

Angry Concepts Team - Team Description Paper

Daniel Hennes Dario Izzo

Advanced Concepts Team

European Space Agency

European Space Research and Technology Centre

Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands

daniel.hennes@esa.int, dario.izzo@esa.int

Abstract

This paper describes a hybrid approach to design an intelligent computer agent able to play the game of Angry Birds (Rovio Trademark). Our approach combines ideas from intelligent search, evolutionary optimization and reinforcement learning. Preliminary results show that our agent outperforms the winner of last years competition.

1 Introduction

Angry Birds is an artillery game where the player takes control of a number of birds. The birds have lost their eggs to a group of pigs which now seek protection in structures made of wood, ice and stone blocks. Using a slingshot the player launches birds with the aim to destruct as much of the structure and hit the pigs. Angry Birds uses the rigid body Box2D physics engine to simulate the interaction of birds with the structure.

The Angry Birds AI Competition takes place for the second time at the 2013 IJCAI conference. The organizers provide a basic game playing software that makes use of the Chrome version of Angry Birds and includes a computer vision, trajectory planning and game interface component.

2 Approach

The Angry Concepts Team agent searches in the space of trajectories. For each object in the scene, i.e., pigs, wood, ice and stone blocks, a trajectory targeted at the center of the object is computed. Subsequently, all trajectories get scored and ranked accordingly.

3 Scoring

We define the fitness (score) of a trajectory based on the following fitness function. We trace along the trajectory and get a resulting sequence of objects. The fitness of the trajectory is calculated as the sum of weights of intercepting objects. A pig gets a nominal weight of 1. Other objects, i.e. stone, wood or ice blocks, get weighted according to their position in the trajectory. Namely, a) either objects are intercepted by the trajectory before the (first) pig in the trajectory, or b) objects are intercepted "after" the first pig, or c) the trajectory

does not intercept a single pig. Furthermore, the weights depend of the current bird as e.g. the yellow bird has a different effect on wood blocks then for instance the red bird. In total this results in 3 (wood, stone, ice) * 3 (before, after, no pig trajectory) * 5 (red, blue, yellow, white or black bird) = 45 possible weights.

The fitness function f for a target point p_t is computed as:

$$f(p_t) = \omega_1 + \omega_2 + \omega_3 + \omega_4 + \dots = \sum_{i=1}^n \omega_i \quad (1)$$

where w_i is the weight of the i th object in the sequence.

The sequence of objects is limited to $n = 6$. Scoring sequences beyond 6 objects showed to be deteriorating performance. This can be explained by the fact that a bird might get deflected - and possibly quite significantly so - by the first object in the sequence. If the sequence intercepts a indestructible of immobile structure (e.g. a mountain) scoring will be terminated before the collision with such structure.

For the 2013 competition, we used an ad hoc definition for weights. Future work will be on evolving weights using evolutionary strategies.

3.1 Ranking

We break ties if two trajectories get evaluated with the same score and have a similar launch angle ($\Delta\alpha < 1^\circ$). This ensures greater variety in the behavioural space and introduces novelty. Probabilistic selection of the "winning" trajectory is done at random from the top five of the pool (other possibilities include for instance roulette wheel selection).

3.2 Discounting

As it is only guaranteed that the bird will hit the first object of the sequence associated with the trajectory and the effect on objects following the first contact is unknown, we use a discount factor to evaluate objects $i > 1$. The fitness function (1) with discounting is:

$$f(p_t) = \gamma^0 \omega_1 + \gamma^1 \omega_2 + \gamma^2 \omega_3 + \dots = \sum_{i=1}^n \gamma^{i-1} \omega_i, \quad (2)$$

where $n \leq 6$.

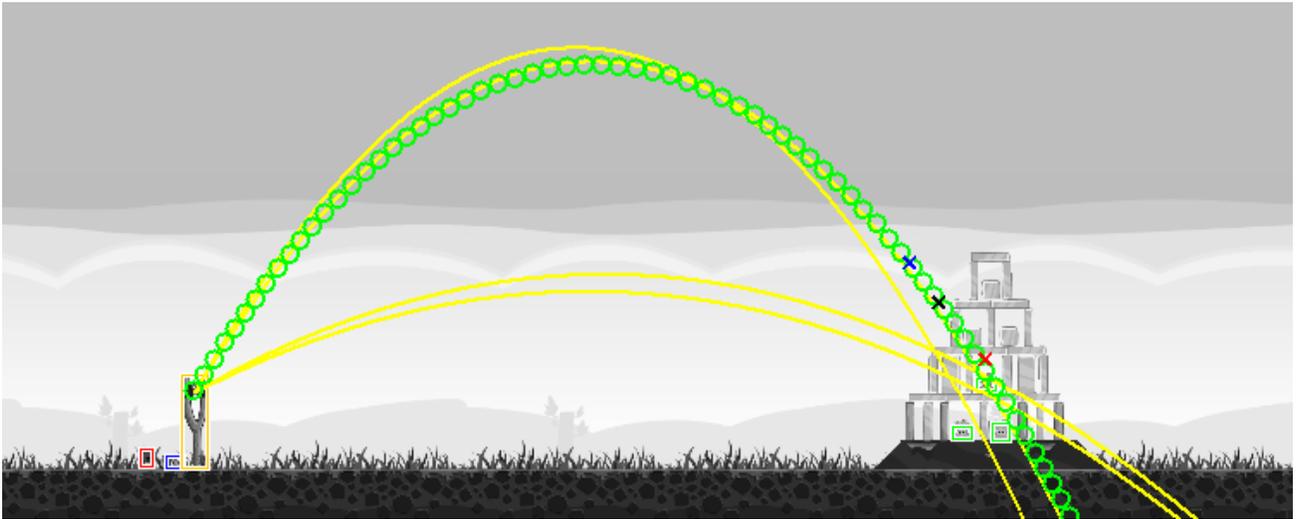


Figure 1: Top five trajectories. The selected trajectory is indicated by green circles. The target point is plotted by a red cross, the estimated collision point in black and the tap point in blue.

3.3 Tapping

For the blue, yellow and white bird we use “just-in-time” tapping. The bounding box of the first object in the sequence of a given trajectory is used to estimate the point of collision. Using the trajectory planner component of the game playing framework we can calculate the corresponding tap time and perform a tap right before the estimated collision point. For the red and black bird no tapping is performed.

4 Improvements of the game framework

We have improved the vision and game interface component of the basic game playing framework in several ways:

1. **Detecting of indestructible or immobile structures (e.g. mountains) is added to the framework:** Bounding boxes are not a fitting representation for these structures as they often have irregular, non-convex shapes. We use a pixel-by-pixel comparison to ensure that trajectories intercepting these structures get terminated at the point of collision.
2. **Detecting of the current bird in the slingshot:** Detecting the current bird in the slingshot is important for correct weighting in the scoring function. The framework did not allow to reliably detect the current bird.
3. **Won screen detection:** The framework allows for two methods of shooting: fast shooting without delay and safe shooting with a delay of 15s after the shot. We use a screenshot to screenshot comparison to detect the exact moment when the scene is stabilized and objects do not move any longer.

5 Preliminary results

We performed preliminary experiments to compare our approach with the Naive Agent provided by the organizers of the competition. The Naive Agent picks a pig at random and

shoots a bird directly at the pig. Our agent was able to outperform the Naive Agent (winner of the 2012 competition) by 111,155 points in the first 21 levels of Poached Eggs. Four runs with a time limit of 60min were performed. The average score was 970,865 with a minimum of 890,800 and a maximum score of 1,034,840. The 2012 Naive Agent scored 859,710 points in 1h12min.

6 Future work

As previously mentioned, future work will be concerned with optimizing the set of weights using evolutionary strategies. With in-game evaluation this optimization process will take a considerable amount of time (up to several hours per iteration). Furthermore, we will extend the trajectory planner to model special cases diverting from a parabola such as for post-tapping trajectories of the yellow, blue and white bird.

7 About the team

The Advanced Concepts Team (ACT) is part of the ESA’s Future Preparation Strategic Studies Office (PPC-PF). The team is essentially a channel for the study of technologies and ideas that are of strategic importance in the long term planning of the Agency. Placed at corporate level, it serves the function of a think tank providing decision makers the support of a highly multidisciplinary research group. Science and engineering research fellows (PhDs working at ESA for 2 years under the already existing fellowship scheme), Young Graduate Trainee and stagiaires form the bulk of this Team. Based at ESTEC, they carry out research work on advanced topics and emerging technologies and perform highly skilled analysis on a wide range of topics.