Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)

Angry Birds as a Challenge for Artificial Intelligence

Jochen Renz XiaoYu Ge, Rohan Verma, Peng Zhang
Research School of Computer Science
The Australian National University
jochen.renz@anu.edu.au

Abstract

The Angry Birds AI Competition[1] has been held annually since 2012 in conjunction with some of the major AI conferences, most recently with IJCAI 2015. The goal of the competition is to build AI agents that can play new Angry Birds levels as good as or better than the best human players. Successful agents should be able to quickly analyze new levels and to predict physical consequences of possible actions in order to select actions that solve a given level with a high score. Agents have no access to the game internal physics, but only receive screenshots of the live game.

In this paper we describe why this problem 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, including a new competition track we introduced recently.

1 Introduction

Angry Birds is one of the most popular games of all times. It has a simple game play and one simple task: destroy all pigs of a given level by throwing different angry birds at them using a slingshot (see Figure 1). The pigs are protected by a structure composed of blocks of different materials with different physical properties such as mass, friction, or density. The actions a player can perform are spegified by the release coordinate $\langle x, y \rangle$ and the tap time t after release when the bird's optional special power is activated. Angry Birds is an example of the *physics-based simulation game* (PBSG) category. The game world in these games simulates Newtonian physics using a game internal physics engine such as Box2D (http://box2d.org) that knows all physics parameters and spatial configurations at all times, which makes the physics of the game play look very realistic.

In the Angry Birds AI Competition (AIBIRDS), the task is to build AI agents that can play new game levels as good as or better than human players. AI agents do not have access to the game internal parameters, they play a live game just like humans and get all their information about the current state of the game using computer vision.

This "human way" of playing makes it very hard for computers to play well, particularly compared to games like

[1]http://aibirds.org

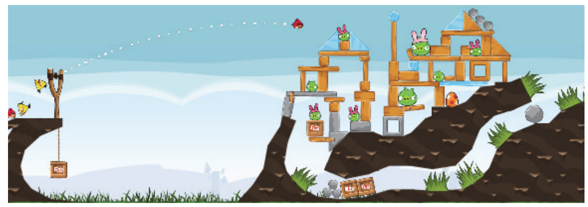

Figure 1: Angry Birds, Easter Eggs level 8 (©Rovio Entertainment). A good shot would hit the three round rocks on the top right which will trigger the destruction of the left half of the structure.

Chess that are difficult for humans, but easy for computers to play. The main difficulty of PBSGs like Angry Birds is related to the problem that the consequences of physical actions are not known in advance without simulating or executing them. This is partly due to the fact that the exact physics parameters are not available, which makes exact calculations impossible. But even if they were available, it would be necessary to simulate a potentially infinite number of possible actions, as every tiny change in release coordinate or tap time can have a different outcome. Knowing the outcome of an action is important for selecting a good action, and particularly for selecting a good sequence of actions.

The capabilities required for accurately estimating consequences of physical actions using only visual input or other forms of perception are essential for the future of AI. Any major AI system that will be deployed in the real world and that physically interacts with the world must be aware of the consequences of its physical actions and must select its actions based on potential consequences. This is necessary for guaranteeing that there won't be any unintended consequences of its actions, that nothing will get damaged and no one will get hurt. If this cannot be guaranteed, it is unlikely society will accept having these AI systems living among them, as they will be perceived as potentially dangerous and threatening. Angry Birds and other PBSGs provide a simplified and controlled environment for developing and testing these capabilities. It allows AI researchers to integrate methods from different fields of AI, such as Computer Vision, Machine Learning, Knowledge Representation and Reasoning, Heuristic Search, Reasoning under Uncertainty, and AI Planning that are required to achieve this.

2 Past competitions and the state of the art

The Angry Birds AI competition attracted interest from participants all over the world. Over 40 teams from 17 countries have participated so far and a multitude of AI approaches have been tried. One common characteristic of many top performing agents is that they try to solve the game via structural analysis. For example, the winner in both 2014 and 2015, DataLab Birds from CTU Prague used multiple strategies based on qualitative structural analysis. The agent analysed the structure by exploring spatial relations between the building blocks and selected the weak part (e.g. a block that supports the entire structure) as the target. Several other agents based their strategies on spatial reasoning and qualitative physics. The winning agent in 2013, Beau Rivage from EPFL aimed to exploit weak parts of the structure using multiple pre-specified strategies. The strategy selection was modelled as a Multi-Armed Bandit problem. The second best agent in 2015, AngryHex from U.Calabria and TU Vienna combined answer set programming (rule-based reasoning) and physics simulation. The agent encoded the structural information about the game scene into logic assertions while using physics simulation to determine the stability and reachability of an object. The second best agent in 2014, AngryBER from U.Ioannina, Greece adopted an ensemble inference mechanism based on an efficient tree structure. It constructs the tree structure of the game scene and calculates the expected reward of each reachable node according to a Bayesian ensemble regression.

Some competitive agents did not rely on structural analysis. For example, TeamWisc from U. Wisconsin who became third in 2013 used the weighted majority algorithm and a Naive Bayesian network to select the most appropriate shot. The agent was trained on a set of successful shots made by agents using different heuristics. IHSEV from ENIB, France determined shot sequences using parallel advanced simulations. Therefore, this agent is often able to find better shots and achieve higher scores than other greedy agents, but requires a larger memory and time consumption than others.

Starting in 2014, the best teams have made the source code of their agents available to allow other teams and newcomers to catch up more easily. Despite publishing their code (minus some secret parameter settings), the unmodified agent who dominated the 2014 competition, DataLab Birds, won again in 2015. But this time the other teams were much closer in performance and DataLab Birds were almost eliminated in the semi finals. It was extremely exciting to watch the different rounds live, particularly the grand final between Angry Hex and DataLab Birds where the lead changed with almost every solved level.

After each AI competition, we hold a Man vs Machine Challenge to test if AI agents are already better than humans. In previous competitions, humans always won with a wide, but shrinking margin. In 2013, half of human participants were better than the best AI, while in 2014 it was a third. The surprise team in 2015 was newcomer Tori from UTN Santa Fe in Argentina who ended up being almost as good as the best human players and was among the best eighth of all human players who participated.

3 A new competition track

It seems that AI agents are now almost as good as they can be without planning ahead multiple shots, i.e., as described above, most agents have a greedy strategy and aim for an optimal shot with the current bird, but do not plan sequences of shots. In order to encourage participants to design techniques that actually allow them to predict consequences of their shots well enough to be able to plan shot sequences and to analyse how a level can be solved with multiple shots, we designed a new compeitition track. In this new track, agents play the same set of game levels pairwise against all other agents, a different pairing in each round. Agents use alternating shots, i.e., agent A makes the first shot, agent B the second shot, agent A the third shot, and so on, until a level is solved or runs out of birds. Agents are bidding for the right to make the first or second shot. The agent with the higher bid gets to make the first shot, negative bids can be made to bid for the second shot. The agent who makes the winning shot gets all the points earned for solving the level. If the winning agent also got to make the shot it was bidding for (either first or second shot) then that agent has to pay the absolute bid amount to the other agent.

While this allows for an interesting game theoretic analysis and opponent modeling, the main idea of this track is to force agents to accurately predict the outcomes of their shots rather than shooting in a greedy manner. This track was trialed for the first time at AIBIRDS 2015. The clear winner was team IHSEV with a modified version of their main competition agent.

4 The future of the competition

Rovio Entertainment, the developer of Angry Birds has recently decided to discontinue the website with the freely available Chrome game on which our competition is based on. Now participants need the offline version of the game in order to develop and run their agents. Despite this inconvenience, we plan to continue our competition with the two existing tracks and also consider modifying these tracks or introducing new tracks to encourage the development of practically useful capabilities. One interesting modification would be to use game levels with unknown objects, where agents need to learn physical properties of these objects, but also need to be able to visually detect them in the first place. A third Angry Birds track could be one which requires agents to solve levels while damaging as little as possible. This version would be more closely related to the task of avoiding undesired consequences of physical actions, which is the essential capability of any AI that can successfully interact with the physical world. In addition to the Angry Birds game, we can also use other physics-based simulation games for our competition. This will be particularly useful once the desired capabilities have been successfully developed for the Angry Birds world. We will then move on to more realistic physics simulations until we can successfully interact with the real world. The benefit of using games is that there is an inherent scoring mechanism that allows us to easily compare the performance of agents and to benchmark them. Further details on the competition can be found at http://aibirds.org.