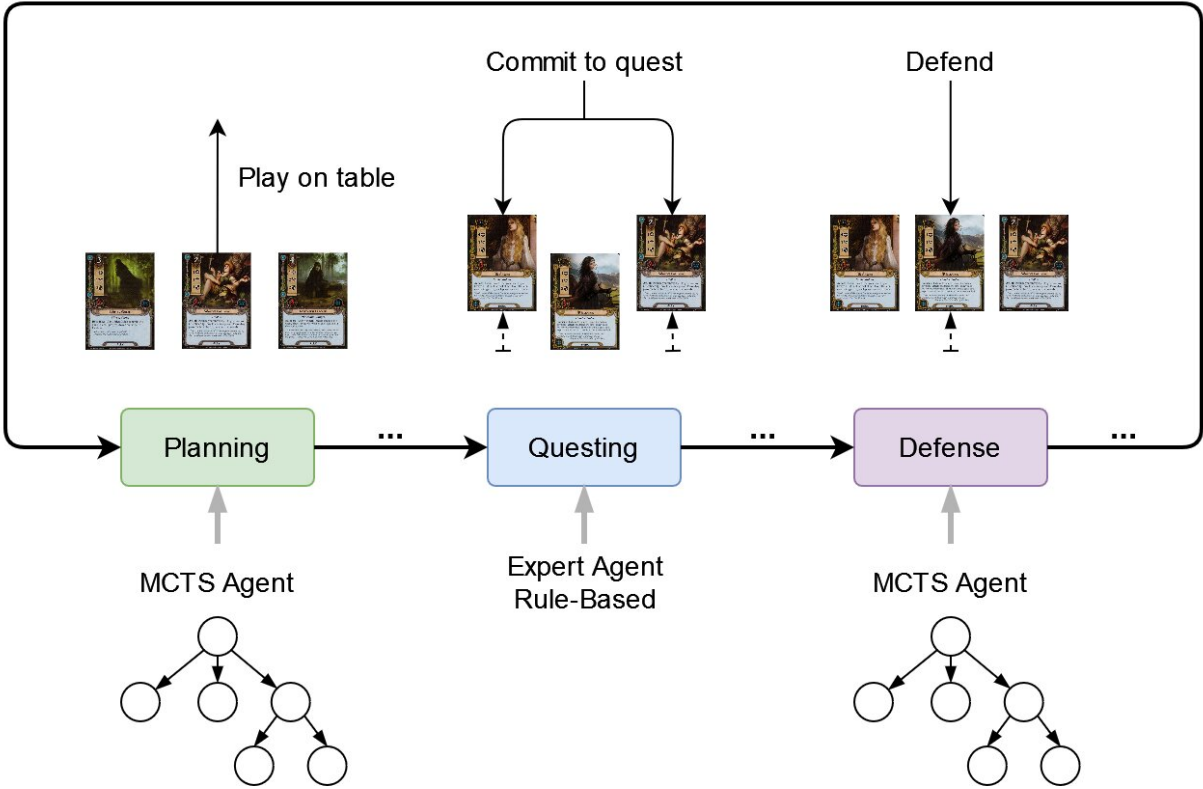# Monte Carlo tree search algorithms that can play the Lord of the Rings card game

October 8 2021, by Ingrid Fadelli

An illustration combining different agents in a single round of play at the card game Lord of the Rings. Credit: Godlewski & Sawicki.

When machine learning algorithms and other computational tools started becoming increasingly advanced, many computer scientists set out to test their capabilities by training them to compete against humans at

different games. One of the most well-known examples is AlphaGo, the computer program developed by DeepMind (a deep learning company later acquired by Google), which was trained to compete against humans at the complex and abstract strategy board game Go.

Over the past decade or so, developers have trained numerous other models to play against humans at strategy games, [board games](), computer games and card games. Some of these artificial agents have achieved remarkable results, beating established human champions and game experts.

Researchers at Warsaw University of Technology have recently set out to develop a technique based on Monte Carlo tree search (MCTS) algorithms that could play the [Lord of the Rings (LotR) classic card game](), released in 2011 by Fantasy Flight Games. An MCTS algorithm is a universal heuristic decision method that can optimize the searching solution space in a given game or scenario, by playing a series of random games, known as 'playouts'. The researchers presented their MCTS technique in a recent paper pre-published on arXiv.

"We are fans of the card game LotR, but we found that there were no existing AI approaches that could play this game," Bartosz Sawicki and Konrad Godlewski, the two researchers who carried out the study, told TechXplore. "Nonetheless, we found applications of tree search methods for similar card games such as Magic: The Gathering or Hearthstone."

The main reason why a computational method that can play the LotR card game did not yet exist is that developing such a method is highly challenging. In fact, LotR is a cooperative card game characterized by a huge space of possible solutions, a complex logical structure and the possibility of random events occurring. These qualities make the game's rules and strategies very difficult to acquire by computational methods.

"The 2016 Go tournament was the last moment when human players had a chance to compete with AI players," Sawicki and Godlewski explained. "The objective of our paper was to implement an MCTS agent for the LotR game."

The LotR card game is difficult to compare to other fantasy and adventure card games, such as Magic the Gathering, Gwent or Hearthstone. In fact, in contrast with these other games, LotR is designed to be played alone or as a cooperative team, rather than in competition with other players. In addition, the decision-making processes in the game are highly complex, as the gameplay includes several stages, most of which depend on the outcome of the previous stage.

Despite these challenges, Sawicki and Godlewiski were able to develop an MCTS-based method that could play LotR. They then evaluated the technique they developed in a series of tests, carried out on a game simulator.

"Our MCTS agent achieved a significantly higher win rate than a rule-based expert player," Sawicki and Godlewski said. "Moreover, by adding domain knowledge to the expansion policy and MCTS playouts, we were able to further improve the model's overall efficiency."

The recent work by Sawicki and Godlewski proves that it is possible to successfully combine different AI and computational techniques to create artificial agents that can play complex and cooperative multi-stage games, such as the LotR card game. Nonetheless, the team found that using MCTS to tackle these complex games can also have significant limitations.

"The main problem is that MCTS merges game logic with the AI algorithm, so you have to know the legal moves when you are building a

game tree," Sawicki and Godlewski explained. "Yet debugging for game trees with significant branching factor is a nightmare. There were many cases in which the program ran smoothly, but the win rate was zero, and we had to examine the whole tree manually."

In the future, the MCTS-based technique developed by this team of researchers could be used by LotR enthusiasts to play the game in collaboration with an AI. In addition, this recent study could inspire the development of other AI tools that can play complex, strategic and multi-stage games. In their current and future studies, Sawicki and Godlewski would like to also explore the potential and performance of deep reinforcement learning (RL) agents trained on the LotR game.

"Our current work focuses on using RL methods to further improve the performance of AI agents in the game," Sawicki and Godlewski added. "In this case, given a game state, the neural network returns an action executed by the environment (i.e., the game simulator). This is tricky, because the number of actions varies in different states and policy networks can only have a fixed number of outputs. So far, our results are promising, and we will explain how we achieved these results in an upcoming article."

More information: Konrad Godlewski, Bartosz Sawicki, Optimisation of MCTS Player for The Lord of the Rings: The Card Game. arXiv:2109.12001v1 [cs.LG], arxiv.org/abs/2109.12001

Konrad Godlewski, Bartosz Sawicki, MCTS based agents for multistage single-player card game. arXiv:2109.12112v1 [cs.AI], arxiv.org/abs/2109.12112