Case-based Decisions in a cooperative Multi-Agent Gaming Scenario

Daniel Schreider[1], Pascal Reuss[1]

[1]Institute of Computer Science, Intelligent Information Systems Lab, University of Hildesheim, Hildesheim, Germany

Abstract
Using artificial intelligence (AI) in competitive games has been widely researched and the AI has proven to be able to win in many games. However, in cooperative games, where the players have to win not against each other but the game itself, has additional challenges to met. This paper describes a digital realization of a complex, cooperative board game called *The Legend of Andor.* We give an overview over the content and rules of the board game and how the digital version is realized. We set a focus on the knowledge modeling of the case-based software agents and describe the case structure and similarity measures. We also present our experimental evaluation and the results as well as out interpretation.

Keywords
Case-Based Reasoning, Software Agents, Cooperative Games, Knowledge Modeling

1. Introduction

Using artificial intelligence (AI) in computer games in not new and has been research and applied in the last decades on a wide scale. Usually, there are three use cases of applying AI in a gaming scenario: the AI is acting as the player, the AI is used to enhance the behaviour of non-player characters, or the AI is used to generate gaming content like maps or items. Of course, these three use cases can be also combined [1][2].

In our research, we focus on the AI as the player of a cooperative gaming scenario. These scenarios require more interaction between the players of a game than competitive scenarios, because not the other players are the enemies, but the game itself. It is required to work together to beat the game. Our point of view is that winning cooperative games is usually not possible within the first few tries, therefore experience is required to be able to beat the game. In addition, the same base scenario can differ in various aspects during individual playthroughs. [3] Therefore, only relying on given rules will not lead to winning the scenarios, but using Case-based Reasoning (CBR) seems a promising approach to capture the experience of human players in cases and let an AI use this experience to play a cooperative game. From our perspective, a combination of applying rules from the given rule book and experience provided by the CBR component will increase the chances of winning cooperative scenarios.

While CBR was already applied to many gaming scenarios it was only once applied to cooperative games in a soccer scenario in a RoboCup tournament in 2008 [4]. From our perspective, the possibilities of applying CBR for cooperative gaming scenarios and incorporation experience into the decision making process between multiple players is very interesting for further research. There are three hypotheses we are investigating in the presented work. The first hypothesis states that CBR is a viable approach for providing experiences from human players to software agents playing a cooperative board game. The second hypothesis says that the provided experience enables the software agents to win the given scenarios. The third hypothesis is that by capturing human player experience in cases, these training cases represent the results of the decision making process of the individual players and is sufficient for the case-based agents to win the game without additional knowledge.

In the next section we give an overview of some related work in the research field and describe the digital realization of the game in Section 3 with the limitation compared to the original board game.

LWDA'24: Lernen, Wissen, Daten, Analysen September 23–25, 2024, Würzburg, Germany
schreide@uni-hildesheim.de (D. Schreider); reusspa@uni-hildesheim.de (P. Reuss)

In Section 3.3 the knowledge modeling of the CBR system is presented in detail. Then we describe the experimental setup and the results of our evaluation and finish the paper with a summary and an outlook on future work.

2. Related Work

During the last decades, many digital variants of board games with artificial intelligence have been created. The research was mainly focused on traditional competitive games and CBR is also used in many competitive games for decision making, for example in poker [5],[6], real-time strategy [7],[8],[9], card games [10], and first-person scenarios [11],[12].

While there are numerous publications in the field of AI development on decision making for conventional board games, there is little activity in this regard in cooperative games. Gaina and colleagues developed a framework in which they used digital versions of various board games, including the cooperative board game Pandemic. Each game was played by four artificial intelligence with different approaches. However, no agent or, in the case of a pandemic, a group of four instances of a used technology was able to win the game because there was no long-term strategic planning. [13]

Sfikas and Liapis developed an AI that controls all decisions made by cooperating players in a simplified version of the game Pandemic by weighing the odds of winning against risks through stochastic factors. It uses an algorithm that belongs to the category of evolutionary algorithms. In short, such algorithms are based on the theory of evolution and implement the idea that only those individuals in a population reproduce who meet a certain selection criterion. Over time, the population converges into a set of individuals best suited to accomplish a desired task. The so-called Rolling Horizon Evolutionary algorithm uses the possible actions of a player to create an individual, which then represents a sequence of actions and thus the behavior of the player. Actions not allowed at a certain time are treated as doing nothing or fitting. The simulated end state of an individual is evaluated using a heuristic fitness function and results in the fitness of the individual. In experiments, the authors test different fitness functions, distinguishing between optimistic (how close players are to victory) and pessimistic (how close players are to defeat). They come to the realization that a combination of optimistic and pessimistic fitness features leads to better results than using only one of the limbs. [14],[15]

Instead of developing individual AI for each player that have to interact with each other, the implementation in this work was done via a single AI that develops a master plan for the game and assigns a plan to each individual player. Thus, the challenge of AI also changes from the individual strategy development and goal alignment of a single player to the holistic consideration and interaction of short-term game moves and long-term strategy considerations. This also reflects the human risk assessment and trade-off between offensive approach to victory and prophylactic measures to avoid sudden defeat due to unpredictable events. This is reinforced by the fact that these types of cooperative games include a single opportunity for victory, but several for defeat.

The challenges are very different for agents, who only have to control a single player and collapse with a human player when they participate in the game. First, it is necessary to anticipate the most likely follow-up actions of the other players, for which the human player is optimally depicted as a model. For this purpose, the authors propose modelling using past actions [16] or procedural personas [17]. Furthermore, the communication of the planned strategy of an AI to humans and vice versa is necessary. Subsequently, a discussion must take place, in which one of the two parties is convinced of the desired course of action of the other and thus can be settled on a strategy. [14]

Our approach differs from other existing approaches in two important ways. On the one hand, we use CBR as a white-box method to map the experiences of human players and make those experiences and the decisions based on them transparent. On the other hand, we do not use any overarching strategic planning or discussions between players, but instead focus on the decisions of individual players. As a result, we use a bottom-up approach where the individual decisions should lead to a common strategy.

3. Cooperative Multi-Agent Gaming Scenario

This section describes the board game that was realized as a gaming scenario. We describe the basic idea and rules of the board game as well as the part that was implemented in our prototype and the integration of the CBR component. The complete board game is very complex an contains several *legends*. We set our focus during the prototypical implementation on the important concepts and rules and the first of five *legends*.

The focus by playing a cooperative board game is set to the group and a sense of community instead to the victory of a single player. Often players loose in cooperative games, because those games only offer one or few ways to victory, but many reasons for defeat. Despite of this ratio between the chances to win or loose, this is the fact that makes cooperative board games interesting and keeps the players motivated. Defeats, especially those that come close to achieving victory, leave players analyzing the game and looking for details of their failure. They also discuss better strategies and procedures after the end of the game and try to identify their weaknesses and mistakes. One of the hallmarks of cooperative games is the dialogue that people have during the game and afterwards. While in competitive board games players must keep their thoughts on strategy and tactics to themselves and communication is limited to the general course of the game, in cooperative board games a collaborative discussion is essential for joint decision-making. The success of cooperative games is based on the intense interaction between the players. The great advantage of a cooperative game in this regard is that the efforts of all players have a positive impact on the overall gaming experience. The challenge lies with the game developer to provide players with a balanced overall experience in terms of difficulty. [3]

3.1. The Legend of Andor Board Game

The "Legends of Andor" is a cooperative, strategic, and turn-based board game in which players act cooperatively and lose or win together. It was designed by Michael Menzel and published in 2012 [18]. The goal of the game is to play a so-called *legend*. A legend is a game scenario with given starting points for the players and victory conditions. In addition, it give a set of scripted events that occur during the legend. Tthe goals are to protect the castle of the players from enemies and perform various additional tasks, like delivering items. As the game progresses, the fictional story of the legend is told, gradually adding more elements to the game while reducing the amount of time available for players to complete the legend. Players play on fictional days, where they can use a limited number of hours per day for actions such as running or fighting. Players take turns performing a main action, which is supplemented by further action options without losing time. Players can interact with each other by exchanging gold or objects. On a legend bar, a narrator character represents the current status of the legend by letters beginning with A to N. If the narrator reaches the letter N and not all tasks have been completed by then or if too many enemies have entered the castle during the course of the game, the legend and thus the game is lost.

This effect is enhanced by the fact that each defeated creature also allows the narrator to continue running. Each player moves a hero on the field, with each hero having a number of so-called will and strength points, as well as different characteristics and unique special abilities. In addition, players can use a dice system and heroes' strength to fight the enemies as they approach the castle over time. Similar to the heroes, there are different types of enemies with different numbers of will and strength points, as well as different rewards upon victory over each enemy. Players can cooperate by fighting an enemy together with united strength. Both, the movement of the heroes across the fields and the fighting with enemies cause cost in terms of hours on the daily bar. At the end of a day, the sunrise takes place, in which, among other things, an event card is opened, all enemies on the game field are moved towards the castle, and the narrator goes one letter further, so that the end of the legend approaches. Defeated enemies provide players a reward, but also move the narrator, forcing players to compromise in order to succeed.

Furthermore, various profitable items can be collected on the fields or purchased from a dealer in exchange for gold. In addition, the strength and will points can be increased there. The number of will points determines how many dice players are allowed to use in battle, with higher will points meaning more dice. There are various hidden fog patches on the playing field, whose activation triggers a certain positive or negative action such as raising will points or spawning an enemy. Another feature of the game is that each playthrough of the same legend is variably constructed by some random factors. The event cards, which are presented in random order during each run, also contribute to this. In this way, enemies can start from different positions, objects and enemies can be integrated into the game at different times, or end bosses have a diced strength.

In the first legend of the game, the focus is on the defense of the castle and another task, the transfer of a parchment to a specific field. Items only play a minor role here and players get to know the basic functions and actions. Throughout the course of the legend, the creatures always appear in predetermined fields, but sometimes the time at which they are integrated into the game varies. Since players only have one strength at the beginning of the game and the weakest creature has two, it is tactically wise to fight the creatures together rather than alone at this stage of the game. When the narrator reaches the letter D, two creatures are placed near the castle, which would be placed on the same field if the creatures moved. However, since each field can only be occupied by one creature at any given time, one of the creatures will already run into the castle. If this event occurs during a period where players have consumed many hours and will soon have to end their days, players may no longer be able to fight the creatures and thus prevent their movement. With four players, the rules dictate that a single creature can run into the castle, so the game is not over yet. Players face a challenge early on. Upon reaching the letter H, players are given the opportunity to win the game by transporting the now placed parchment to a given field. By setting up more creatures at the same time and breaking up a crucial connection between two fields, players must fight their way to this field, as the parchment must not be on a field with a creature at any time. The rules of the game can be read in full in the official game guide. [19]

3.2. Digital Implementation

For the implementation of our digital game, adaptations to the game flow were made to reduce the complexity and effort of the prototypical implementation. One adjustment is the omission of the introduction phase, which only serves to help players learn the basics of the game. This means that our digital version starts directly at the start of the first legend. We focus on the main part of the game and also get the possibility that the AI can play the first legend several times in a row. In the board game, the introduction is meant to be played every time before the first legend is started. A second adaption was made regarding the turn-based actions of the players. As mentioned in the subsection before, players have a main action per turn and several supplement actions. While the main action has to be performed during the specific players turn, the supplement actions can be performed outside their own turn. For our current prototype, all actions can only be performed during the players own turns to reduce the decision making complexity of the AI. The last adaptation to the usual game loop regards the groups decision making process. In the board game the players would discuss the situation and find a mutual decision for the given problem. In our prototype, group decisions are represented as experience from the past in the CBR system rather than having an active discussion between the agents. In future versions of the game it is planned to realize the mentioned game mechanics in their original intend and add more content.

The game was implemented using the Unity 3D framework [20] for the game itself with four software agents and the open source tool myCBR [21][22] for the CBR component of the agents. Figure 1 shows the game board which was transferred into the digital version. The graphical representation of the board acts mainly as background graphic in the game. The game elements are represented using a graph structure with individual spaces as vertices and the connections as edges. Because the players can move freely between the spaces along the connections, we use an undirected graph. Overall, there are 84 vertices in the graph.

In the upper left corner, player information is displayed in four colored rectangles. This includes the interaction opportunities available to players. In the first line are the main actions **Running**, **Fighting**

Figure 1: The game board

and **Joint Fighting**, each of which consume at least an hour on the daily bar. Next to the last main action **Skip** are the buttons to end the day and roll the dice. In the lower half, the free options are positioned. These are depositing and picking up objects and gold, emptying a well and acquiring strength points. Each button is active and can only be pressed if the respective game situation allows it, otherwise it is greyed out. The top of the page shows which day of the game the players are in.

The main actions for the players are implemented like in the board game, while the additional actions are realized with the aforementioned restrictions. The movement of the enemies and the combat system are corresponding to the board game. Legend cards and event cards are also implemented, but the legend cards only for the first legend.

3.3. CBR Integration

While the game can also be played by human players, the focus was set on the realization of AI-supported game play with the help of CBR. The goal was to enable a fully automated, cooperative gameplay of the first legend for the four software agents. Each agent has access to the CBR system and performs a retrieval to get suggestions on how to act in their specific individual solution.

The knowledge model of the CBR system was build to provide the AI with all required information about the current game state. After an analysis of all available knowledge in our digital version, three different situations were identified: the basic situation, the fight situation, and the reward situation. The basic situation represents all situations that can occur outside a fight, while the reward situation takes place directly after a fight and can be seen as the conclusion of a fight situation.

The basic situation contain fifteen problem attributes that describe individual information about the players properties as well as information about the environment. In addition, one solution attribute was modeled to represent the plan of suggested actions. These attributes and their explanations can be found in Table 1. The attributes one to seven contain information about the players individual state, while the other attributes contain information about the environment, e.g. the state of the game board.

The fight situation occurs, when a player has initiated a battle with an enemy, for example based on the suggested solution of a basic situation retrieval. While the battle itself is based on the dice roll and therefore not experience-based, the question whether a battle should be continued after the first round or not, can be answered by experience. The second case structure contains eleven attributes. The attributes *dice*, *hour*, *strength*, and *willpower* are the same a in the basic situation. Additional attributes represent the current *battle round*, *the players position*, the *strength* and *willpower* of the enemy and the *willpower* lost by the player and the enemy in the current round. The solution is stored in the *plan*

Table 1
The attributes for the basic situation

Attribute	Data Type	Explanation
(1) canFight	boolean	is an enemy in reach for combat?
(2) canFightTogether	boolean	is an enemy in reach for more than player?
(3) dice	integer	the player's available number of dices
(4) gold	integer	the player's amount of gold
(5) strength	integer	the player's strength points
(6) willpower	integer	the player's willpower points
(7) hour	integer	the player's used hours on the current day
(8) fogtokens	integer	the number of inactive fog tokens
(9) letter	integer	the current letter on the legend bar
(10) shields	integer	the number of available shields
(11) parchmentIsActive	boolean	is the parchment delivery goal active?
(12-15) position{color}	symbolic	position of the {blue, green, purple, or yellow} player
(16) plan	String	the suggested actions

attribute.

The reward situation occurs after a successful fight, when the players have to decide which kind of reward they want to choose. Based on the enemies level, the amount of reward points can range from two to four points. These points can distributed between gold or willpower. The case structure for this situation contains the attributes *gold*, *letter*, *parchmentActive*, *strength*, *willpower* and as an additional attribute the *amount* of reward points. The solution can again be found in the *plan* attribute. The additional attributes in the fight and reward situation have integer data types.

The boolean attributes have a binary similarity measure, while the integer attributes have a linear similarity measure based on the distance of the values. The symbolic attributes for the players position requires a special similarity measure. The spaces on the game board are not in a specific order and only linear connect, but one space can have several connections to other spaces. Therefore, the movement of the players is not linear, but they can freely move around the board along the connections. Figure 2 shows the spaces as a graph.

Two space positions are similar if they are close to each other on the game board or if they are direct neighbors. In this game, the spaces with adjacent numbers are not necessarily also adjacent fields. Since our tool myCBR cannot work directly with graphs, we needed a different approach to calculate similarity. Therefore, we populate a similarity matrix based on the number of edges between two nodes. This distance is then converted into a similarity value. An edge between two nodes corresponds to a similarity of 0.9 while the maximum distance of 7 edges corresponds to a similarity of 0. For the global similarity of the cases, a weighted sum of the local similarities is used. The weights of the individual attributes can be found in Table 2. These weights were determined by three experienced players of the game based on their decision making process during the game and adapted during experimental retrievals until the returned cases matched the experts expectations.

After modeling the necessary concepts, concepts and case structures, the initial case bases were created. Overall, five case bases were created, one for each player and one that contains all recorded cases. The cases were recorded by observing human game plays of our prototype. At the beginning of each players turn, a case is created with the current situation as the problem description. At the end of the turn, the performed actions are stored as the solution. This way, we are not using only a few constructed cases, but the experience from real situation during the game play. In addition,

Table 2
Global similarity weights.

Attribute	Weight	Attribute	Weight
canFight	1.25	canFightTogehter	1.25
dice	1.0	gold	0.75
strength	1.0	willpower	1.0
hour	1.5	fogtokens	0.5
letter	1.25	shields	1.0
parchmentActive	1.5	position{color}	1.0

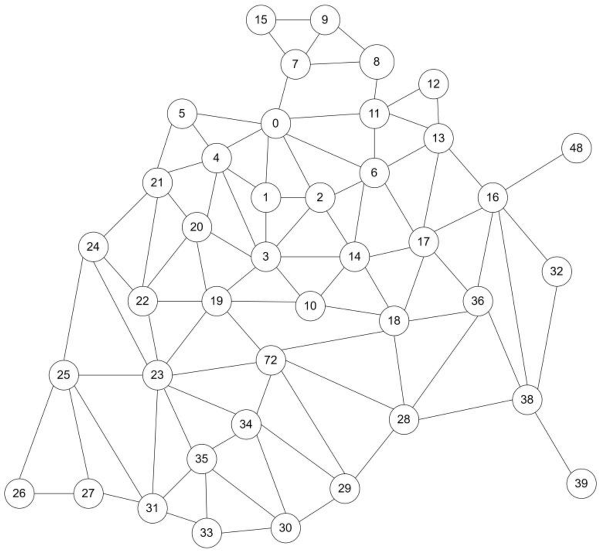

Figure 2: The game board as a graph

new cases from human game play can also be recorded while testing the AI with the already existing ones. Another major thought behind this idea is to capture the decision making of the human game play into the case. The human players interact and argue with each other during the training games. The arguments itself are not represented in the cases, but the result of the decision making process, represented as actions in the plan. The game and the CBR system can be found via github: https://github.com/PascalReuss/DieLegendenVonAndor

4. Evaluation - Setup and Results

This section describes our experimental evaluation of the prototypical game. For our experimental setup, three different training scenarios were defined based on the number of training games. In the first scenario, three training games with human players were played and a total amount of 193 cases were recorded. The second scenario contains six human-played games with a total amount of 393 cases recorded. In the last scenario, eight training games were played and the CBR system recorded 503 cases for all players. Table 3 shows the use cases with their training games and the cases distributed over the three types of cases. The cases are not equally distributed over all four players, because if a player decides to take no action in his turn, no case was recorded. I shall be mentioned, that the human-players were experience players and all eight training games have been won by them.

In our experiments, we let the AI play 7 games for each training scenario. For each game we observed the defeated enemies, the day and the letter that were reached, the status of the parchment and whether

Table 3
Recorded cases based on the training scenarios

Scenario	Total Cases	Basic Cases	Fight Cases	Reward Cases
Three training games	193	163	3	27
Six training games	393	434	15	56
Eight training games	503	424	7	72

victory was achieved or not. The results of this experiment can be found in Table 4. For the enemies, we also take into account, whether the AI defeated the standard enemies or the stronger ones, the so-called Skral. For the parchment, we distinguish between three states: a - means that the parchment was not in the game at all, because the corresponding progress was not achieved. a x means that the parchment appeared in the game, but was not taken by the players. *Collected* means the the parchment was taken by one of the players, but not transported towards the castle. *Transported* means that the parchment was taken and transported towards the castle.

The results show that AI has won only one of the 21 games in total. The main reason from our perspective, is that the AI failed to pick up the parchment in the correct time frame. This means that while the AI was able to find the plan to pick up the parchment in the retrieval, the parchment was not in play at that point. In the next step, the AI moved towards the target field, even though the parchment is not in the possession of a player. As a result, the AI lost the game shortly afterwards. In only three games has the AI managed to pick up the parchment and carry at least one field. In many games, the AI was unable to defend the castle and lost the game early before the parchment even came into play.

We took a closer look into the decisions of the AI to find the factors that contributed to the results. Therefore, we considered the following factors: the improvement of the heroes strength and willpower, the positioning of the heroes to interact with each other, the defense of the castle, and the interaction of the two winning factors in the first legend.

The AI increases its stats both by using water fountains and by purchasing strength points. Additionally, the AI increases stats by defeating creatures. An interesting point is that, contrary to how humans play, the AI often works overtime and loses willpower points as a result. With regards to interaction, the AI often used group fighting to defeat creatures. Although it was rare for the AI to pick up previously placed gold, it was much more common for the left-behind gold not to be picked up and left there for the rest of the game. This gold was therefore lost for the players. Regarding the defense of the castle, the AI have been observed to often ignore creatures near the castle. In rare cases, the AI would attempt to fight such a creature, lose the fight, and then move away from it. In the final stages of the game, we could observe that it was difficult for the AI to maintain both victory conditions. Most of the time, either the parchment was picked up by a player, but the path to field 57 was not cleared, or the players were able to clear the path in question, but no player picked up the parchment and took it with them. This interaction of the players only worked once.

Based on the evaluation results, the three hypotheses stated in the beginning could be answered. The first Hypothesis that CBR could be used to capture human player experience for the use by AI players has proven correct with some restrictions. The experience captured in the cases has enabled the software agents to play the game and win at least one time. The agents could deal successfully with various situations during the game play like fighting enemies or empowering the heroes. Especially defeating a Skral is not an easy task and the agents managed to summon and win a fight at least five times. This shows that the captured experience is successfully used to solve individual situations during a legend on their own or as a team, but the superior goals of the legend were often not considered. The second hypotheses is therefore proven partially correct. The case-based agents was able to win one game. While this shows it is possible to win, the majority of the games were lost. This hypothesis needs further research and improvement of the CBR component to enable the software agents to win at least half of the games played. The third hypothesis has also proven to be partially correct. The

Scenario	Game	Defeated Enemies	Reached Day	Reached Letter	Parchment	Victory
Three training games	1	4	3	G	-	No
Three training games	2	4	3	G	-	No
Three training games	3	5	3	H	x	No
Three training games	4	3	3	F	-	No
Three training games	5	5(1 Skral)	3	H	collected	No
Three training games	6	6	4	L	collected	No
Three training games	7	4	3	H	x	No
Six training games	1	8(1 Skral)	3	K	x	No
Six training games	2	2	3	E	-	No
Six training games	3	5	2	G	-	No
Six training games	4	0	3	C	-	No
Six training games	5	4	3	G	-	No
Six training games	6	4	3	G	-	No
Six training games	7	10(1 Skral)	4	K	transported	Yes
Eight training games	1	2	3	E	-	No
Eight training games	2	5	3	H	transported	No
Eight training games	3	8(1 Skral)	4	L	x	No
Eight training games	4	4	3	G	-	No
Eight training games	5	5	3	H	collected	No
Eight training games	6	5	4	K	collected	No
Eight training games	7	7 (1 Skral)	3	K	transported	No

Table 4
Results of the evaluation

idea of waiving overall strategic knowledge that has an impact on the decision of the AI has not been fully demonstrated to be a success, but was also no failure. The agents was able to win one time only using the captured experience, but the use of additional knowledge could have an impact on playing to game more focused on the superior goals of a legend. The results and answers to the hypotheses show from our perspective that either more case are required, that capture a wider range of situations or an adaptation process that can adapted the suggested plan. In addition, the case structure should be extended to consider the impact of a case to the superior goals. Considering this impact in the similarity measures could then lead to a more focused game play with respect to the superior goals.

5. Conclusion and Outlook

In this paper we describe a cooperative gaming scenario based on the board game *The Legend of Andor* and the realization of a case-based AI to play the game. We conducted an experimental evaluation to analyze the game-playing behaviour of the AI to answer the question, if a case-based AI can play and win a complex cooperative game. While the AI loses most time, it was able to win one time. Interpreting this result, the fact that many games played by human players are also lost, should be considered. The results show, that there is a lot of room for improvement, especially for the cooperative situation in the game and working towards the superior goals.

There are several limitations of our approach. While capturing the human experience gives an adequate initial case base, there is no learning approach for learning new cases during game play or adapt the existing solutions to new situations. Also the number of captured cases is limited to the situations that occurred during the game session of the human players. Another limitation is no quality information about the individual cases is stored. While the humans won their games and therefore their stored decisions could be seen beneficial for the game play, no information about the impact of individual situations and decisions on the overall strategy is considered.

There is planned a lot of future work. At first, the case-based AI shall be improved regarding the knowledge model, the retrieval and the interaction between the agents. An approach for learning additional cases during game play or rate used cased based on a reward function will also be considered. In addition, the game will be connected to our visualization tool VISAB [23], to get more insights and explanations on the agents behaviour. The evaluation will also be extended to compare different knowledge models and similarity measures.

References

[1] I. Millington, J. Funge, Artificial Intelligence for Games, Third Edition, 2nd ed., Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2019.

[2] G. N. Yannakakis, J. Togelius, Artificial Intelligence and Games, Springer, 2018. http://gameaibook.org.

[3] D. Maynard, J. Herron, The allure of struggle and failure in cooperative board games, Analog Game Studies 3 (2016).

[4] J. Y. Kuo, H. Z. Lin, Cooperative robocup agents using genetic case-based reasoning, in: 2008 IEEE International Conference on Systems, Man and Cybernetics, IEEE, 2008, pp. 613–618.

[5] A. Sandven, B. Tessem, A case-based learner for poker, Publications of the Finnish Artificial Intelligence Society (2006).

[6] I. Watson, J. Rubin, Casper: A case-based poker-bot, volume 5360, 2008, pp. 594–600. doi:10.1007/978-3-540-89378-3_60.

[7] S. Ontanon, Case acquisition strategies for case-based reasoning in real-time strategy games, Proceedings of the 25th International Florida Artificial Intelligence Research Society Conference, FLAIRS-25 (2012) 335–340.

[8] B. Weber, M. Mateas, Case-based reasoning for build order in real-time strategy games, Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment 5 (2009) 106–111. URL: https://ojs.aaai.org/index.php/AIIDE/article/view/12360. doi:10.1609/aiide.v5i1.12360.

[9] J. M. Schönborn, K.-D. Althoff, Towards case-based reasoning in real-time strategy environments with seasalt, in: D. Trabold, P. Welke, N. Piatkowski (Eds.), Lernen, Wissen, Daten, Analysen. GI-Workshop-Tage "Lernen, Wissen, Daten, Analysen" (LWDA-2020), September 9-11, Online, CEUR, 2020, pp. 154–161.

[10] M. Woolford, I. Watson, Scout: A case-based reasoning agent for playing race for the galaxy, in: D. W. Aha, J. Lieber (Eds.), Case-Based Reasoning Research and Development, Springer International Publishing, Cham, 2017, pp. 390–402.

[11] P. Reuss, J. Hillmann, S. Viefhaus, K.-D. Althoff, Case-based action planning in a first person scenario game, in: R. Gemulla, S. Ponzetto, C. Bizer, M. Keuper, H. Stuckenschmidt (Eds.), LWDA 2018 - Lernen, Wissen, Daten, Analysen - Workshop Proceedings. GI-Workshop-Tage "Lernen, Wissen, Daten, Analysen" (LWDA-2018), August 22-24, Mannheim, Germany, Universität Mannheim, 2018.

[12] M. Kolbe, P. Reuss, J. Schoenborn, K.-D. Althoff, Conceptualization and implementation of a reinforcement learning approach using a case-based reasoning agent in a fps scenario, in: CEUR LWDA 2019 Workshop Proceedings 2454, CEUR, 2019.

[13] R. D. Gaina, M. Balla, A. Dockhorn, R. Montoliu, D. Perez-Liebana, Design and implementation of tag: A tabletop games framework, 2020. URL: https://arxiv.org/abs/2009.12065. arXiv:2009.12065.

[14] K. Sfikas, A. Liapis, Collaborative agent gameplay in the pandemic board game, in: Proceedings of the 15th International Conference on the Foundations of Digital Games, FDG '20, Association for Computing Machinery, New York, NY, USA, 2020. URL: https://doi.org/10.1145/3402942.3402943. doi:10.1145/3402942.3402943.

[15] D. Perez, S. Samothrakis, S. Lucas, P. Rohlfshagen, Rolling horizon evolution versus tree search for navigation in single-player real-time games, in: Proceedings of the 15th Annual Conference on

Genetic and Evolutionary Computation, GECCO '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 351–358. URL: https://doi.org/10.1145/2463372.2463413. doi:10.1145/2463372.2463413.

[16] G. N. Yannakakis, P. Spronck, D. Loiacono, E. André, Player modeling, in: Artificial and Computational Intelligence in Games, 2013. URL: https://api.semanticscholar.org/CorpusID:590230.

[17] C. Holmgård, A. Liapis, J. Togelius, G. N. Yannakakis, Evolving personas for player decision modeling, in: 2014 IEEE Conference on Computational Intelligence and Games, 2014, pp. 1–8. doi:10.1109/CIG.2014.6932911.

[18] M. Menzel, 2012. URL: https://legenden-von-andor.de.

[19] M. Menzel, 2012. URL: https://www.brettspiele-report.de/images/d/die-legenden-von-andor/Spielanleitung-Die-Legenden-von-Andor.pdf.

[20] UnityTechnologies, Unity 3d overview, 2021. URL: https://unity.com.

[21] A. Stahl, T. Roth-Berghofer, Rapid prototyping of cbr applications with the open source tool mycbr, in: Advance in Case-Based Reasoning, Proceeding of the 9th European Conference on Case-Based Reasoning, 2008.

[22] A. Hundt, P. Reuss, C. S. Sauer, T. Roth-Berhofer, Knowledge modeling and maintenance in mycbr 3, in: Proceedings of the FGWM workshop at Learning, Knowledge, Adaptation conference 2014, Springer, 2014.

[23] J.-J. Bartels, S. Viefhaus, P. Yasrebi-Soppa, P. Reuß, K.-D. Althoff, Visualizing the behaviour of cbr agents in a fps scenario, in: D. Trabold, P. Welke, N. Piatkowski (Eds.), Lernen, Wissen, Daten, Analysen. GI-Workshop-Tage "Lernen, Wissen, Daten, Analysen" (LWDA-2020), September 9-11, Online, CEUR, 2020, pp. 130–141.