



Playing with AI

Konstantinos Sfikas

By 2017, AI had advanced far enough for AlphaGo, a specialised AI that can play the highly complex board game Go, to beat the major Go players in the world and be awarded professional 9-dan by the Chinese Weiqi Association. Go, however, is a fully deterministic game like Chess, with no random elements. Probabilistic games like *Pandemic*, on the other hand, are even trickier for AI to play efficiently, as the randomness of dice rolls or shuffled cards makes it much harder for computers to crack them. This problem inspired me (Konstantinos Sfikas) to attempt to create an AI that can play the *Pandemic* board game.

In the summer of 2018, I started working on this problem as part of my thesis for the M.Sc. in Digital Games (Institute of Digital Games, University of Malta), under the supervision of Dr Antonios Liapis.

At the core of our methodology lies Rolling Horizon Evolution (RHE), a planning algorithm that makes decisions by optimising action sequences through artificial evolution (introduced by University of Essex researchers in 2013). In order to make a single decision, RHE initially composes a population of random action sequences and evaluates them by simulating their potential result. Then an iterative process of optimisation takes place: the action sequences are randomly mutated, generating a set of offspring. The offspring will either replace their parents or be discarded based on a quality comparison. While this process repeats, the overall quality of the population tends to increase. After a predefined number of iterations, the agent simply selects the first action of the best-found sequence and applies it to the actual game.

Based on RHE, we designed the Policy-Based Rolling Horizon Evolution Agent (PB-RHEA), which operates on a higher level of abstraction, using a set of 'policies' (artificial behaviours) as an indirect encoding of action sequences. When composing or mutating sequences, PB-RHEA does not consider the full amount of potential single actions (as

RHE does), but rather selects among a much smaller set of possible behaviours that translate into specific actions and approximates their probable outcome through repeated randomised simulations. Through this technique, the agent's operation was greatly enhanced in terms of computational efficiency and overall performance.

During my thesis and the two publications that followed (both co-authored with my supervisor Dr Antonios Liapis), we performed a large number of computational experiments, analysing the agent's behaviour and optimising its performance. One of the most challenging aspects of our research was to design a set of heuristics that approximate the quality of any given game state, thus allowing the agent to evaluate the outcome of an action-sequence. Another challenge was to define the set of policies that the agent would use as building blocks in such a way that they are both efficient and expressive. Finally, fine-tuning the algorithm's parameters through trial and error was another critical aspect of the agent's degree of success. The results overall showcase that our proposed methodology exhibits a good performance against a hard problem and leaves clear avenues for further improvement.

From an academic perspective, the main contribution of our research is that it clearly expanded the knowledge on planning algorithms like RHE and, more precisely, their applicability on complex problems like *Pandemic*. Agents like the PB-RHEA can be used to play alongside human players in the digital versions of board games or even be used in the context of automated playtesting during the development phase of board games. Although gamers have been playing alongside AI for a long time, will game developers also adopt AI as a partner when designing their games? **T**

This research was carried out as part of an M.Sc. in Digital Games at the Institute of Digital Games, University of Malta, under the supervision of Dr Antonios Liapis.