

# AIBIRDS: The Angry Birds Artificial Intelligence Competition

**Jochen Renz**

Research School of Computer Science  
The Australian National University  
jochen.renz@anu.edu.au

## Abstract

The Angry Birds AI Competition ([aibirds.org](http://aibirds.org)) has been held in conjunction with the AI 2012, IJCAI 2013 and ECAI 2014 conferences and will be held again at the IJCAI 2015 conference. The declared goal of the competition is to build an AI agent that can play Angry Birds as good or better than the best human players. In this paper we describe why this is a very difficult problem, why it is a challenge for AI, and why it is an important step towards building AI that can successfully interact with the real world. We also summarise some highlights of past competitions, describe which methods were successful, and give an outlook to proposed variants of the competition.

## Angry Birds as a Challenge for AI

Angry Birds is a very popular example of the *physics-based simulation game* (PBSG) category. The game world in these games is typically completely parameterized, i.e., all physics parameters such as mass, friction, density of objects, gravity, as well as all object types and their properties and location are known internally. Any chosen action can be perfectly simulated using an underlying physics simulator such as Box2D ([box2d.org](http://box2d.org)) which makes the execution and the consequences of actions look very real. In Angry Birds, the actions a player can perform are very simple: the player can decide (1) on the release point  $\langle x, y \rangle$  where a bird is released from the slingshot and (2) on a time point  $\langle t \rangle$  during flight of the bird when its optional special power is activated. Both the release point and the time point are in theory continuous, in practice very large, so that the resulting number of different actions is huge. A game level is solved if executing a selected sequence of actions  $\langle x, y, t \rangle$  leads to a game state that satisfies certain victory conditions. In Angry Birds all green pigs need to be destroyed in order to solve a given game level (see Figure 1).

This simple game play and the simple actions mean that even very young children are able to play the game successfully. The challenge of the AIBIRDS competition is to build an AI player that can play new game levels as good or better than the best human players. While this may sound simple, particularly in comparison to seemingly hard games



Figure 1: Angry Birds, Easter Eggs level 8 (©Rovio Entertainment). A good shot would hit the three round rocks on the top right which will trigger the destruction of the left half of the structure.

such as Chess, it is surprisingly difficult for a number of reasons. Assuming all parameters of the game world are known, we could simulate a number of actions and based on the outcome of the simulation, we could simulate a number of follow-up actions and so on until the victory condition is reached. If we select the actions we simulate in an intelligent way, this could lead to a successful solution strategy. However, the main issue in physics-based simulation games is that the outcome of an action is only known once we simulate it, which in turn requires us to know all parameters necessary for the simulation. This is very different from games such as Chess where the outcome of each action is known in advance. Accurately predicting or approximating the outcome of actions is one of the main challenge in this type of games, which is required before we can even look into determining good action sequences. Combining this with the potentially infinite action space and the possible lack of complete information about all required parameters means that we are in for a big challenge. Humans are very good at predicting consequences of physical actions (if I do this then this will happen, which in turn will trigger that, etc.), while this form of reasoning in unknown environments is still unsolved in its generality.

In the AIBIRDS competition, we play Angry Birds using the web version, publically available at [chrome.angrybirds.com](http://chrome.angrybirds.com). The competition server interfaces

with the website using a Chrome browser extension which allows us to take screenshots of the live game and to execute different actions using simulated mouse operations. Participating agents run on a client computer and can only interact with the server via a fixed communication protocol. This allows agents to request screenshots and to submit actions and other commands which the server then executes on the live game. Therefore, the only information participants obtain are sequences of screenshots of the live game. Hence, AI agents have exactly the same information available as human players. In particular, they don't know the exact location and other parameters of objects or the game world.

What is provided to participants is a computer vision module that detects known (=hard-coded) objects and gives an approximation of the objects boundary, their location and type. In addition, a trajectory planning module is provided which allows agents to specify which point they want to hit with a bird and if they want to shoot with a high or a low trajectory. This module then returns the approximate release point that hits the given target point. Since this depends on the scale of the game world, which can be different for every level, the trajectory planning module automatically adjusts trajectories in subsequent shots. In order to demonstrate the use of these modules, a sample agent called the Naive Agent is provided, which selects a random pig as the next target, and selects a random trajectory and tap point depending on the bird type.

During the competition, participants receive a number of unknown game levels they have to solve within a given time limit. Game levels can be played and re-played in any order. Agents are ranked according to the overall points they obtain in the given time (=sum of maximum points scored per solved game level) and after several rounds of elimination a winner is determined. During the Human vs Machine Challenge, we then test whether the best AI agents are already better than humans (=typically conference participants).

In order to achieve the goal of the competition, i.e., to develop an AI agent as good as or better than human players, we need to efficiently solve a number of problems in an environment that behaves according to the laws of physics:

1. detect and classify known and unknown objects
2. learn properties of (unknown) objects and the game world
3. predict the outcome of actions
4. select good actions in a given situation
5. plan a successful action sequence
6. plan the order in which game levels are played

These problems can be covered by different areas of AI such as Computer Vision, Machine Learning, Knowledge Representation and Reasoning, Planning, Heuristic Search, and Reasoning under Uncertainty, but due to the physical nature of the game world and the unknown outcome of actions, these are largely open problems. Progress and contributions in each of these areas will improve the performance of an agent, but in order to reach the goal, we need to jointly develop solutions to these problems across different AI areas. What makes research on PBSG such as Angry Birds so

important, is that the same problems need to be solved by AI systems that can successfully interact with the physical world. Humans have these capabilities and are using them constantly, AI is a long way away in this respect. This competition and other PBSG games offer a platform to develop these capabilities in a simplified and controlled environment.

## An overview of past competitions

So far 36 teams from 16 countries have participated in the competition and we have seen a multitude of different techniques. The best approaches so far used logic programming, qualitative reasoning, advanced simulation, structural analysis, analysis of predicted damage, and different machine learning techniques such as Bayesian ensemble learning. Interestingly, the winning agents in 2013 and in 2014 both used a number of different strategies, some of them very simple, and selected one of them depending on the game level and the current situation.

The performance of agents is improving significantly, as measured in the Human vs Machine challenge. In 2013, agents were clearly better than beginners, while in 2014 the best agents are already in the top third of human players. Interestingly, the two best teams in 2014 were both new teams, so newcomers have a good chance of doing well and are encouraged to participate. We also benchmarked all teams using the standard game levels. Team descriptions and other information are available at the [aibirds.org](http://aibirds.org) website.

## New variants of the competition

We have developed a version of the basic Angry Birds AI game playing software using Snap! ([snap.berkeley.edu](http://snap.berkeley.edu)), a simple visual programming language also used by [code.org](http://code.org). The goal of the AIBIRDS Snap! implementation ([aibirds.org/snap](http://aibirds.org/snap)) is to teach kids about AI and programming in a playful way. It allows kids to implement and test their own Angry Birds strategies while acquiring basic programming skills. We envisage a separate competition for school children based on the AIBIRDS Snap! implementation.

At the upcoming competition at IJCAI 2015, we plan to have a new competitive track where two agents try to solve game levels with alternating shots. At the beginning the two agents each submit a concealed offer as to how many points they are willing to pay for the right of the first shot. The agent with the higher offer starts and the agent with the winning shot gets all the points of the level. If the agent with the higher offer wins the level, the offered points are paid to the other agent. This competitive variant is easily possible with a small modification to the client/server communication protocol. The aim is to encourage agents to analyze game levels and to evaluate actions in advance rather than winning by a lucky shot. In addition it could make the competition interesting for game theory researchers.

## Acknowledgements

Thanks to my co-organizers XiaoYu (Gary) Ge, Peng Zhang, Stephen Gould, and the many students, supporters and participants who contributed to this competition.