Solving Patience

Patience is a simple card game and is similar to Solitaire. The outcome of Patience is caused entirely by the arrangement of the pack of cards the game is played with. In theory, if a player can analyse the deck before playing they should be able to determine the score they will achieve. This research aims to analyse the process of playing Patience programmatically and whether the final score can be predicted given any shuffled pack of cards. A time series interpretation of the data is presented and further research topics suggested, such as the K-Mean clustering algorithm to detect features and trends in the underling time series distribution.

Contents

1. Introduction

Patience is a simple card game that is played by a single player. This game involves the player shuffling a pack of cards (without the Jokers). The objective of the game to obtain the lowest score possible by removing cards of the same suit. The outcome of the game is almost entirely determined by the ordering of these cards. Analysis of this game is therefore interesting as, theoretically, if a player knew the ordering of the cards they could determine the outcome of the game.

The game is played in the following way.

1. A pack of 52 cards are shuffled and placed face down in one pile, known as the main pile.

2. Four cards are then taken from the top of the pack and placed in four separate piles, known as bins, face up so the player can see their value.

3. For each suit that is repeated more than once, the lowest value is removed from the game. For example, if the two of hearts and the seven of hearts were observed, the two of hearts would be removed. Aces are considered high.

4. Repeat the above step until no matching suits are observed. If one of the four bins is empty and an Ace is on the top of one of the bins, the Ace can be moved to one of the empty bins.

5. Place four more cards onto the bins from the main pile on top of the cards present in the bin.

6. Repeat step 3 – 5 until there are no more left in the main pile.

7. Calculate the score for the game as the number of cards in the bins minus four.

The outcome of the game is almost entirely determined by the ordering of the cards. The only edge case is when Aces can be moved from one of the bins to an empty bin. To begin analysing the game pragmatically we need to come up with some notation to model the game.

1.1. Notation

Each card can be represented with a number (the value on the card) and a suit. For example the seven of hearts could be represented as,

$$S(C) = 7 \cdot H, \tag{1}$$

where the seven represents the card value and H the suit. The $S(C)$ notation represents the state of the card, C. More generally, a card can be represented as,

$$S(C) = a \cdot B, \tag{2}$$

With $a \in \{0, 1, ..., 12\}$ are the possible values of 13 cards counted from zero and $B \in \{H, D, C, S\}$ are the possible suits that a card state S can occupy.

As cards can only interact with each other if their suits match we can define the following relation for card states:

$$X \cdot Y = \begin{cases} 1 : X = Y \\ 0 : X \neq Y \end{cases}$$

Where $X, Y \in \{H, D, C, S\}$. We can then model a shuffled pack of cards as a sequence of these card states provided the state is not in the sequence already (as identical cards cannot appear twice). Keeping track of the Deck of cards D we can recursively add cards to the D provided $S(C) \notin D$. By maintaining the ordering as items are added to D and randomly selecting a, B this will result in a random ordering of cards.

1.2. Approach

We will proceed by creating a randomly shuffled Deck D and playing the game of patience. The ordering of the cards in the Deck will be recorded after each passage of play, with cards being removed, and with possible re-ordering of the Aces, from D without using Bins directly. This will allow us to record analyse the contents of the deck D after each turn. We have interpreted each game to be a random walk, with the walk being the player's total score as the game continues. Fig. 1 shows these random walks and the distribution of player's final scores.

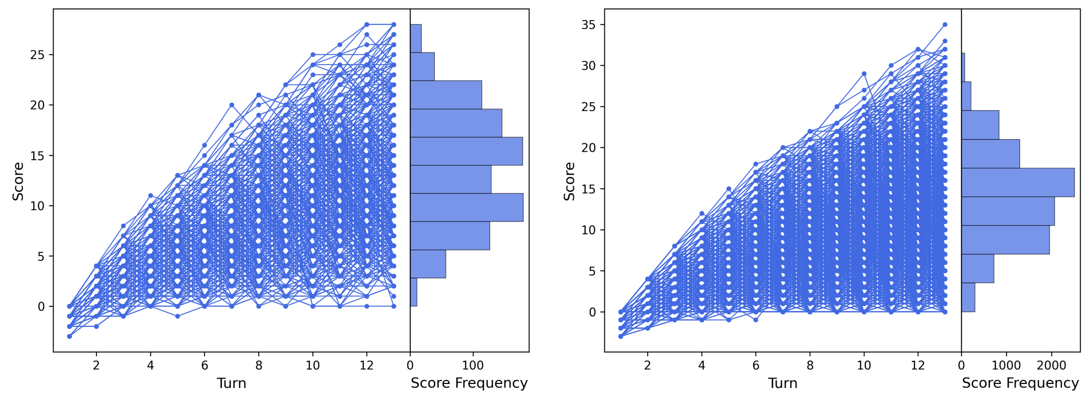

FIG. 1: Time series analysis showing the score at every turn of patience games. The lefthand plot shows 10^3 games and the righthand plot shows 10^4. The final scores have a mean value of 13.5 with a standard deviation of 5.6 with minimal variation between the two cases. Neither of the distributions are normally distributed. However, players of patience can anticipate to get a score in the region of 8 to 19 around 67% of the time. A perfect score of zero was only observed 2 and 34 for the left and righthand cases respectively. At these rates players can expect to get a perfect score roughly 0.2 – 0.4% of the time.

2. Results

We anticipate the endgame scores to be normally distributed; however, it is unclear what the mean or the standard deviation will be. After writing a program to play patience and record the deck D configuration and the score after every turn we can review the resultant scores.

According to the Shapiro-Wilks test, the final distribution of scores was not normally distributed at the 0.01 confidence level. Fundamentally, this indicates that the deviation of the players score from one turn to the next did not proportionally go upward and downward by the same amount. Indeed this property seemed to be random on a surface review, with no clear indicator for these movements. The plot in Fig.1 clearly shows this behaviour, with scores jumping by a maximum of 4 points per turn and potentially decreasing rapidly on some turns, with the upper limit of this decrease being the number of cards on the table (minus four).

The normalised residuals of the plots are not normally distributed. This is primarily because a player's score can only increase by a maximum of four but can decrease to a minimum score of -4 depending on the stage of the game. This disrupts the underlying distribution preventing a normal distribution. Features of this distribution could be analysed with classification algorthims, such as K-mean clustering.

3. Conclusions

Whilst Patience is a relatively simple problem to solve the techniques required to understand and analyse the game is not simplistic. Machine or deep learning could provide a means to analyse relationships between properties of the shuffled deck and the likelihood of obtaining a particular score. However, it is unclear which techniques should be used to analyse the information. Future research could begin by trying to understand relationships between the shuffled deck of cards and the users' final score with a unsupervised learning models. This will help with visualisation and initial analysis of the problem. This will require features of the shuffled deck to be extracted from each game and normalised between the lowest and highest possible values. A simplistic linear normalisation could be used in each case or a more complex log uniform normalisation. Unsupervised machine learning models should be applied, such as the K-Mean Clustering algorithm, to begin analysing relationships in the data. More complex algorithms, such as Markov chains, could provide insight in the time-series evolution of the Patience score.