

---

# Hearts Gym: Learning Reinforcement Learning as a Team Event

---

Jan Ebert<sup>1</sup> Danimir T. Donevic<sup>2,3</sup> Ramona Kloß<sup>4</sup> Stefan Kesselheim<sup>1</sup>

## Abstract

Amidst the COVID-19 pandemic, the authors of this paper organized a Reinforcement Learning (RL) course for a graduate school in the field of data science. We describe the strategy and materials for creating an exciting learning experience despite the ubiquitous Zoom fatigue and evaluate the course qualitatively. The key organizational features are a focus on a competitive hands-on setting in teams, supported by a minimum of lectures providing the essential background on RL. The practical part of the course revolved around *Hearts Gym*, an RL environment for the card game *Hearts* that we developed as an entry-level tutorial to RL. Participants were tasked with training agents to explore reward shaping and other RL hyperparameters. For a final evaluation, the agents of the participants competed against each other.

## 1. Introduction

In the last decades, many subfields of Machine Learning (ML) have made impressive progress and drawn the attention of many scientists who are transforming their disciplines by applying and developing ML methods. However, Reinforcement Learning (RL), while producing notorious headlines about breakthroughs of artificial intelligence, lacks broader application to ordinary scientific problems because its implementation and adaptation require proficient practitioners. While teaching material and self-teaching opportunities for Supervised Learning in particular are omnipresent, opportunities to get started and collect experience with RL in a guided setting are scarce.

With *Hearts Gym*, we created an opportunity for such hands-

---

<sup>1</sup>Jülich Supercomputing Centre, Forschungszentrum Jülich, Jülich, Germany <sup>2</sup>Institute of Energy and Climate Research – Energy Systems Engineering (IEK-10), Forschungszentrum Jülich GmbH, 52425 Jülich, Germany <sup>3</sup>RWTH Aachen University, 52062 Aachen, Germany <sup>4</sup>IBG-1: Biotechnology, Forschungszentrum Jülich, Jülich, Germany. Correspondence to: Stefan Kesselheim <s.kesselheim@fz-juelich.de>.

on experience in the form of a training course for the students of the graduate school HDS-LEE (Helmholtz School for Data Science in Life, Earth and Energy) located at Forschungszentrum Jülich (FZJ, Jülich Research Centre). This course has been established during the COVID-19 pandemic, a time during which virtually all academic staff has been confined to home office, resulting in low engagement and tiredness from long hours of video conferencing, so-called *Zoom fatigue* (independent of the tool) (McCulloch, 2020). As an antidote to this Zoom fatigue, we decided to maximize the focus on hands-on experience and to further spice up the planned training course by adding the element of competition. The course concept was proposed by the HDS-LEE students, and implemented in the frame of a *voucher*, a support function offered by the AI consultants of Helmholtz AI, a platform fostering the use of AI in science.

The Hearts Gym workshop is an RL course based on the multi-player card game Hearts. Around 40 course participants were distributed in groups and trained RL agents that would play Hearts against agents based on manually implemented rules and other RL agents. For evaluation, at the end of the event, the agents were put to competition in two semifinals and a final, and an overall winning team was picked. In this paper, we describe the concept of the training event, the experience, the feedback, and provide an outlook to future events. We hope that this description motivates others to use similar educational concepts. Our teaching material is available under a GPLv2 license on GitHub<sup>1</sup>.

This paper is organized as follows. In Sections 2 and 3, we review related RL teaching materials and the basic rules of the game Hearts. The format of the Hearts Gym event is outlined in Section 4. Section 5 provides an overview over the applied teaching material. Observations and feedback from the event are summarized in Section 6. We conclude with an outlook on future development and application of the course concept.

## 2. Related Approaches

A large number of online resources exist to get familiar with the theoretical background on RL, e.g., teaching materials

---

<sup>1</sup><https://github.com/HelmholtzAI-FZJ/Hearts-Gym>

shared by universities or tutorials provided on the websites of ML software packages. Here, we focus on practical materials, i.e., (libraries of) RL environments made available to the public such that the user can focus on specifying the agent and the reward and training the agent, and sometimes specifying the environment. Typically such environment libraries are offered by universities, big IT companies with strong ties to AI research, or independent projects.

Most similar to Hearts Gym is *RLCard* (Zha et al., 2019). *RLCard* provides open-source tools for RL research in card games. The project highlights the relevance of card games for current RL research questions, as card games constitute a challenging but ultimately intuitively understandable multi-agent, imperfect information setting. Hearts is not readily available through *RLCard*. Other RL libraries that feature an interface for card game environments are *OpenSpiel* and *PettingZoo*. *OpenSpiel* (Lanctot et al., 2019) is a collection of RL environments and algorithms for training agents for a plethora of games, including trademark board games. Regarding card games, several well-known games such as *Bridge* are featured. *PettingZoo* (Terry et al., 2020) is another library featuring many RL environments that focuses on environments allowing for *multi-agent* RL, including some card games. In some cases, *PettingZoo* acts as a multi-agent extension to *Gym*. *OpenAI Gym* is one of the most famous and fundamental RL libraries, implementing a common environment API (Brockman et al., 2016). *Gym* provides close to a hundred RL environments, even without counting third-party environments. Notable categories are the famous Atari games (Mnih et al., 2015), environments built on the physics engine *MuJoCo* (Todorov et al., 2012), e.g., to simulate robots, and *cart pole*, a typical baseline environment and first project in RL courses.

Other noteworthy public RL environments from purely educational providers are a 1 vs. 1 snowball fight from a deep RL course by HuggingFace (Siminoni, 2021) and *Deep Traffic* (Fridman et al., 2018). *Deep Traffic* appears to have a comparable scope to Hearts Gym. It provides an environment and an agent template to train car agents that navigate a busy highway road as fast as possible. Users have to tune hyperparameters, among which the most important is the neural network for Deep Q-Learning. Further, *Deep Traffic* features a 2D visualization of the environment and a leader board in its Git repository.

Especially for self-driving cars, RL environments with modern computer graphics exist. For instance, *AirSim* (Shah et al., 2017) is a simulator for autonomous vehicles built with a game design engine. In *AirSim*, researchers can try out algorithms for computer vision and reinforcement learning simultaneously. Further, *DeepRacer* by Amazon (Balaji et al., 2019) is a cloud-based 3D race simulator that has its own global competitive leagues. The element of competition

is also central in *AIArena* (Kurbatov et al., 2021), an environment and platform for competition between AI agents in the computer game *StarCraft II*. The participating agents can be trained, e.g., with DeepMind’s environment *PySC2* (Vinyals et al., 2017).

### 3. About the Game of Hearts

Hearts is a simple 4-player trick-taking card game played with a regular deck of 52 cards in which the goal of a player is to *avoid* collecting penalty cards. The essential rules are stated here. Following suit is mandatory and the highest card played in the same suit as the first card of a trick wins the trick. Penalty cards include all 13 cards of hearts (1 penalty each) and the queen of spades (13 penalty!). The suit hearts can only be the leading suit once it has been played in a preceding trick in the game by a player who was not able to follow suit. However, no penalty cards, including hearts, can be played in the first trick, in which the 2 of clubs always leads. We note for the sake of completeness that Hearts has some additional complicating rules which enrich the strategic options of players and therefore offer challenges for advanced RL agents. One of these opens up a high-risk, high-reward strategy with extremely sparse reward information for RL agents.

The basic strategies of Hearts reveal themselves intuitively to human players, providing starting points for reward shaping. Additionally, the implementation of competitive rule-based agents that can serve as benchmarks is a rather straightforward practice. As a further advantage, with just 13 tricks to play to finish a game, computing a full game’s policy gradients, used for training, has manageable cost for personal computers. Moreover, Hearts can be played by participants *in vivo*, in preparation for and during the event, to get to know each other and to get a first feeling for the game. Finally, being a 4-player game, various competitions between participants can conclude the course. For these reasons, Hearts is a great game for a hands-on RL course.

### 4. Event Format

To enable strong engagement and enable a positive social experience, the chosen course format is largely based on a hackathon with few introductory lectures spread across the initial phase. During four core hours in the mornings of four subsequent days, the students worked on the RL algorithms. The afternoons were deliberately left vacant as the strong focus of real-world hackathons cannot be established in an online format. This free time was used for work on the hackathon only by a subset of the participants. The participants of the event were graduate students with a focus in applied machine learning. Therefore, it was possible to assume a decent background in mathematics and general ML.

The aim behind the selected teaching material was to convey basic concepts of RL, such as what environments, actions, rewards, and the ideas behind RL algorithms are, e.g., maximizing cumulative future reward. It is fair to say that the priority of the course was to provide a hands-on and social experience rather than in-depth understanding of the applied RL methods.

To this end, the participants watched a 15-minute video *An introduction to Reinforcement Learning* as a quick-start (Arxiv Insights, 2018). This knowledge was deepened in a reading phase of the section *Part 1: Key Concepts in RL* of OpenAI’s deep RL introduction *Spinning Up* (OpenAI, 2018). Furthermore, the participants were referred to the freely available, standard RL textbook of Sutton and Barto (Sutton & Barto, 2018), two UCL courses on RL (Silver, 2015; van Hasselt, 2018), and two sections with practical RL tips and additional resources of the online documentation of the software package *Stable Baselines3* (Raffin et al., 2021). This selection, in combination with the ready-to-go software environment, enabled students to get started quickly, without over-simplifying the field.

In addition to the self-study material, we demonstrated the source code of our software environment (explained in more detail in Section 5) and its structure in a *code safari*. To get familiar with the card game, the rules were explained to the participants. As an ice breaker and to reinforce understanding, participants were encouraged to start their group work by playing an online Hearts game. We scheduled the remainder of the available time for preparing the Hearts-playing agents in a team setting. To this end, the participants teamed up in eight groups of 4–6 people. The groups were formed before the workshop started and were composed of at least one PhD student with advanced experience in ML and one who is proficient in Python.

The two mentors of the course were available via a group chat set up for the hackathon. In addition, the mentors occasionally checked into each group’s video conferencing room. This provided the mentors with a picture of how the hackathon is going while not overwhelming participants, but also enabled support for groups that were less inclined to ask for help on their own. Following the hackathon convention, participants could and did use the time between the scheduled sessions for further work on their agents. Submissions of agents were due during the fourth and final session. Each team could submit one RL agent and one rule-based agent, which were pitted against agents of other teams in respective tournaments with a knockout format. In each tournament round, 4 agents played several matches of Hearts against each other and the two best-performing agents advanced to the next round. Overall winners in the RL and rule-based categories were determined in the final.

## 5. Learning Environment Implementation

The learning environment was constructed following the principle that simplicity and ease of use are clearly favored over complexity and performance. Therefore, we chose a Python implementation using widespread RL libraries and minimal further dependencies. To get started on the practical side of the course, participants only required a PC, Git, and a Python distribution. Then, Hearts Gym could be retrieved from Git and its dependencies automatically downloaded.

Hearts Gym was implemented using the multi-agent environments of RLlib (Liang et al., 2018). The concept of RL frameworks entails that users provide implementation of the key elements of RL, such as environment, policy, etc. The framework defines the relevant interfaces between these modules, in order to obtain a composable architecture. Dependencies were deliberately kept minimal in order to have an easily maintainable environment with the least amount of possible failure modes. To prevent technical issues during the event, two test runs with volunteer participants were done in preparation. These helped sort out portability errors and resulted in first feedback and according improvements.

The learning environment thus provided the required implementations for a wide range of possible policies. Hence, the participants were able to explore a lot of different strategies simply by adjusting a configuration file. This also allows participants with a weaker background or strong limitations regarding the available time to take full participation in the activity. With intuitive alterations of the configuration file, participants were already able to make successful improvements to the agents. More advanced learners could dive deeper into the code and significantly change the training procedures, implement new rule-based policies, or make even more complex alterations. One complex modification would be the implementation of new algorithms. This could be encompassed as a learning exercise in a more advanced version of the hackathon.

The most important and fundamental code modification was the implementation of a custom reward function. This reflects the fact that *reward shaping* is a key ingredient for successful RL. In addition, participants were encouraged to implement their own rule-based agents. This was possible with relative ease as we also implemented a simple deterministic policy which could serve as a skeleton for participant implementations.

For the remote evaluation, a client-server structure was created. This also allowed teams to set up their own servers and evaluate their agents against one another. One challenge was the collection of results in a relatively short time frame, while still allowing for internet connections with some latency. To this end, on the one hand, we implemented an auto-kick-system that replaces players who take too long to

respond with randomly acting bots. On the other hand, many games are played and communicated simultaneously. The servers stayed active for the duration of the course. Some teams used the servers as “sparring grounds” to coordinate evaluations of their agents against each other even during the hackathon, which improved communication between groups. At the end of the event, the servers were used as an arena for the finals, for which evaluation was carried out live in a video conference.

## 6. Event Evaluation

One of the authors participated in the course and the hackathon and we interviewed four additional participants from two groups. Here, we qualitatively summarize their feedback. This qualitative feedback closely matches the quantitative feedback collected from 17 participants.

All of the interviewed participants found the interactive group setting of the course appropriate, given the pandemic-induced home office situation. Even though the online format allowed minor distractions from the hackathon, the engagement level was high. The competitive element of the course contributed positively to this engagement level. However, the limited time frame led some participants to disregard the competitive element. Participants highly approved of the learning experience provided by the Hearts Gym hackathon with regard to high-level understanding of RL. Studying the materials and participating in the hackathon, participants believed that they learned enough about RL to competently discuss ideas of application of RL to their field. In addition, some participants consider the experience gathered sufficient to plan or even conduct an RL project on their own. At least, the event has revealed existing gaps in the knowledge of individuals. The setup of Hearts Gym worked for all users. However, participants felt that proficiency in Python was a requirement to navigate the code base in the limited time. Further, differing expertise in Python tended to decrease collaboration within the groups. In consequence, the practical learning success during the hackathon varied from group to group. One of the teams invested more than the allotted time of four half days on the practical task and stated a noticeable day-to-day progress in handling the tools available in Hearts Gym as a result of their effort. Other participants experienced flatter learning curves due to long feedback loops of RL training and a high amount of options that could be manipulated. On top of that, inefficient group dynamics including screen-sharing of code and resulting pair programming slowed down progress.

As a strategic approach to the challenge posed by Hearts Gym, the interviewed participants reported initially focusing on the logic behind clever play in Hearts. Experimentation with rule-based agents was common. In fact, the best agent developed during the hackathon was a sophisti-

cated rule-based agent. Later during the hackathon, most teams focused on reward shaping and hyperparameter tuning for RL in a very exploratory fashion. With the default evaluation setting being play versus randomly acting agents, we observed a tendency among participants to target small tactical improvements for learning that would lead to better-than-random play. For instance, if an agent learned to drop isolated high cards in the first trick without penalty, they would typically perform better than average even with random play in the remaining tricks. Once the submission deadline approached, exploitation in the form of longer training times or larger policy networks for promising configurations was conducted by all teams. Some groups have used high-performance computing (HPC) resources to reduce the training times of agents and reported that the faster feedback enabled a more stringent exploration. None of the teams spent time on analyzing the individual games played by their agents in more depth. Therefore, reward shaping was exclusively driven by trial and error based on the intuition of the participants. In the final evaluation, all teams succeeded to submit agents that played distinctively better than random and a deserved winner could be crowned among them.

## 7. Conclusion and Outlook

Based on the experience and the feedback, the event can be considered a success. With moderate time effort, the course participants became familiar with the key concepts of RL and had a group experience that was as immersive as the COVID pandemic allowed for.

With small tuning, the event concept can be adjusted for in-presence hackathon-like events. The duration can be altered to 2–3 full days. In-presence activities allow for much better team-building activities to strengthen participant interactions. In addition, scientific interaction, such as a poster session, have proven to be positive in other events organized by the authors. Especially in-person, it is important to speed up the feedback cycle. Providing access to HPC center resources with many-core GPUs is then a must. We hope that the release of our material helps other teams to set up such activities. Interested persons are encouraged to contact us.

Another exciting aspect would be integrating the challenge into a challenges platform such as the Helmholtz Data Challenges<sup>2</sup>. Here, we imagine that agents submitted as code to such a challenges platform would be evaluated automatically with new submissions. This convenient long-time evaluation environment could create a setting of continuous progress just as other data challenges have provided.

---

<sup>2</sup><https://helmholtz-data-challenges.de>



## 8. CRediT Author Statement

This section states the authors main roles following Contributor Role Taxonomy CRediT (Allen et al., 2019).

**Jan Ebert:** Conceptualization, Validation, Methodology, Software, Resources, **Danimir Doncevic:** Writing – Original Draft, Conceptualization, Evaluation, Literature Review, **Ramona Kloß:** Coordination, Administration, **Stefan Kesselheim:** Conceptualization, Supervision

## Acknowledgements

This work was partially performed as part of the Helmholtz School for Data Science in Life, Earth and Energy (HDS-LEE) and received funding from the Helmholtz Association of German Research Centres.

We would like to thank Ji Gao, Daniel Mayfrank, Felix Rauh, and Luisa Schuhmacher for participating in trial runs ahead of the event. Their feedback was extremely valuable for making the event run more smoothly.

## References

- Allen, L., O’Connell, A., and Kiermer, V. How can we ensure visibility and diversity in research contributions? how the contributor role taxonomy (credit) is helping the shift from authorship to contributorship. *Learned Publishing*, 32(1):71–74, 2019.
- Arxiv Insights. An introduction to reinforcement learning, 2018. URL <https://www.youtube.com/watch?v=JgvyzIkgxF0>.
- Balaji, B., Mallya, S., Genc, S., Gupta, S., Dirac, L., Khare, V., Roy, G., Sun, T., Tao, Y., Townsend, B., et al. Deep-racer: Educational autonomous racing platform for experimentation with sim2real reinforcement learning. *arXiv preprint arXiv:1911.01562*, 2019.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym, 2016.
- Fridman, L., Terwilliger, J., and Jenik, B. Deeptraffic: Crowdsourced hyperparameter tuning of deep reinforcement learning systems for multi-agent dense traffic navigation. In *Neural Information Processing Systems (NIPS 2018) Deep Reinforcement Learning Workshop*, 2018. doi: 10.5281/zenodo.2530457. URL <http://arxiv.org/abs/1801.02805>.
- Kurbatov, A., Clarke, M., van Schoor, D., Chin, N., and Urdesic, V. Sc2 ai arena, 2021. URL <https://aiarena.net/>.
- Lanctot, M., Lockhart, E., Lespiau, J.-B., Zambaldi, V., Upadhyay, S., Pérolat, J., Srinivasan, S., Timbers, F., Tuyls, K., Omidshafiei, S., Hennes, D., Morrill, D., Muller, P., Ewalds, T., Faulkner, R., Kramár, J., Vylder, B. D., Saeta, B., Bradbury, J., Ding, D., Borgeaud, S., Lai, M., Schrittwieser, J., Anthony, T., Hughes, E., Danihelka, I., and Ryan-Davis, J. OpenSpiel: A framework for reinforcement learning in games. *CoRR*, abs/1908.09453, 2019. URL <http://arxiv.org/abs/1908.09453>.
- Liang, E., Liaw, R., Nishihara, R., Moritz, P., Fox, R., Goldberg, K., Gonzalez, J., Jordan, M., and Stoica, I. RLlib: Abstractions for distributed reinforcement learning. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3053–3062. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/liangl18b.html>.
- McCulloch, G. A mission to make virtual parties actually fun, Nov 2020. URL <https://www.wired.com/story/zoom-parties-proximity-chat/>.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533, 2015.
- OpenAI. Spinning up, 2018. URL <https://spinningup.openai.com>.
- Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., and Dormann, N. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL <http://jmlr.org/papers/v22/20-1364.html>.
- Shah, S., Dey, D., Lovett, C., and Kapoor, A. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics*, 2017. URL <https://arxiv.org/abs/1705.05065>.
- Silver, D. Lectures on reinforcement learning, 2015. URL <https://www.davidsilver.uk/teaching/>.
- Simonini, T. MI agents snowball fight 1v1, 2021. URL <https://huggingface.co/ThomasSimonini/ML-Agents-SnowballFight-1vs1>.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Terry, J. K., Black, B., Grammel, N., Jayakumar, M., Hari, A., Sullivan, R., Santos, L., Perez, R., Horsch, C., Diefendahl, C., Williams, N. L., Lokesh, Y., Sullivan, R., and Ravi, P. Pettingzoo: Gym for multi-agent reinforcement learning. *arXiv preprint arXiv:2009.14471*, 2020.

Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.

van Hasselt, H. Reinforcement learning lecture series 2018, 2018. URL <https://www.deepmind.com/learning-resources/reinforcement-learning-lecture-series-2018>.

Vinyals, O., Ewalds, T., Bartunov, S., Georgiev, P., Vezhn-evets, A. S., Yeo, M., Makhzani, A., Küttler, H., Agapiou, J., Schrittwieser, J., et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.

Zha, D., Lai, K.-H., Cao, Y., Huang, S., Wei, R., Guo, J., and Hu, X. Rlcard: A toolkit for reinforcement learning in card games. *arXiv preprint arXiv:1910.04376*, 2019.