Towards playing AIs for 7 Wonders: main patterns and strategies for 3-player games

Rafael Bettker
Graduate Program in Computer Science
Universidade Federal do Rio Grande do Sul
Porto Alegre, Brazil
rvbettker@inf.ufrgs.br

Pedro Minini
Graduate Program in Computer Science
Universidade Federal do Rio Grande do Sul
Porto Alegre, Brazil
ppminini@inf.ufrgs.br

Gabriel Pereira
Undergraduate Program in Information Systems
Universidade Federal de Santa Maria
Santa Maria, Brazil
ggpereira@inf.ufsm.br

Joaquim V. C. Assunção
Department of Applied Computing
Universidade Federal de Santa Maria
Santa Maria, Brazil
joaquim@inf.ufsm.br

***Abstract*—Artificial Intelligence (AI) in games may be greatly supported by a robust set of patterns, rules, and strategies that contribute to securing a win. In this work, we made a Knowledge Discovery pipeline – involving data selection, association, classification, and clustering – to successfully identify important factors that help to achieve a win in the board game 7 Wonders, focusing on 3-player matches. Our results show strong patterns and main strategies used by the best players in the world. This knowledge narrows the search for a Nash equilibrium for the game, getting us closer to create a top-tier AI.**

***Index Terms*—KDD, Data mining, Nash Equilibrium, board games, 7Wonders**

I. Introduction

Finding the balance for a game is a challenging undertaking. Balancing a game is a task that involves the game design and its many numerical attributes [1], [2]. Aiming to keep the players' involvement and the competition fair, current video games rely on constant updates to keep the balance among characters and strategies. However, it is not possible to update board games. Once they are on the market, without expansions, the balance will remain the same. On the other hand, most board games do not depend on a player's ability to react quickly and press buttons in synchrony. The main factors are strategy, tactics, and luck. From all sets of possible strategies in a game, one cannot be improved in the sense of average gains. If no player can deviate to a different strategy to achieve an average high score, such a strategy is considered a Nash equilibrium [3]. For instance, in a Rock-Paper-Scissors game, playing only any one of the three, with an equal probability, is a Nash equilibrium strategy; because in the long run, the player will win about $1/3$ of the matches and lose about $1/3$ of the matches. If an opponent plays paper often, the player can switch to scissors to explore the opponent's behavior. However, the player ends up vulnerable to exploitation (by rock), a nonequilibrium strategy.

Theoretically, the Nash equilibrium exists in all finite games and many infinite games. Recent Artificial Intelligence (AI) improvements try to exploit such equilibrium rather than explore opponent weakness [3]. However, finding the equilibrium for a non-trivial game is a challenging task. In fact, the category of finite games narrows down to two-player zero-sum games and mostly deterministic games. Even for these games, such a task can be difficult considering the calculation for the combinations of game states; given any non-trivial game, the amount of power required is beyond personal computers. Therefore, the first step is to obtain viable strategies; then identify the main strategies according to the likelihood of success.

Regarding game AI, the state-of-the-art uses neural networks with reinforcement learning supported by an initially supervised learning step [4]. Such a combination of techniques means that even deep learning models with super-computers need a first step of supervised learning, so the model can learn the essential rules for a complex game and make a filter of viable strategies. Furthermore, deep learning models often rely heavily on GPUs and are time-consuming, making the use of initial supervised learning useful for identifying basic patterns among viable strategies.

In a game of checkers [5], chess [6], or go [7], we have a relatively simple structure and tons of logs to be learned from algorithms. In more complex electronic games, such as StarCraft II, we have vast amounts of data, and we can use bots to play thousands of matches in a few days [4]. However, we do not have enough data for feeding complex models for many other board and card games. Furthermore, we also lack a fast forward machine *vs.* machine for deep learning training; thus, even with a deep learning model to learn tactics by playing, it is not possible to achieve a robust model. In fact, regarding the board game 7 Wonders, we do not even have a game log (hand-by-hand), making it difficult to find a proper set of winning strategies to explore. We only have match statistics, which describe matches as a whole (not action-by-action).

This work shows a complete Knowledge Discovery Process (KDD) to achieve patterns, rules, and winning strategies regarding 7 Wonders, the board game. Specifically, we explore large amounts of data from the best-ranked players from Board Game Arena (BGA)[1]. Due to less randomness – and therefore, more weight to strategy –, the best players tend to choose 3-player games. Although proving a Nash equilibrium is beyond our scope, this study shows results of an extensive mining process, filtering patterns, and strategies, which are statistically robust and used for a lightweight AI capable of challenging high-level players. As we highlight the leading strategies and patterns in 7 Wonders, our results are important to the development of top-tier AI for 3-player matches.

II. Related Work

There is a myriad of works that uses data mining to achieve patterns in games. However, only a few explore board games with a partial view and stochastic factors, such as 7 Wonders. Most are performed in digital games with an abundance of data generated by game logs. Furthermore, our process uses not only computational methods but also empirical validation from hundreds of pages available in guides made by some of the best 7 Wonders players (see Section IV). Although our contribution might become evident due to the unique aspects of our work, there are a few works that have a similar approach.

Consider the game of Texas hold'em poker, similar to 7 Wonders in having partial visibility, randomness by drawing cards, and stochastic scenarios in which players can infer an output with a certain probability. Silva and Reis [8] showed that data mining techniques could be used to create models of real players. Data from matches of professional Poker players were used to feed different algorithms. These algorithms were competing against each other to have a better game performance. The final model could make decisions close to professional players. Watson and Rubin [9] used case-based reasoning with different approaches to train and find a bot closer to defeat human players. Brown *et al.* [3], created Pluribus, a AI capable of defeating elite human professionals in six-player Texas hold'em poker.

Siqueira *et al.* [10] and Odierna *et al.* [11] used data mining to analyze and identify patterns in the players' behavior in World of Warcraft. They used clustering and regression models to identify patterns of players and indicate if the player will stop playing soon.

Oliveira *et al.* [12] worked to form team compositions to increase the victory rate in the game League of Legends. The data were collected from professional matches, observing the choice of characters and the game's outcome. They generated a tree with several combinations of possible suitable compositions. Using linear regression, they identified teams with a higher chance of winning, helping the player put together a composition that increases a team's overall chances of winning.

[1]www.boardgamearena.com

Similarly, Araujo *et al.* [13] used data mining to build an item recommendation system for League of Legends, and then compared the two approaches utilized: a recommender system based on association rules mining (comparing the Apriori and Eclat algorithms) and a recommender system based on classifiers (comparing decision trees, logistic regression, and artificial neural networks).

In order to find a player's profile, Benmakrelouf *et al.* [14] used multiple linear regression and the K-means algorithm. The first is used to perform data analysis, and the second is to extract players' groups to identify common characteristics. Considering that the characteristics of the players influence their performance, the clustering algorithm was able to find patterns between the variables analyzed and identified different types of players.

Robilliard *et al.* [15] wrote about the Monte-Carlo tree search algorithm to create an AI for the 7 Wonders game. They implemented using a Monte-Carlo tree search with susceptible levels, in which the nodes correspond to the possibilities of plays. A second AI was implemented deterministically, using fixed rules. Concluding the experiment, they compared the two AIs, showing that the first one provided better results.

Similar to our approach, most of these works use data mining to achieve winning patterns from players' behavior [10]–[14]. However, unlike these works, we performed a complete KDD process for a board game with no hand-by-hand log. We also used a combination of classification, clustering, and association rules to find solid patterns and the main strategies in the game. Finally, instead of exploring ordinary players' behavior and find the best scenarios, our process has a goal to learn from the data of the top players and get closer to a Nash equilibrium, similar to Brown *et al.* [3].

III. 7 Wonders overview

7 Wonders is a board game where 3 to 7 players receive a board representing one of the seven wonders of the ancient world. The game is composed of boards (Wonders) and cards (such as buildings and productions). The game is split into three ages, each using its card deck. In each age, seven cards from the deck are randomly distributed to each player, who in each round chooses a card to play and passes the others to the next player until only one card remains, which indicates the end of an age. At the end of each age, there are military conflicts between adjacent players that can give or take victory points from each, depending on the outcome and the age of the battle. Each card can represent a structure or a resource, being divided into seven types:

– Military structures (red cards) produce shields for the player, which increases their military power and is used to win conflicts at the end of each age.
– Commercial structures (yellow cards) have varied effects, such as giving the player extra coins or reducing the cost of buying resources from neighbors.
– Scientific structures (green cards) give the player one of the three scientific symbols, which generate victory points through combinations.

- Civilian structures (blue cards) give the player a fixed number of victory points.
- Raw materials (brown cards) and Manufactured goods (gray cards) are the resources.
- Guild (purple cards) produce victory points according to a determined condition, mainly linked to the status of the neighbor players.

To build a structure (card), a player needs to have the necessary resources for its construction, which can be raw materials or manufactured goods. It is also possible to buy resources from neighbors (using coins) or build a structure for free, in the case of chains (one card "allows" another). Figure 1 displays how is the board of a player during a match from Board Game Arena (BGA)[1].

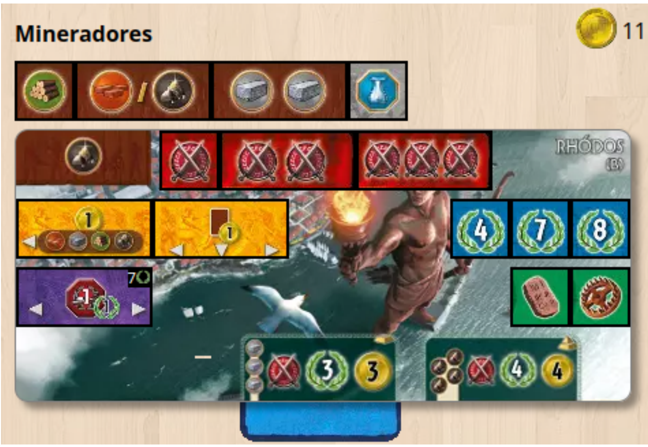

Fig. 1. Colossus of Rhodes' board along with the cards played. A screenshot from BGA showing only the part of the cards that matters; productions, shields, commerce, points, etc.

Besides building structures, a player can also discard one to earn three coins or build a Wonder stage on his board, paying its cost. The Wonder of the board can have 2 to 4 stages and can produce, among others, coins, victory points, and military shields.

Like most card games, 7 Wonders is a stochastic strategic game where players deal with other players' actions. After each age, the players have to ponder or change their strategies based on the situation in the first age. The winner is the one who accumulated the most victory points at the end of the third age. Detailed rules can be found in the game manual[2].

IV. Mining 7 Wonders match logs

As far as we know, there are only two implementations of the game digitally available, one from the game's developer for mobile devices and another on the Board Game Arena (BGA) website[1]. The match logs were extracted from BGA, which has hundreds (perhaps thousands) of games of 7 Wonders played every day. At the end of each game, a set of statistics is available for each player, making it possible to know the source of the winner's points or where the losers made mistakes.

The game statistics are composed of 25 different columns, such as victories points from each type of card, the amount of each type of card played, and the number of Wonder's stages built. Also, BGA provides a history of matches and a ranking system, giving access to the matches logs of high-level players that can generate more consistent results. Since they are not hand-by-hand logs, we need to find helpful information using only the data corresponding to the final statistics of each match. Figure 2 illustrates our process in a high level view.

To define a minimum number of rules, categories, and data attributes, we previously created a series of questions regarding the game [16]. These questions were based on a few analyses made by experienced players and the board game community. In terms of depth, the best guides are available in BGA forums[3] and in the BoardGameGeek forums[4], being the first[3] made by the current (2020) top player in BGA (known as "Pistolero").

On the other hand, the attributes used for classification and clustering were based only on the victory points by cards. Classification answered the following question, "*given a distribution of points, by type of card, which position shall a player fall?*". In other words, "*What are the main attributes, leading to 1st, 2nd, and 3rd place?*". Clustering was performed to know the average distribution of points for the main strategies, "*given the two main strategies, what is the average distribution of points by each type of card?*". Each of these tasks are described in the Sections IV-C, IV-D, and IV-E. Also, for each task, the results are detailed in Section VI.

A. Data selection

To extract and format the data, we implemented a script that generates a CSV file. We selected the best players in the BGA ranking as a source for collecting the data. We discarded the games played with other players below a predetermined level since the platform allows low-level players to enter public rooms of high-level players, resulting in an unbalanced game. Thus, we have the results belonging to the highest-level games in the entire platform.

Since the game configuration depends on the number of players and our goal is matches of 3 players, a dataset was made with all the 3-player matches of the top five players in the BGA (ranking in March 2020). A dataset of 13699 lines was generated[5] corresponding to 4566 3-player matches, which is 53.92% of all matches played by the selected players. Each row in the dataset corresponds to one player so that a 3-player match will fill three lines of data in the set. Each column represents an attribute corresponding to a match. Table I shows an example.

[2]https://cdn.1j1ju.com/medias/c8/d6/88-7-wonders-rule.pdf

[3]https://boardgamearena.com/forum/viewtopic.php?f=192&t=14557

[4]https://boardgamegeek.com/thread/691370/some-complex-strategies-7-wonders

[5]The data and scripts are publicly available at https://github.com/dmag-ufsm/Mining

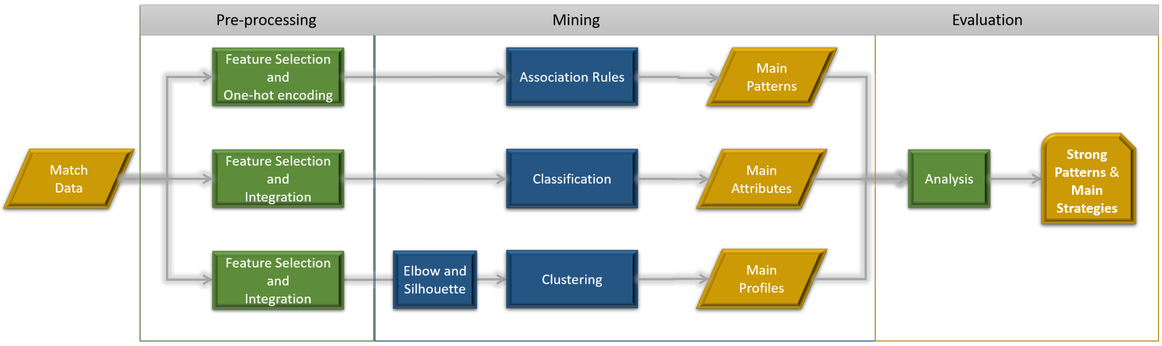

Fig. 2. Workflow to generate strong patterns and strategies.

TABLE I
EXAMPLE OF THE DATA COLLECTED

Place	VP total	Thinking time	VP from Military Victories	...
1	51	4:03	10	...
2	45	5:33	13	...
3	45	5:24	4	...

B. Pre-processing and Integration

Not all extracted data was ready to be used. When a player leaves the game before it is finished, the data is broken. When such interruption occurs in the BGA, the game is canceled, and a tie is declared between all the other players. However, it is kept in the game history, leading to incomplete and incoherent statistics. Since these data are irrelevant, we verified the dataset, removing the matches that finished before the time, resulting in 4504 matches. Thus, our dataset is composed of only complete games with consistent data.

For using the Apriori and Eclat algorithms, we transformed the data into presence matrices, also known as one-hot encoding. To fill each matrix and use a different amount of cards according to our questions and previously established criteria [16], we defined parameters such as an arithmetic mean, quantile, as well as fixed values. Thus, values greater or equal to a parameter were converted to 1; 0 otherwise.

C. Getting association rules

Using the generated matrices (one-hot encoding), we applied the Apriori and Eclat algorithms to generate the association rules. The first algorithm works from two given parameters: the minimum support and the minimum confidence. A support for an attribute (or a set of attributes) A implying attribute (or a set of attributes) B, is given by: $Sup(A \Rightarrow B) = P(B|A) = \frac{A \cap B}{total(T)}$. Confidence, $A \Rightarrow B$ is given by: $\frac{A \cap B}{A}$ [17]. A and B are also commonly represented as Left Hand Side (LHS) and Right Hand Side (RHS). First, Apriori checks to see if the items meet the minimum support. After that, the support is used again to validate the generated combinations. Finally, the rules are strong if they are above minimal confidence. The Apriori property guarantees that infrequent items eliminate the computation for any set that involves these items. This makes the algorithm feasible to compute large datasets.

The second algorithm generates a frequent itemset from a given dataset. As we only have the most frequent set of items, we still need to find the association rules from it; the *ptree* method is used for such a task. In this method, transactions will be counted in a prefix tree, and the rules will be selectively generated using the counts in the tree. This approach is generally faster than Apriori.

We generated two larger sets of rules, from different matrices, one based on generated symbols, for instance, military cards may hold one to three shields; and the other based on the number of cards. In each of these sets, we used as a discrimination metric the constants 1 and 2. As parameters for analysis, we used the first quantile, the average, and the third quantile. The constants are merely for marking the use of a card, one or two times, while the average and the quantiles seek to obtain patterns on the number of uses for each card or type of card. For instance, if the average number of yellow cards was 2, the presence matrix will generate 1 for every player in a match in which they used more than 2 yellow cards. Using these categories, we got a total of 11122 rules with a support of 50%, being 12% of them based on the average and the others based on quantiles.

D. Generating classification trees

More than strong rules, we need to know the main attributes that lead to 1st place. To achieve this goal, we used classification trees [18] to find essential attributes and which are the central values separating these attributes. The trees generated a flow of attributes and victory points, which leads players to victory, second position, or defeat. Before generating the trees, the data was shuffled to avoid bias from the selected data. Afterward, the data was split between training (80%) and testing datasets (20%). We then proceeded to generate the trees, predicting the class *Place* (y, the position of a player, 1st, 2nd, or 3rd) with different inputs (x) such as the number of shields, raw and manufactured goods, and discarded cards. However, the best results were found by joining different types of victory points as inputs. The results can be seen in Section VI-F.

E. Clustering

The classification gave us the main attributes, which are also the main strategies for the game (empirical verification[3]). Now the question is "*what are the most common profiles for each strategy?*". Such information was extracted using clustering on the subset of winners (1st place). With the results, we made a small player profiling ("game styles") based on each of their total victory points at the end of the game. The results ended up being closely associated with our findings based on the classification trees. The results can be seen in Section VI-G.

V. Dataset Statistics

In order to have an overview of the dataset and understand the distribution and the results, we made a statistical analysis of the attributes. Table II shows the minimum (Min), maximum (Max), mean ($\overline{x}$), standard deviation (σ), and the first (Q_1), second (Q_2) and third (Q_3) quantiles obtained from each attribute from the matches of our dataset. In addition to these attributes, there is also each player's thinking time, which is not relevant and was removed.

TABLE II
Top player matches statistics

	Min	Max	$\overline{x}$	σ	Q_1	Q_2	Q_3
Total VP	26	90	53.83	7.36	49	54	59
Military VP (victory)	0	18	8.14	6.25	1	9	13
Military VP (defeat)	-6	0	-2.56	2.00	-4	-2	-1
Treasury VP	0	15	3.60	2.48	2	3	5
Wonder VP	0	20	6.90	4.89	3	7	10
Civil VP	0	48	15.63	8.56	9	15	22
Science VP	0	85	11.87	16.53	0	1	21
Commerce VP	0	15	2.57	2.80	0	3	4
Guild VP	0	35	7.68	6.68	0	7	12
Wonder stages built	0	4	2.57	0.85	2	3	3
Discarded cards	0	8	0.71	0.88	0	0	1
Chained constructions	0	10	2.69	1.89	1	2	4
Coins spent (commerce)	0	36	9.87	4.61	6	10	13
Coins given (commerce)	0	34	9.87	5.26	6	9	13
Shields	0	12	4.66	2.80	3	5	7
Civilian structures	0	10	3.20	1.78	2	3	4
Scientific structures	0	12	2.58	3.14	0	1	5
Guilds	0	5	1.16	0.91	0	1	2
Military structures	0	6	2.34	1.31	1	3	3
Commercial structures	0	8	2.44	1.49	1	2	3
Raw materials	0	8	2.58	1.14	2	3	3
Manufactured goods	0	3	0.97	0.92	0	1	2

The major sources of victory points (VP) are *civil*, *scientific*, and *military cards*, averaging 15.63, 11.87, and 8.14 points, respectively. The victory points through scientific cards are those with the highest standard deviation (16.53) since they are cards that generate more points when played in large numbers; thus, players who successfully play these cards get a high number of points, while many players do not even try to play these cards (close to 0 points). This can also be noticed through the quantiles – while the median is 1, the third quantile is 21. Guild VPs also have a high standard deviation, 7.68, as guilds are cards with more specific effects, and not all players end up playing them. The average for total victory points was 53.83, and the third quantile was 59 points; thus, scores close to 60 were enough to win the match.

Regarding each type of structure, the most played is *civil*, as they give direct victory points and can be used with any strategy. *Scientific structures* have the most significant standard deviation of all (same reasons as science VP). *Military structures* may be important, but they are the third least played type, probably because the goal is to have more than the other players, no matter how many; thus, ideally, players play just enough cards to be ahead. *Guilds* are used less, averaging just 1.16 per game, partially because there are only five cards, and they are only available in age III. *Manufactured goods* have an average of less than 1, as few cards need this type of resource compared to raw materials.

As for the other attributes, *Wonder stages built* is how many stages on the board were built by a player; most boards have 3 stages to be built, exceptions are The Colossus of Rhodes B (2) and The Pyramids of Egypt B (4). *Discarded cards* is how many cards have been discarded by the player, an action that can be done in exchange for 3 coins, but which has an average of less than 1 card and a median of 0, showing that this choice is not popular with high-level players.

If the player has a specific card already played, the current card (some) can be played for free without expending resources (chained constructions). In high-level matches, players tend to use this technique about 3 times during a match. However, the distribution is unequal since science cards are usually used as chained constructions. Spent and given coins correspond to commercial exchanges, which is how many coins a player spent buying resources or making his resources available to other players. Both attributes have the same average (9.87) since each coin given by one player is received by another, which is 5 resources obtained (2 coins each). Shields are used in military conflicts at the end of each age. On average, players have 4.66 shields with a standard deviation of 2.8; Table III has a more detailed of how they can influence military victory points. For instance, a player with 5 shields will receive an average of 6.9 VPs in conflicts, considering the victories and defeats – a value that can be achieved with two military structures (one from the second age and one from the third age).

TABLE III
More shields, more Victory Points

Shields	VP Victory	VP Defeat	Balance	VPs / shield
0	0	-5.83	-5.83	-5.83
1	1.1	-4.07	-2.97	-2.97
2	2.5	-4.35	-1.85	-0.93
3	4.25	-3.4	0.85	0.28
4	7.56	-2.4	5.16	1.29
5	9.16	-2.26	6.9	1.38
6	9.98	-1.51	8.47	1.41
7	14.08	-1.03	13.05	1.86
8	14.49	-0.96	13.53	1.69
9	14.62	-0.8	13.82	1.54
10	15.09	-0.9	14.19	1.42
11	15.02	-0.77	14.25	1.3
12	16.5	-0.5	16	1.33

The hot-spot for military is around 7 shields. This quantity leads to an average balance of 13.05 VPs, almost 5 VPs more than having a shield less, and a better $VPs/shields$ rate.

Hence, players can obtain more than 7 shields to secure more military points, but the $VPs/shields$ rate drops.

Players who do not want to spend their moves on military structures lose an average of 5.83 VPs (maximum is 6, -1 for each of the 6 conflicts during the match), but playing only one military card in the first age (receiving 1 shield) results in approximately half the loss, 3 VPs. The boards are another critical aspect of the gameplay: there are 7 boards, each with two sides (A and B) and having different free resources and Wonder's stages. Table IV shows the game boards' main statistics, such as how many times each board was used, how many wins, and the average VPs earned from each type of card.

TABLE IV
TOP PLAYERS' BOARD STATISTICS

Board	Uses	Wins	Win Rate	Military VP	Science VP	Wonder VP
Giza A	373	96	25.74	7.59	3.91	12.30
Babylon A	210	53	25.24	6.18	21.49	5.86
Olympia A	543	208	38.31	9.55	4.44	9.44
Rhodes A	185	45	24.32	11.80	2.34	7.58
Ephesos A	41	9	21.95	5.24	8.39	8.14
Alexandria A	64	6	9.38	4.54	13.35	7.90
Halikarnassus A	164	19	11.59	5.66	18.18	6.39
Giza B	1488	448	30.11	7.80	2.47	16.41
Babylon B	1731	478	27.61	5.73	23.90	2.97
Olympia B	1401	393	28.05	9.15	4.79	4.86
Rhodes B	1749	718	41.05	11.97	3.19	5.70
Ephesos B	1912	656	34.31	7.02	14.18	8.61
Alexandria A	1847	677	36.65	8.80	9.42	6.08
Halikarnassus B	1804	711	39.41	6.65	25.19	2.73

Regarding player preference of the boards' side, B is clearly the chosen side for all the boards, with 88.31% of use. This value is likely higher since some matches in BGA may have been made with the chosen side at random (this is a game configuration option, defined by the host), thus using more times the A side. The side choice also reflects the overall win rate, while side A had 27.59% of victories, side B had 34.20%. Olympia A was the only side A with a higher victory rate than its side B (and third-highest among all boards). It was also the most played side A. The best board (considering the win rate) is Rhodes B with 41.05%, and the worst is Alexandria A with only 9.38%.

Among the boards, the ones that most obtained military VPs are both Rhodes, 11.97 in B side and 11.8 in A side. This is expected, since it is the only Wonder that gives shields in its stages, being able to earn 2 shields on both sides.

For scientific VPs, Halikarnassos B (25.19), Babylon B (23.90) and Babylon A (21.49) stand out. Both sides of Babylon have a Wonder stage that provides an extra scientific piece. Halikarnassos B has three Wonder stages, all with an effect that gives a player the chance to play a discarded card at the end of an age. This can be useful when the player cannot play a card, so the player can sell the card and play it at the end of the age – in the best scenario, allowing the player to add three science cards. All boards containing their stage effects can be consulted in the game manual.

VI. RESULTS

The following subsections A-E contain the strong rules achieve by association. Subsection VI-F shows the main attributes for optimal strategies achieved through classification trees. Subsection VI-G shows the players' profile and distribution of card points regarding the two dominant strategies.

A. *If war is inevitable, face it or perish*

Starting with a shield (or two) is a simple and effective tactic for a military-driven strategy. When it comes to a 3-player game, the player's two opponents are involved in the conflict resolution, making war inevitable. Thus, winning a conflict means getting points and removing points from the direct opponents (neighbors); the opposite is also valid. The association rules related to military cards are shown in Table V.

TABLE V
MORE MILITARY CARDS THAN THE AVERAGE

LHS	RHS	Support	Confidence	Lift
victory	militaryStr=1	0.22	0.65	1.25
militaryStr=0	defeat	0.37	0.76	1.14
raw material=1	militaryStr=1	0.28	0.55	1.06

Considering the table above, the average use of military cards (structures), among the players who won the game, 65% invested in military cards, with a support of 22%. At the same time, those who played few (or none) military have 76% confidence for losing a game. In addition, 76% of players who have not invested in military might have lost the match. Compared to our previous study [16], we find the military confidence dropping (78% $\rightarrow$ 65%). This indicates that top players are less susceptible to early intimidation; *i.e.*, when a player plays two military cards in the early game to make the others give up on military. Nonetheless, 0.65 confidence is a strong pattern for high-level competitive players. All military structures need raw materials to be played, making this resource necessary for players to invest in this strategy. According to the association rules, 55% of those with raw materials above the average also have military cards above average.

By winning all conflicts it is possible to earn a maximum of 18 points, which is about one-third of the points that the player will have in total. However, winning a conflict means preventing an opponent from receiving these points and reducing 1 point. If an opponent gives up on the conflict, 21 points can be achieved with a few cards. In a perfect scenario (does not happen among good players), an initial card can be worth 21 points. The average scenario among top players can be seen in Table III.

B. *Coins should to be spent*

Experienced players understand that holding a large number of coins to "transform" them into victory points is not a good strategy. However, being out of coins can make a player lose an opportunity to construct an important card. The data show us that it is better to stay below the average regarding the number of coins. Table VI shows the rule regarding the average amount of coins on the matches.

TABLE VI
MORE COINS THAN THE AVERAGE

LHS	RHS	Support	Confidence	Lift
treasure=1	defeat	0.29	0.64	0.95

Considering the players' average, those who accumulate more than half of the total value (we call it "treasure") tend to lose. The apparent reason is that every 3 coins score only 1 point. The countereffect makes it not so evident as large amounts of coins imply debt for the other players (point transferring), mitigating the effect. The balance is subtle, but it is more advantageous to spend with base resources for building cards that grant more victory points at the end of the game.

C. Trading leads to money

Coins are an essential attribute in the game; not having them means not being able to buy resources from opponents, decreasing the range of cards available to be played, and causing the player to lose the opportunity to play more profitable cards. An important factor in having a coin reserve is commercial cards.

TABLE VII
MORE COINS THAN THE AVERAGE

LHS	RHS	Support	Confidence	Lift
treasure=1	commercialVP=1	0.32	0.69	1.35
commercialVP=1	treasure=1	0.32	0.62	1.35
treasure=0	commercialVP=0	0.35	0.64	1.31
commercial=0	treasure=0	0.35	0.71	1.31

Table VII shows strong rules regarding the relationship between commercial structures and a large number of coins. Players who finish the game with a large number of coins are the ones who have a high number of commercial structures (69%), while those who have few commercial structures also have few accumulated coins (71%). These rules also apply to the reverse with reasonable confidence and support of more than 30%.

This rule is bonded with the previous one, which is not a good strategy since coins above the average, in most cases, leads to defeat, as seen in Section VI-B. Therefore, commercial structures should be played in moderation, trying to stay at a minimum to guarantee enough resources.

D. A science career opens doors, but demands sacrifices

Being able to play a high number of science cards makes it possible to win a game with only science points without depending on other types of cards. Science cards are known for their cumulative property, where a combination of identical symbols scores n^2. Also, each set of all three symbols gives more 7 points. However, as a unique card is only worth 1 point, the real advantage is large sets. These properties imply that a player going for a science strategy should take at least 3 of a kind to take advantage of science cards. Our rules detected that successful players, in a 3-player game, usually tend to sacrifice everything else to focus on science. This is a known strategy in the 7 Wonders community known as the *mad scientist*, which we focus on later. The necessary production and raw material for playing the cards are replaced with cards chain and commerce, even when that means selling cards for 3 coins. Table VIII shows these rules, which are the ones with the highest confidence in our sets. An exception to this rule is the manufactured goods, which are necessary to play the first three cards. As each card allows two other cards, the player can use the science chain. This holds for the matches with 3 to 7 players, but the strongest rules were found in 3 player matches, which is the focus of this work.

TABLE VIII
MORE SCIENCE CARDS THAN THE AVERAGE

LHS	RHS	Support	Confidence	Lift
scientificStr=1	raw materials=0	0.27	0.66	1.33
scientificStr=1	militaryStr=0	0.29	0.70	1.44
scientificStr=1	commercialStr=0	0.29	0.72	1.30
scientificStr=1	civilianStr=0	0.33	0.81	1.40
scientificStr=1	guildStr=0	0.32	0.79	1.17
scientificStr=1	manufacture=1	0.31	0.75	1.18

Table VIII shows that the "mad scientist" strategy is successful even among the best players. Also, it is the only strong rule that counters a military-driven strategy. The lower usage of other cards is due to the use of science chain and the sacrifice of other constructions to get as much science as possible. Table VIII shows the strong rules for these cases, only manufactured goods are played above the average, which implies that most successful science strategies are driven from the science opening to new science cards (chain).

E. Use woods for buildings

Table IX describes the action of overproduction, *i.e.,* players that create too many raw materials (above the third quantile) tend to lose. Both the support and confidence are strong for these rules. Players perform such actions to guarantee the best cards or incoming money on a trade. However, this tactic does not pay off because it is better to pay a few coins for the best cards than lose two or three cards to create raw material.

TABLE IX
RAW MATERIAL ABOVE THE THIRD QUANTILE

LHS	RHS	Support	Confidence	Lift
raw materials=1	defeat	0.35	0.69	1.02

Another case is the underproduction, which the best players know well and solve this problem with commerce (see Table X). Otherwise, when investing in materials (near the average), players do not invest heavily in commerce, *i.e.*, an average production of raw material implies lower usage of commerce. These rules hold for all combinations of players but were slightly stronger on an odd number of players. This can be explained due to the distribution of the cards.

F. Optimal strategies

Figure 3 shows the classification tree generated from the processed data of 3-player matches. To predict the class

TABLE X
RAW MATERIAL AND COMMERCE CONSIDERING THE AVERAGE

LHS	RHS	Support	Confidence	Lift
raw materials=0	commercial=1	0.24	0.50	1.12
commercial=1	raw materials=0	0.24	0.55	1.12
raw materials=1	commercial=0	0.31	0.60	1.09
commercial=0	raw materials=1	0.31	0.55	1.09
manufacture=1	commercial=0	0.23	0.63	1.42
commercial=0	manufacture=1	0.23	0.52	1.42

"Place", we combined the different types of victory points as input: $fit(Y = Place, X = WonderVP + ConflictVicVP + CivilVP + CommerceVP + GuildVP + ScienceVP + TreasuryVP)$.

A training and testing split $(0.8, 0.2)$ was made with shuffled instances from 10809 observations. Among a few trials, this was the best classification tree we could find regarding details and accuracy[6].

Specifically, at each node, we have the class labeled 1, 2, or 3, indicating a placement. In Figure 3, with the root node as an example, the full set of placements is divided between those players that had $conflict.victory >= 16$ (17%) and those that did not pass this condition (83%). The Gini index is the metric behind the split that continues until the leaf nodes, which indicates the final classification with its respective percentage of observations and the predicted probability of the classification.

The result shows that the military and scientific strategies are the most important factors to divide the tree, thus the most important to secure first place for high-level players. These main strategies can be seen both in the percentage of observations and the predicted probability of the classification.

Although the predicted probability is far from optimal, in a theoretical sense, considering such high-level matches from a well-balanced game, any deviation from the ordinary $1/3$, for each class is an interesting finding. Furthermore, the classification confirms the strong rules (see the Sections VI-A and VI-D) and expand our understanding of the main attributes, as *guilds* and *treasury* appear in the middle nodes.

In summary, the military VP went over 60% for 17% of the instances, undoubtedly one of the main strategies considering the scenario. The same can be stated for science, over 60% for 7% of the instances; considering that science is a counter for military, it is defined as another main strategy. The tactics involved in each strategy will vary along with the game. The tree shows us a few variations where guilds and coins are the second and third most decisive factors.

Finally, we need to get the main profile for each of the main strategies, *i.e.*, the most common division of scores for players following the main strategies.

G. Main profiles

Using the K-means algorithm, along with the Elbow [19] and the Silhouette [20] methods, we determined the optimal number of clusters. As a continuation of the previous subsection, we selected only the observations from 1st place in each match and generated the clusters based on the resulting data.

[6]The tree was generated with the R language package *rpart*, with default parameters, except Complexity Parameter (CP), which was changed to 0.005

TABLE XI
CONTRIBUTION (APPROX. %) OF EACH VICTORY POINT TYPE PER CLUSTER

#	Size	Civil	Commerce	Guilds	Military	Science	Treasury	Wonder
1	1270	8.036	0.806	4.523	4.228	41.850	2.622	3.902
2	3247	18.122	3.205	10.800	12.924	2.486	4.257	8.238

From Table XI and Figure 4, we obtained two profiles for winners, the same main strategies found in the classification, one from a science-based strategy (cluster 1), another from a military-driven strategy (cluster 2).

In 7 Wonders, scientific resources have the potential to give players the greatest number of VPs out of any VP type, as a consequence of the science scoring equation: $sci\_VPs = tablet^2 + gear^2 + compass^2 + 7 \times completed\_sets$.

The $tablet$ is the total number of tablets, $gear$ is the total number of gears, $compass$ is the total number of compasses, and $completed\_sets$ is the number of complete sets of scientific cards the player has (one set is composed by one gear, one tablet, and one compass). Thus, the science scoring equation explains why the science-based strategy focuses only on obtaining scientific resources (known by the players as "mad scientist"). However, the scientific strategy is more difficult because other players can actively sabotage the strategy by discarding important cards.

On the other hand, the military-driven strategy needs to complement its military victories with other types of VP; from the clustering results, this is achieved mainly by civilian and guild VPs. The military strategy is also the most used between winners, with 3247 observations in cluster 2. While the clustering results do not provide us any unexpected results, they reinforce our findings, providing extra knowledge on the competitive game dynamics. Furthermore, the average distribution for each strategy is a new knowledge that drives us closer to the Nash equilibrium. The equilibrium itself is the best average strategy, in the long run, avoiding exploitation. Another interesting piece of information is shown in Table XII, each Wonder and its possible 18 VPs by using military cards.

TABLE XII
MILITARY VP CONTRIBUTION FOR EACH GAME BOARD (1ST PLACE)

	Number of Military Victory Points									
Board	10	11	12	13	14	15	16	17	18	Total
Giza A	5	2	8	9	4	3	8	10	6	55
Babylon A	2	3	5	0	1	2	2	1	4	20
Olympia A	20	6	20	2	8	7	22	16	49	150
Rhodes A	5	3	3	1	0	0	8	4	12	36
Ephesus A	0	1	1	0	0	0	0	1	1	4
Alexandria A	0	0	0	0	0	0	0	0	1	1
Halikarnassus A	0	0	2	1	2	0	0	0	2	7
Giza B	32	22	42	12	18	11	41	27	50	255
Babylon B	23	16	28	3	19	3	18	3	39	152
Olympia B	33	15	35	15	18	10	33	34	84	277
Rhodes B	40	28	60	39	37	12	134	60	170	580
Ephesos B	36	13	45	15	17	12	50	29	70	287
Alexandria B	48	32	61	18	33	9	86	33	109	429
Halikarnassus B	34	15	43	11	36	14	41	30	63	287
Total	278	156	353	126	193	83	443	248	660	2540
Rate	10.94	6.14	13.9	4.96	7.6	3.27	17.44	9.76	25.98	100

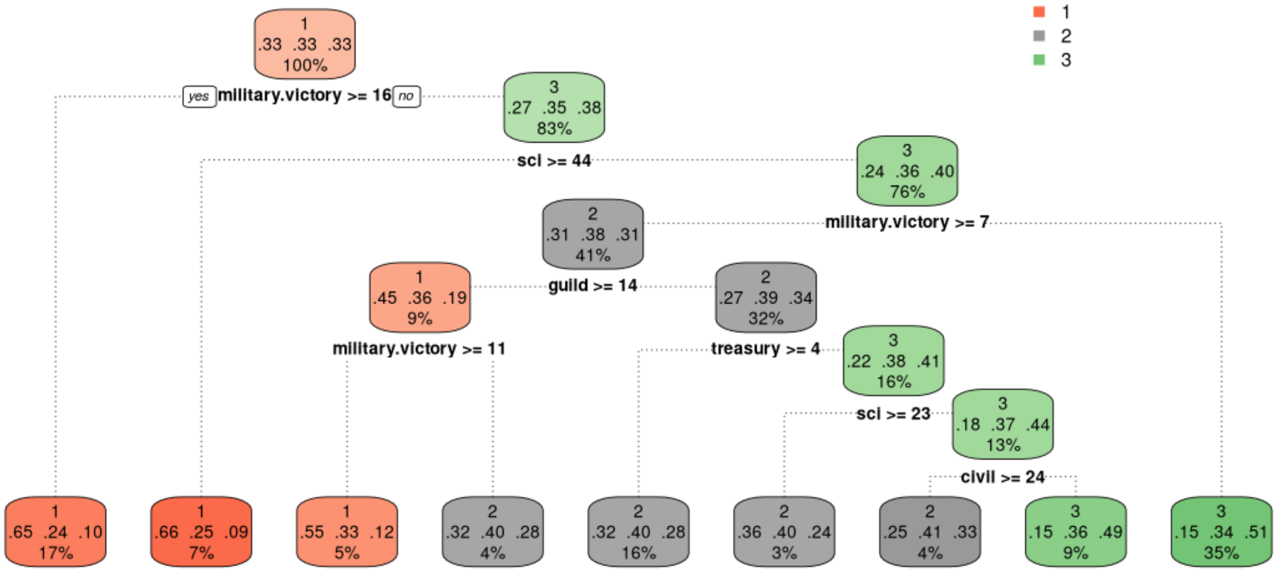

Fig. 3. 3-player matches classification based on the different types of victory points.

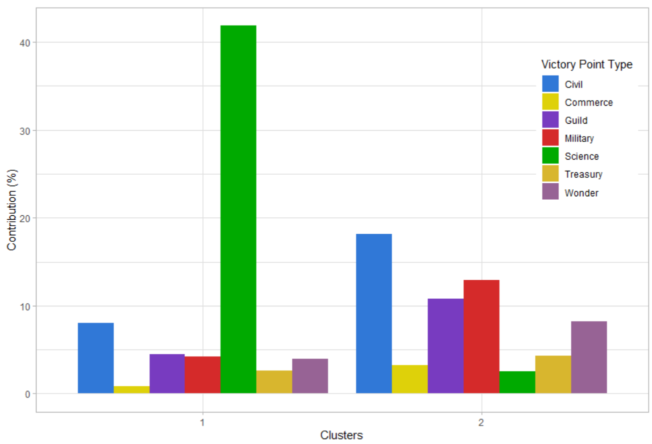

Fig. 4. Contribution of VPs on each cluster. 1st place players from 3-player matches.

From Table XII, we can see that between top players, the game boards Rhodes B and Alexandria B are the ones that are typically considered to give the most military VPs, and except for Olympia A, no A side boards tent to end with a reasonable amount of military VPs. Also, there is no unique hot-spot for military points; the most common is the maximum (18), but many players successfully win with 12 and 16 military points as well. On the other hand, 15 is a "cold-spot" where only in 83 matches players finished in 1st. Since there are many possible combinations of cards and boards, we do not have enough data to answer this phenomenon.

Although, we know that the only possible combination of points (see Table XIII) is: two victories in the first age (2), one victory in the second age (3), and two victories in the third age (10). The second "cold-spot" is 13 points, which can only be achieved with a defeat in the second age. Hence, winning the military in the second age is important.

TABLE XIII
POSSIBLE COMBINATIONS FOR MILITARY VP

	10	10	11	11	12	12	13	13	14	15	16	17	18
Age 1	0	2	1	0	2	1	0	2	1	2	0	1	2
Age 2	0	3	0	6	0	6	3	6	3	3	6	6	6
Age 3	10	5	10	5	10	5	10	5	10	10	10	10	10

Regarding the science VPs, Table XIV shows that between top players the game boards Halikarnassus B and Babylon B tend to be the most profitable choices for *mad scientists*, and not a single A side board is considered good enough. The hot-spot for most Wonders (e.g., Babylon B with 92) seems to be between 30 and 39 points, probably due to the opponents denying science cards.

TABLE XIV
SCIENCE VP CONTRIBUTION FOR EACH GAME BOARD (1ST PLACE)

	Number of Science Victory Points								
Board	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80-89	Total
Giza A	0	2	2	2	0	2	0	0	8
Babylon A	0	11	12	6	5	0	1	0	35
Olympia A	1	5	2	1	1	0	1	0	11
Rhodes A	1	1	0	0	1	0	0	0	3
Ephesos A	0	1	1	0	0	1	0	0	3
Alexandria A	0	1	0	0	1	0	0	0	2
Halikarnassus A	0	1	4	3	2	1	0	0	11
Giza B	0	4	7	4	2	1	0	0	18
Babylon B	0	56	92	67	62	31	8	1	317
Olympia B	0	5	4	8	4	0	0	0	21
Rhodes B	3	15	10	1	5	0	0	0	34
Ephesos B	1	48	59	55	38	22	0	0	223
Alexandria B	1	43	33	17	16	2	1	0	113
Halikarnassus B	0	62	125	108	103	48	13	0	459
Total	7	255	351	272	240	108	24	1	1258
Rate	0.56	20.27	27.9	21.62	19.08	8.59	1.91	0.08	100

Compared to the military-driven strategy, the *mad scientist* is restrictive regarding board choice. While technically all B sides and one A side are feasible for the military strategy (albeit not always recommended), only 3 (4, at most) boards are seen as viable by the top players for the scientific strategy. The following fact may explain this: fewer players are *mad scientists* ($< 50\%$ in relation to military-driven ones), as it is a more risky strategy. If a military dispute is a loss, it is easier to change tactics. On the other hand, the *mad scientists* requires abandoning raw material and focus only on science and its chains, making it more difficult to change tactics. Board-wise, the scientific strategy is just less flexible compared to the military one. When chosen at random, the player will most likely get a board where the military strategy is safer.

VII. Comparing Results

The results were obtained from the dataset collected in March 2020 of the top 5 players in BGA, which has 4504 3-players matches. For comparison, the previous work [16] was done with a dataset of 50 3-player matches from the top 20 players that were collected in February 2019. Despite the significant difference in the size of the dataset, the results remained similar. Almost all rules have kept their values or had a slight change [16]. The data also is in accordance with the "Pistolero's" guide[3] (best ranked BGA player), where the player makes a deep analysis of the game's tactics. Even informal, based on experience, this guide is a valuable comparison to verify our interpretations and better understand the results. Most findings are a match; point exceptions are a few positions in the Wonder's ranking (shown in Table XV), and the common use for a few Wonders (such as Halikarnassus B, mostly successfully used in science-driven strategies).

TABLE XV
Top 7 Boards by win rate (average values)

Wonder	Win Rate	Military VPs	Science VPs
Rhodes B	41.05	11.97	3.19
Halikarnasses B	39.41	6.65	25.19
Olympia A	38.31	9.55	4.44
Alexandria B	36.65	8.80	9.42
Ephesos B	34.31	7.02	14.18
Giza B	30.11	7.80	2.47
Olympia B	28.05	9.15	4.79

VIII. Final Remarks

This work described a KDD process, which retrieved the main strategies by discovering patterns among the data from top-tier global players. As far as our knowledge goes, this is the only data mining work with large amounts of data from the best 7 Wonders players. The patterns found are unique and allow us to go further in the pursuit of an optimal global strategy. Our data is conclusive that are two elite strategies, one military-based and another science-based. Different Wonders have a different focus; hence, the Nash equilibrium should be found for each combination of Wonders (taking the neighbors into account). The data also indicates a few strong rules and patterns that lead to a set of tactics that are better than others. However, finding a Nash equilibrium in a strategy is necessary to find the best tactics for each scenario, which requires a hand-by-hand log and considerable amounts of data. To achieve such a goal, we developed our implementation of the game [21], which already has a set of heuristic-based AIs playing against each other to produce enough data to partially feed a neural-network-based AI. Thus, as future work, we will deliver the platform along with the first AI to play against humans, with self-improvement and generation of a hand-by-hand log.

References

[1] K. Oxland, *Gameplay and Design*. Addison-Wesley, 2004.
[2] P. Schuytema, *Game Design: A Practical Approach*. Charles River Media Game Development, Charles River Media, 2007.
[3] N. Brown and T. Sandholm, "Superhuman ai for multiplayer poker," *Science*, vol. 365, no. 6456, pp. 885–890, 2019.
[4] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser, J. Quan, S. Gaffney, S. Petersen, K. Simonyan, T. Schaul, H. van Hasselt, D. Silver, T. P. Lillicrap, K. Calderone, P. Keet, A. Brunasso, D. Lawrence, A. Ekermo, J. Repp, and R. Tsing, "Starcraft II: A new challenge for reinforcement learning," *CoRR*, vol. abs/1708.04782, 2017.
[5] J. Schaeffer, *One Jump Ahead: Challenging Human Supremacy in Checkers*. Springer New York, 2013.
[6] M. Campbell, A. J. Hoane Jr, and F.-h. Hsu, "Deep blue," *Artificial intelligence*, vol. 134, no. 1-2, pp. 57–83, 2002.
[7] D. Silver, A. Huang, C. Maddison, A. Guez, L. Sifre, G. Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, pp. 484–489, 01 2016.
[8] N. Silva and L. P. Reis, "Professional poker players' modeling using data-mining," in *2016 11th Iberian Conference on Information Systems and Technologies (CISTI)*, IEEE, 2016.
[9] I. Watson and J. Rubin, "Casper: A case-based poker-bot," in *AI 2008: Advances in Artificial Intelligence* (W. Wobcke and M. Zhang, eds.), (Berlin, Heidelberg), pp. 594–600, Springer Berlin Heidelberg, 2008.
[10] E. S. Siqueira, C. D. Castanho, G. N. Rodrigues, and R. P. Jacobi, "A data analysis of player in world of warcraft using game data mining," in *2017 16th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*, pp. 1–9, IEEE, 2017.
[11] B. A. Odierna and I. F. Silveira, "Player game data mining for player classification," in *Proceedings of SBGames*, 2018.
[12] V. da Costa Oliveira, B. J. Placides, M. d. F. O. Baffa, and A. F. da Veiga Machado, "A hybrid approach to build automatic team composition in league of legends," in *Proceedings of SBGames*, 2017.
[13] V. Araujo, F. Rios, and D. Parra, "Data mining for item recommendation in MOBA games," in *Proceedings of the 13th ACM Conference on Recommender Systems - RecSys '19*, ACM Press, 2019.
[14] S. Benmakrelouf, N. Mezghani, and N. Kara, "Towards the identification of players' profiles using game's data analysis based on regression model and clustering," in *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, IEEE, 2015.
[15] D. Robilliard, C. Fonlupt, and F. Teytaud, "Monte-carlo tree search for the game of "7 wonders"," in *Workshop on Computer Games*, pp. 64–77, Springer, 2014.
[16] J. Assunção, G. Pereira, J. Acosta, R. Vales, and L. Rossato, "Data mining 7 wonders, the board game," in *SBC - Proceedings of SBGames*, 2019. ISSN:2179-2259.
[17] C. Kaur, "Association rule mining using apriori algorithm: a survey," *International Journal of Advanced Research in Computer Engineering & Technology (IJARCET)*, vol. 2, no. 6, 2013.
[18] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Monterey, CA: Wadsworth and Brooks, 1984.
[19] R. Thorndike, "Who belongs in the family?," *Psychometrika*, vol. 18, no. 4, pp. 267–276, 1953.
[20] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53 – 65, 1987.
[21] J. Jardim, R. Bettker, P. Minini, G. Pereira, J. Acosta, and J. Assunção, "An implementation of the 7 wonders board game for ai-based players," in *19th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames 2020)*, IEEE, 2020.