

Pommerman: A Multi-Agent Playground

Cinjon Resnick*
NYU

Wes Eldridge
Rebellious Labs

David Ha
Google Brain

Denny Britz
Stanford University

Jakob Foerster
University of Oxford

Julian Togelius
NYU

Kyunghyun Cho and Joan Bruna
NYU, FAIR

Abstract

We present Pommerman, a multi-agent environment based on the classic console game Bomberman. Pommerman consists of a set of scenarios, each having at least four players and containing both cooperative and competitive aspects. We believe that success in Pommerman will require a diverse set of tools and methods, including planning, opponent/teammate modeling, game theory, and communication, and consequently can serve well as a multi-agent benchmark. To date, we have already hosted one competition, and our next one will be featured in the NIPS 2018 competition track.

Why Pommerman

In this section, we provide our motivation and goals for both the Pommerman benchmark and the NIPS 2018 competition. Currently, there is no consensus benchmark involving either general sum game settings or settings with at least three players. Instead, recent progress has focused on two player zero-sum games such as Go and Chess. We believe that the Pommerman environment can assume this role for multi-agent learning. Additionally, we are organizing competitions for Pommerman because we believe that they are a strong way to push forward the state of the art and can contribute to lasting results for years to come.

Multi-Agent Learning

Historically, a majority of multi-agent research has focused on zero-sum two player games. For example, computer competitions for Poker and Go over the past fifteen years have been vital for developing methods culminating in recent superhuman performance (Moravčík et al. 2017; Noam Brown 2017; Bowling et al. 2017; Silver et al. 2016). These benchmarks have also lead to the discovery of new algorithms and approaches like Monte Carlo Tree Search (Vodopivec, Samothrakis, and Šter 2017; Browne et al. 2012; Kocsis and Szepesvári 2006; Coulom 2006) and Counterfactual Regret Minimization (Zinkevich et al. 2008).

We believe that an aspect restraining the field from progressing towards general-sum research and scenarios with

more than two players is the lack of suitable environments. We propose Pommerman as a solution.

Pommerman is stylistically similar to Bomberman (Bomberman 1983), the famous game from Nintendo. At a high level, there are at least four agents all traversing a grid world. Each agent’s goal is to have their team be the last remaining. They can plant bombs that, upon expiration, destroy anything (but rigid walls) in their vicinity. It contains both adversarial and cooperative elements. The Free-For-All (FFA) variant has at most one winner and, because there are four players, encourages research directions that can handle situations where the Nash payoffs are not all equivalent. The team variants encourage research with and without explicit communication channels, including scenarios where the agent has to cooperate with previously unseen teammates. The latter is a recently burgeoning subfield of multi-agent learning (Foerster et al. 2016; Resnick et al. 2018d; Evtimova et al. 2017; Foerster et al. 2017; Lewis et al. 2017; Mordatch and Abbeel 2017; Lazaridou et al. 2018) with established prior work as well (Steels 1999; 2003; Levy and Kirby 2006; Fehervari and Elmenreich 2010), while the latter has been underexplored.

We aim for the Pommerman benchmark to provide for multi-agent learning what the Atari Learning Environment (Bellemare et al. 2013) provided for single-agent reinforcement learning and ImageNet (Deng et al. 2009) for image recognition. Beyond game theory and communication, Pommerman can also serve as a testbed for research into reinforcement learning, planning, and opponent/teammate modeling.

RoboCup Soccer (Nardi et al. 2014) is a similar competition that has been running since 1997. There, eleven agents per side play soccer. Key differences between Pommerman and RoboCup Soccer are:

1. Pommerman includes an explicit communication channel. This changes the dynamics of the game and adds new research avenues.
2. Pommerman strips away the sensor input, which means that the game is less apt for robotics but more apt for studying other aspects of AI, games, and strategy.
3. Pommerman uses low dimensional, discrete control and input representations instead of continuous ones. We believe this makes it easier to focus on the high level strate-

*Correspondence: cinjon@nyu.edu

gic aspects rather than low level mechanics.

4. In team variants, the default Pommerman setup has only two agents per side, which makes it more amenable to burgeoning fields like emergent communication which encounter training difficulties with larger numbers of agents.
5. Pommerman’s FFA variant promotes research that does not reduce to a 1v1 game, which means that a lot of the theory underlying such games (like RoboCup Soccer) is not applicable.

The second, third, and fourth differences above are a positive or negative trade-off depending on one’s research goals.

Another, more recent, benchmark is Half-Field Offense (Hausknecht et al. 2016), a modification of RoboCup that reduces the complexity and focuses on decision-making in a simplified subtask. However, unlike the FFA scenario in Pommerman, Half-Field Offense is limited to being a zero-sum game between two teams.

In general, the communities that we want to attract to benchmark their algorithms have not gravitated towards RoboCup but instead have relied on a large number of one-off toy tasks. This is especially true for multi-agent Deep RL. We think that the reasons for that could be among the five above. Consequently, Pommerman has the potential to unite these communities, especially when considering that future versions can be expanded to more than four agents.

High Quality Benchmark

There are attributes that are common to the best benchmarks beyond satisfying the community’s research direction. These include having mechanics and gameplay that are intuitive for humans, being fun to play and watch, being easy to integrate into common research setups, and having a learning problem that is not *too* difficult for the current state of method development. Most games violate at least one of these. For example, the popular game Defense of the Ancients (OpenAI 2018) is intuitive and fun, but extremely difficult to integrate. On the other hand, the card game Bridge is easy to integrate, but it is not intuitive; the gameplay and mechanics are slow to learn and there is a steep learning curve to understanding strategy.

Pommerman satisfies these requirements. People have no trouble understanding basic strategy and mechanics. It is fun to play and to watch, having been developed by Nintendo for two decades. Additionally, we have purposefully made the state input based not on pixel observations but rather on a symbolic interpretation so that it does not require large amounts of compute to build learning agents.

Research game competitions disappear for two reasons - either the administrators stop running it or participants stop submitting entrants. This can be due to the game being ‘solved’, but it could also be because the game just was not enjoyable or accessible enough. We view Pommerman as having a long life ahead of it. Beyond the surface hyper-parameters like board size and number of walls, early forays suggest that there are many aspects of the game that can be modified to create a rich and long lasting research challenge and competition venue. These include partial observability

of the board, playing with random teammates, communication among the agents, adding power-ups, and learning to play with human players.

These potential extensions, and the fact that N-player learning by itself has few mathematical guarantees, suggest that Pommerman will be a challenging and fruitful testbed for years to come.

There are, however, limitations to this environment. One difficulty is that a local optimum arises where the agent avoids exploding itself by learning to never use the bomb action. In the long term, this is ineffective because the agent needs to use the bomb to destroy other agents. Players have successfully solved this challenge (Resnick et al. 2018a), but it is an aspect of basic gameplay that has to be handled in order for the multi-agent research benefits to become apparent.

Description

In this section, we give details of the Pommerman environment. Note that all of the code to run the game and train agents can be found in our git repository (Resnick et al. 2018b), while our website (pommerman.com) contains further information on how to submit agents.

Game Information




Figure 1: Pommerman start state. Each agent begins in one of four positions. Yellow squares are wood, brown are rigid, and the gray are passages.

As previously mentioned, Pommerman is stylistically similar to Bomberman. Every battle starts on a randomly drawn symmetric 11x11 grid (‘board’) with four agents, one in each corner. Teammates start on opposite corners.

In team variants, the game ends when both players on one team have been destroyed. In FFA, it ends when at most one agent remains alive. The winning team is the one who has remaining members. Ties can happen when the game does not end before the max steps or if the last agents are destroyed on the same turn. If this happens in competitions, we will rerun the game. If it reoccurs, then we will rerun the game with collapsing walls until there is a winner. This is a variant where, after a fixed number of steps, the game board becomes smaller according to a specified cadence. We have a working example in the repository.

Besides the agents, the board consists of wooden and rigid walls. We guarantee that the agents will have an accessible

Game	Intuitive?	Fun?	Integration?
Bridge	1	3	5
Civilization	2	3	1
Counterstrike	5	5	2
Coup	4	5	5
Diplomacy	1	4	3
DoTA	3	5	2
Hanabi	2	3	5
Hearthstone	1	4	1
Mario Maker	4	5	3
Pommerman	5	4	5
PUBG	5	5	1
Rocket League	5	4	1
Secret Hitler	4	4	3
Settlers of Catan	4	3	3
Starcraft 2	3	5	5
Super Smash	5	5	1

Table 1: Comparing multi-agent games along three important axes for uptake beyond whether the game satisfies the community’s intended research direction. Attributes are considered on a 1-5 scale where 5 represents the highest value. Fun takes into account both watching and playing the game. The Intuitive and Fun qualities, while subjective, are noted because they have historically been factors in whether a game is used in research.

path to each other. Initially, this path is occluded by wooden walls. See Figure 1 for a visual reference.

Rigid walls are indestructible and impassable. Wooden walls can be destroyed by bombs. Until they are destroyed, they are impassable. After they are destroyed, they become either a passage or a power-up.

On every turn, agents choose from one of six actions:

1. *Stop*: This action is a pass.
2. *Up*: Move up on the board.
3. *Left*: Move left on the board.
4. *Down*: Move down on the board.
5. *Right*: Move right on the board.
6. *Bomb*: Lay a bomb.

Additionally, if this is a communicative scenario, then the agent emits a message every turn consisting of two words from a dictionary of size eight. These words are passed to its teammate in the next step as part of the observation. In total, the agent receives the following observation each turn:

- *Board*: 121 Ints. The flattened board. In partially observed variants, all squares outside of the 5x5 purview around the agent’s position will be covered with the value for fog (5).
- *Position*: 2 Ints, each in [0, 10]. The agent’s (x, y) position in the grid.
- *Ammo*: 1 Int. The agent’s current ammo.
- *Blast Strength*: 1 Int. The agent’s current blast strength.
- *Can Kick*: 1 Int, 0 or 1. Whether the agent can kick or not.
- *Teammate*: 1 Int in [-1, 3]. Which agent is this agent’s teammate. In non-team variants, this is -1.
- *Enemies*: 3 Ints in [-1, 3]. Which agents are this agent’s enemies. In team variants, the third int is -1.

- *Bomb Blast Strength*: List of Ints. The bomb blast strengths for each of the bombs in the agent’s purview.
- *Bomb Life*: List of Ints. The remaining life for each of the bombs in the agent’s purview.
- *Message*: 2 Ints in [0, 8]. The message being relayed from the teammate. Both Ints are zero only when a teammate is dead or if it is the first step. This field is not included for non-cheap talk variants.

The agent starts with one bomb (‘ammo’). Every time it lays a bomb, its ammo decreases by one. After that bomb explodes, its ammo will increase by one. The agent also has a blast strength that starts at two. Every bomb it lays is imbued with the current blast strength, which is how far in the vertical and horizontal directions that bomb will effect. A bomb has a life of ten time steps. Upon expiration, the bomb explodes and any wooden walls, agents, power-ups or other bombs within reach of its blast strength are destroyed. Bombs destroyed in this manner chain their explosions.

Power-Ups: Half of the wooden walls have hidden power-ups that are revealed when the wall is destroyed. These are:

- *Extra Bomb*: Picking this up increases the agent’s ammo by one.
- *Increase Range*: Picking this up increases the agent’s blast strength by one.
- *Can Kick*: Picking this up permanently allows an agent to kick bombs by moving into them. The bombs travel in the direction that the agent was moving at one unit per time step until they are impeded either by a player, a bomb, or a wall.

Early results

The environment has been public since late February and the competitions were first announced in late March. In that time, we have seen a strong community gather around

the game, with more than 500 people in the Discord server (<https://discord.gg/mtW7kp>) and more than half of the repository commits from open source contributors.

There have also been multiple published papers using Pommerman (Resnick et al. 2018a; Zhou et al. 2018). These demonstrate that the environment is challenging and we do not yet know what are the optimal solutions in any of the variants. In particular, the agents in (Resnick et al. 2018a) discover a novel way of playing where they treat the bombs as projectiles by laying, then kicking them at opponents. This is a strategy that not even novice humans attempt, yet the agents use it to achieve a high success rate.

Preliminary analysis suggests that the game can be very challenging for reinforcement learning algorithms out of the box. Without a very large batch size and a shaped reward (Ng, Harada, and Russell 1999), neither of Deep Q-Learning (Mnih et al. 2013) nor Proximal Policy Optimization (Schulman et al. 2017) learned to successfully play the game against the default learning agent (‘SimpleAgent’). One reason for this is because the game has a (previously mentioned) unique feature in that the bomb action is highly correlated with losing but must be wielded effectively to win.

We also tested the effectiveness of DAgger (Daumé, Langford, and Marcu 2009) in bootstrapping agents to match the SimpleAgent. We found that, while somewhat sensitive to hyperparameter choices, it was nonetheless effective at yielding agents that could play at or above the FFA win rate of a single SimpleAgent ($\sim 20\%$). This is less than chance because four simple agents will draw a large percentage of the time.

Competitions

In this section, we describe the Pommerman competitions. This includes both the upcoming NIPS 2018 event and the FFA competition that we already ran.

FFA competition

We ran a preliminary competition on June 3rd, 2018. We did not advertise this widely other than within our Discord social group (<https://discord.gg/mtW7kp>), nor did we have any prizes for it. Even so, we had a turnout of eight competitors who submitted working agents by the May 31st deadline.

The competition environment was the FFA variant (Resnick et al. 2018c) where four agents enter, all of whom are opponents. The top two agents were submitted by Görög Mrton and a team led by Yichen Gong, with the latter being the strongest.

Görög’s agent improved upon the repository’s baseline agent through a number of edits. On the other hand, Yichen’s agent was a redesign implementing a Finite State Machine Tree-Search approach (Zhou et al. 2018). They respectively won 8 and 22 of their 35 matches (with a number of the remaining being ties).

NIPS Competition

The NIPS competition will be held live at NIPS 2018 and competitors are required to submit a team of two agents by

November 21st, 2018. The featured environment will be the partially observable team variant without communication. Otherwise, we will be reusing the machinery that we developed to run the FFA competition.

Submitting Agents

We run the competitions using Docker and expect submissions to be accompanied by a Docker file that we can build on the game servers. For FFA competitions, this entails submitting a (possibly private) repository having one Docker file representing the agent. For team competitions, this means the submission should have two Docker files to represent the two agents. Instructions and an example for building Docker containers from trained agents can be found in our repository (Resnick et al. 2018b).

The agents should follow the prescribed convention specified in our example code and expose an ‘act’ endpoint that accepts the dictionary of observations. Because we are using Docker containers and http requests, we do not have any requirements for programming language or framework.

The expected response from the agent will be a single integer in $[0, 5]$ representing which of the six actions that agent would like to take. In variants with messages, we also expect two more integers in $[1, 8]$ representing the message. If an agent does not respond in an appropriate time limit for our competition constraints (100ms), then we will automatically issue them the Stop action and, if appropriate, have them send out the message (0, 0). This timeout is an aspect of the competition and not native to the game itself.

Conclusion

In this paper, we have introduced the Pommerman environment, detailed why it is a strong setup for multi-agent research, and described early results and competitions.

All of the code is readily available at our git repository (github.com/MultiAgentLearning/playground) and further information about competitions, including NIPS 2018, on our website (pommerman.com).

Acknowledgments

We are especially grateful to Roberta Raileanu, Sanyam Kapoor, Lucas Beyer, Stephan Uphoff, and the whole Pommerman Discord community for their contributions, as well as Jane Street, Facebook AI Research, Google Cloud, and NVidia Research for their sponsorship.

References

- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *J. Artif. Int. Res.* 47(1):253–279.
- Bombardier, W. 1983. Wikipedia bomberman. <https://en.wikipedia.org/wiki/Bombardier>.
- Bowling, M.; Burch, N.; Johanson, M.; and Tammelin, O. 2017. Heads-up limit hold’em poker is solved. *Commun. ACM* 60(11):81–88.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.;

- Samothrakis, S.; and Colton, S. 2012. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games* 4(1):1–43.
- Coulom, R. 2006. Efficient selectivity and backup operators in monte-carlo tree search. In van den Herik, H. J.; Ciancarini, P.; and Donkers, H. H. L. M., eds., *Computers and Games*, volume 4630 of *Lecture Notes in Computer Science*, 72–83. Springer.
- Daumé, H.; Langford, J.; and Marcu, D. 2009. Search-based structured prediction. *Machine learning* 75(3):297–325.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Evtimova, K.; Drozdov, A.; Kiela, D.; and Cho, K. 2017. Emergent language in a multi-modal, multi-step referential game. *CoRR* abs/1705.10369.
- Fehervari, I., and Elmenreich, W. 2010. Evolving neural network controllers for a team of self-organizing robots. *Journal of Robotics*.
- Foerster, J.; Assael, I. A.; de Freitas, N.; and Whiteson, S. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, 2137–2145.
- Foerster, J. N.; Nardelli, N.; Farquhar, G.; Torr, P. H. S.; Kohli, P.; and Whiteson, S. 2017. Stabilising experience replay for deep multi-agent reinforcement learning. *CoRR* abs/1702.08887.
- Hausknecht, M.; Mupparaju, P.; Subramanian, S.; Kalyanakrishnan, S.; and Stone, P. 2016. Half field offense: An environment for multiagent learning and ad hoc teamwork. In *AAMAS Adaptive Learning Agents (ALA) Workshop*.
- Kocsis, L., and Szepesvári, C. 2006. Bandit based monte-carlo planning. In *Proceedings of the 17th European Conference on Machine Learning, ECML’06*, 282–293. Berlin, Heidelberg: Springer-Verlag.
- Lazaridou, A.; Hermann, K. M.; Tuyls, K.; and Clark, S. 2018. Emergence of linguistic communication from referential games with symbolic and pixel input. In *International Conference on Learning Representations*.
- Levy, S. D., and Kirby, S. 2006. Evolving distributed representations for language with self-organizing maps. In Vogt, P.; Sugita, Y.; Tuci, E.; and Nehaniv, C. L., eds., *EELC*, volume 4211 of *Lecture Notes in Computer Science*, 57–71. Springer.
- Lewis, M.; Yarats, D.; Dauphin, Y. N.; Parikh, D.; and Batra, D. 2017. Deal or no deal? end-to-end learning for negotiation dialogues. *arXiv preprint arXiv:1706.05125*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. cite arxiv:1312.5602Comment: NIPS Deep Learning Workshop 2013.
- Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. H. 2017. Deepstack: Expert-level artificial intelligence in no-limit poker. *CoRR* abs/1701.01724.
- Mordatch, I., and Abbeel, P. 2017. Emergence of grounded compositional language in multi-agent populations. *arXiv preprint arXiv:1703.04908*.
- Nardi, D.; Noda, I.; Ribeiro, F.; Stone, P.; von Stryk, O.; and Veloso, M. 2014. RoboCup soccer leagues. *AI Magazine* 35(3):77–85.
- Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, 278–287.
- Noam Brown, T. S. 2017. Libratus: The superhuman ai for no-limit poker. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 5226–5228.
- OpenAI. 2018. Dota 2. <https://blog.openai.com/dota-2/>.
- Resnick, C.; Raileanu, R.; Kapoor, S.; Peysakhovich, A.; Cho, K.; and Bruna, J. 2018a. Backplay: “Man muss immer umkehren”. *ArXiv e-prints*.
- Resnick, C.; Eldridge, W.; Britz, D.; and Ha, D. 2018b. Playground: Ai research into multi-agent learning. <https://github.com/MultiAgentLearning/playground>.
- Resnick, C.; Eldridge, W.; Britz, D.; and Ha, D. 2018c. Pommerman ffa competition environment. <https://github.com/MultiAgentLearning/playground/blob/master/pommerman/configs.py#L20>.
- Resnick, C.; Kulikov, I.; Cho, K.; and Weston, J. 2018d. Vehicle community strategies. *CoRR* abs/1804.07178.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *CoRR* abs/1707.06347.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587):484–489.
- Steels, L. 1999. *The Talking Heads Experiment. Volume 1. Words and Meanings*. Antwerpen: Laboratorium.
- Steels, L. 2003. Evolving grounded communication for robots. *Trends in Cognitive Sciences* 7:308–312.
- Vodopivec, T.; Samothrakis, S.; and Šter, B. 2017. On monte carlo tree search and reinforcement learning. *J. Artif. Int. Res.* 60(1):881–936.
- Zhou, H.; Gong, Y.; Mugrai, L.; Khalifa, A.; Andy, N.; and Togelius, J. 2018. A hybrid search agent in pommerman. In *The International Conference on the Foundations of Digital Games (FDG)*.
- Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2008. Regret minimization in games with incomplete information. In Platt, J. C.; Koller, D.; Singer, Y.; and Roweis,

S. T., eds., *Advances in Neural Information Processing Systems 20*. Curran Associates, Inc. 1729–1736.