students struggled initially to define concrete goals for their projects, they eventually pushed themselves to produce diverse and technically-advanced projects. All the project, with their videos, websites, and source code, is available at <http://sameersingh.org/courses/malmo>.

Conclusions

We argue for using Minecraft for project-based learning of AI in undergraduate courses. The flexibility and open-world nature of Minecraft, which involves building, crafting, combat, and navigation, combined with the different granularity of control and sensing provided by Malmo, results in a platform that supports many data types (language, vision, symbolic states, and real-time control) and technical approaches (supervised learning, reinforcement learning, and planning). We highlight the variety of student-designed projects in two undergraduate offerings, totaling 238 students and 82 projects.

Acknowledgements

I would like to thank Microsoft Research for providing and supporting Malmo, Moshe Lichman for his help in the first offering, and other staff that helped made this course possible, Zhe Wang, Stephen McAleer, and Yasaman Razeghi.

References

- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI gym. In *arXiv:1606.01540*.
- DeNero, J., and Klein, D. 2010. Teaching introductory artificial intelligence with Pac-Man. In *AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*.
- Guss, W. H.; Codel, C.; Hofmann, K.; Houghton, B.; Kuno, N. S.; Milani, S.; Mohanty, S.; Liebana, D. P.; Salakhutdinov, R.; Topin, N.; Veloso, M.; and Wang, P. 2019. The MineRL Competition on Sample Efficient Reinforcement Learning using Human Priors. In *NeurIPS*.
- Johnson, M.; Hofmann, K.; Hutton, T.; Bignell, D.; and Hofmann, K. 2016. The Malmo Platform for Artificial Intelligence Experimentation. In *IJCAI*.
- Juliani, A.; Berges, V.-P.; Vckay, E.; Gao, Y.; Henry, H.; Mattar, M.; and Lange, D. 2018. Unity: A general platform for intelligent agents. In *arXiv:1809.02627*.
- Markham, T. 2011. Project based learning a bridge just far enough. *Teacher librarian* 39(2):38.
- McManus, J. W., and Costello, P. J. 2019. Project based learning in computer science: A student and research advisor’s perspective. *J. Comput. Sci. Coll.* 34(3):38–46.
- Perez-Liebana, D.; Hofmann, K.; Mohanty, S. P.; Kuno, N. S.; Kramer, A.; Devlin, S.; Gaina, R. D.; and Ionita, D. 2019. The Multi-Agent Reinforcement Learning in MalmÖ (MARLÖ) Competition. In *arXiv:1901.08129*.