

The Angry Birds AI Competition

*Jochen Renz, Xiaoyu Ge,
Stephen Gould, Peng Zhang*

■ *The aim of the Angry Birds AI competition (AIBIRDS) is to build intelligent agents that can play new Angry Birds levels better than the best human players. This is surprisingly difficult for AI as it requires similar capabilities to what humans need for successfully interacting with the physical world, one of the grand challenges of AI. As such the competition offers a simplified and controlled environment for developing and testing the necessary AI technologies, a seamless integration of computer vision, machine learning, knowledge representation and reasoning, reasoning under uncertainty, planning, and heuristic search, among others. Over the past three years there have been significant improvements, but we are still a long way from reaching the ultimate aim, and thus, there are great opportunities for participants in this competition.*

The aim of the Angry Birds AI competition (AIBIRDS) is to build intelligent agents that can play novel Angry Birds levels better than the best human players. The competition was initiated in 2012 by the authors of this report and is held in collocation with some of the major AI conferences such as the International Joint Conference on Artificial Intelligence in 2013 and again in 2015, and the European Conference on Artificial Intelligence conference in 2014.

Angry Birds is a popular physics-based puzzle game developed by the company Rovio that requires players to use a slingshot to shoot birds toward green pigs protected by a physical structure (see figure 1). The actions available to players are simple, the release point (x, y) of the bird from the slingshot and the time (t) after release at which the special power of the bird is activated. A level is solved once all pigs are destroyed; most levels can be solved using up to five birds. Different birds have different behaviors and special powers, and while the player knows the order in which birds will appear on the slingshot, the player cannot manipulate this order.

While this intuitively sounds very simple, it is a very hard problem for AI as the action space is continuous and the exact outcome of each action is unknown without simulating it. The built-in physics simulator can deterministically

calculate and display the outcome of every action in real time as all game parameters are internally known. However, agents do not have access to the internal state of the game; they only get visual information in the same way a human player would. The competition uses the official Chrome browser version of Angry Birds,¹ and agents play the game like humans would. Therefore, the outcome of actions has to be estimated using methods other than brute force simulation, and a successful sequence of actions has to be selected and revised without knowing the true outcome of each action in advance.

This way of solving problems is very similar to what humans have to do every day when interacting with the physical world and what humans are very good at; that is, physical properties of entities are unknown, most information is obtained visually, the action space is continuous, and the exact outcome of actions is not known in advance. It requires an effective integration of several AI areas such as computer vision, machine learning, knowledge representation and reasoning, reasoning under uncertainty, planning, and heuristic search. In order to develop AI that can interact with the physical world, and in order to conquer this grand challenge of AI, we have to be able to solve these problems as well as humans can solve them.

The AIBIRDS competition provides a simplified and controlled environment for developing and evaluating AI for the physical world. It allows AI researchers to focus on the core problems without distractions that are present in the real world. We see this competition as the start of a series of physics simulation domains that become more and more realistic over time and that encourage the development of relevant and practical AI methods. Approaches that work in these domains would constitute an important stepping stone for developing intelligent agents that perform well in other real-world tasks.

Competition Overview

The task at the competition is to play a set of Angry Birds levels within a given time limit, typically about 3 minutes per level on average. Levels can be played and replayed in any order. Importantly, the levels are new and have not been seen during the development and training of the agents. The agents are ranked according to their combined high scores over all solved levels.

In order to make it easier for participants to build their own agents, we provide basic game-playing software² and also encourage all participants to open-source their agents. Our software only requires a Chrome browser and a Java environment and can be installed on most popular operating systems. It includes an interface to the official Angry Birds game that can take screenshots and execute mouse actions. It also includes a basic computer-vision module that

detects and categorizes all relevant game objects and a basic trajectory-planning module that calculates the release point based on a specified target point. Our software includes a sample agent (the Naive agent) that targets random pigs using random trajectories. Having a bit of randomness seems beneficial as it avoids being trapped in unsuccessful strategies. It can also help making a lucky shot. Interestingly, the Naive agent was the winner in 2012 and still outperforms about one-third of the agents in the current benchmark.

The main AIBIRDS competition runs on a client-server system where agents are not directly connected to the game but have to communicate through a given protocol with the server that runs the live games. This only allows agents to obtain screenshots and current high scores across all agents, and to submit game actions that will be executed by the server. In the past, we have offered participants the opportunity to test compatibility of their agent with our competition environment a few weeks before the competition. However this was not always successful, and therefore we will release our competition environment so participants can test it themselves and researchers and educators can run their own competition.

The AIBIRDS competition runs over several rounds, a qualification round where we select the best agents for the final rounds, followed by group stages of four agents where the two best agents of each group progress to the next round until only two agents are left. They then compete in a grand finale to determine the champion. While remote participation is possible, attendance is encouraged and well worth it. All games are streamed live with a replay of successful solutions and a live leaderboard. This makes the competition very exciting to watch for the audience and particularly for attending teams, with several heart-stopping moments and outbreaks of joy.

At the end of every competition, we invite the general public to compete against the winning agents in a human versus machine challenge where we test whether the goal of beating humans has been achieved. The performance of agents is clearly improving every year. In 2014, the best agent was already better than two-thirds of the human players.

Competition Entrants

So far, 36 teams from 16 countries have participated in our competitions. Most are academic participants and students, some from research institutes and even AI amateurs. We have seen various approaches in past competitions including logic programming (Calimeri et al 2013), qualitative reasoning (Walega, Lechowski, and Zawidzki 2014), advanced simulations (Polceanu and Buche 2013), heuristics,³ and machine-learning techniques.⁴ The winning agent in

2014 was mainly based on qualitative structural analysis.⁵ To destroy the physical structure formed by connected blocks, the agent identifies and targets the structure's weak part, which is determined by spatial relations between the building blocks. P. Zhang and J. Renz (2014), and P. Walega, T. Lechowski, and M. Zawidzki (2014) used similar agents that used spatial reasoning or qualitative physics to estimate the outcome of a shot. The best agent in 2013 was a metaagent trying to match different strategies to different levels. The strategy-level matching is modeled as a multiarmed bandit problem.⁶ In 2013, A. Narayan-Chen, L. Xu, and J. Shavlik (2013) and, in 2014, N. Tziortziotis, G. Papagiannis, and K. Blekas (2014) used Bayesian inferences and were ranked third in 2013 and second in 2014, respectively. We have also received several agents that determine shot sequences based on approximate simulation. While most simulation agents do not perform well, the best agent using advanced simulations⁷ ranked fourth in 2014.

Next Competition and New Developments

The next AIBIRDS competition will be held at the 2015 International Joint Conference on Artificial Intelligence (IJCAI 2015) in Buenos Aires, Argentina, from July 29–31, 2015. We encourage new teams to participate, as the two best teams of the 2014 competition also participated for the first time. To encourage participants to focus on estimating the outcome of actions rather than relying too much on randomness, we are introducing a new track this year where two agents play alternating shots for the same level and bid for the right to have the first or second shot. The agent with the winning shot receives all points of the level.

We have also released a version of our basic game-playing software written in Snap!, a simple visual programming language. Using this version of our framework, it is very easy for anyone to develop his or her own Angry Birds AI agent and experiment with different strategies even without prior programming experience. We envisage that it will be used to excite kids about

computing, programming, and AI and even to hold competitions in schools.

More information about all aspects of the future and past competitions can be found on our website.⁷

Notes

1. Available at chrome.angrybirds.com.
2. Angry Birds Game Playing Software Version 1.3: Basic Game Playing Software (www.aibirds.org).
3. See the 2014 team description by J. Lee, H. Seon, J. Kim, S. Joo, and K. Kim. Angry Plan A+ Team. ECAI 2014 Symposium on AI in Angry Birds (aibirds.org/2014-papers/PlanA+.pdf).
4. See the 2013 team description by A. Jutzeler, M. Katanic, and J. J. Li. Managing Luck: A Multi-Armed Bandits Meta-Agent for the Angry Birds Competition (aibirds.org/2013-Papers/Team-Descriptions/beaurivage.pdf).
5. See the 2014 team description by T. Borovicka, R. Spetlik, and K. Rymes. Datalab Birds Angry Birds AI (aibirds.org/2014-papers/datalab-birds.pdf).
6. See the 2013 team description by A. Jutzeler, M. Katanic, and J. J. Li. Managing Luck: A Multi-Armed Bandits Meta-Agent for the Angry Birds Competition (aibirds.org/2013-Papers/Team-Descriptions/beaurivage.pdf).
7. See www.aibirds.org.

References

- Calimeri, F.; M. Fink, M.; Germano, S.; Ianni, G. Redl, C.; and Wimmer, A. 2013. AngryHEX: An Artificial Player for Angry Birds Based on Declarative Knowledge Bases. In *Popularize Artificial Intelligence: Proceedings of the Workshop Popularize Artificial Intelligence*, CEUR Workshop Proceedings, volume 1107. Aachen, Germany: RWTH Aachen University (ceur-ws.org/Vol-1107/paper10.pdf).
- Polceanu, M., and Buche, C. 2013. Towards a Theory-of-Mind-Inspired Generic Decision-Making Framework. Paper presented at the IJCAI 2013 Symposium on AI in Angry Birds, Beijing, China, 3–9 August (aibirds.org/2013-Papers/Symposium/IHSEV.pdf).
- Narayan-Chen, A.; Xu, L.; and Shavlik, J. 2013. An Empirical Evaluation of Machine Learning Approaches for Angry Birds. Paper presented at the IJCAI 2013 Symposium on AI in Angry Birds, Beijing, China, 3–9 August (aibirds.org/2013-Papers/Symposium/teamwisc.pdf).
- Tziortziotis, N.; Papagiannis, G.; and Blekas, K. 2014. A Bayesian Ensemble Regression Framework on the Angry Birds Game. Paper presented at the ECAI 2014 Symposium on AI in Angry Birds, Prague, Czech Republic, 18–22 August. (aibirds.org/2014-papers/AngryBER-2014.pdf).
- Walega, P.; Lechowski, T.; and Zawidzki, M. 2014. Qualitative Physics in Angry Birds: First Results. Paper presented at the ECAI 2014 Symposium on AI in Angry Birds, Prague, Czech Republic, 18–22 August. (aibirds.org/2014-papers/ImpactVector-2014.pdf).
- Zhang, P., and Renz, J. 2014. Qualitative Spatial Representation and Reasoning in Angry Birds: The Extended Rectangle Algebra. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference*. Palo Alto, CA: AAAI Press.
- Jochen Renz** is a professor in the Research School of Computer Science at the Australian National University (ANU) and the head of the ANU Artificial Intelligence group. He received his PhD degree from the University of Freiburg in 2000 and his Habilitation degree from the Vienna University of Technology (Technische Universität Wien) in 2003. His main research interests are in qualitative spatial reasoning, both its theoretical foundations and its applications.
- XiaoYu Ge** is a Ph.D. candidate at the Australian National University. He received a bachelor degree in information technology from Monash University and a bachelor honors degree in artificial intelligence from the Australian National University. His research interests include physics reasoning, qualitative spatio-temporal reasoning, and developing AI that can safely and successfully interact with the physical world.
- Stephen Gould** is an associate professor in the Research School of Computer Science in the College of Engineering and Computer Science at the Australian National University. He earned his PhD degree from Stanford University in 2010. His research interests include computer and robotic vision, machine learning, probabilistic graphical models, and optimization.
- Peng Zhang** is a Ph.D. candidate under the supervision of Professor Jochen Renz in the Research School of Computing Science at the Australian National University. He received a BE degree in intelligence science and technology from Beijing University of Posts and Telecommunications in 2011 and a MC degree in artificial intelligence from the Australian National University in 2013. His research interest is qualitative spatial representation and reasoning.