Intelligent Agent for Playing Casino Card Games

Sanchit Goyal
Department of Computer Science
University of North Dakota
Grand Forks, ND 58202
sanchitgoyal01@gmail.com

Eric Kjeldergaärd
Department of Computer Science
University of North Dakota
Grand Forks, ND 58202
kjelderg@gmail.com

Ruchitha Deshmukh
Department of Computer Science
University of North Dakota
Grand Forks, ND 58202
ruchitha.deshmukh@und.edu

Eunjin Kim
Department of Computer Science
University of North Dakota
Grand Forks, ND 58202
ejkim@aero.und.edu

Abstract

Casinos are a place for people to participate in irrational and risky behaviour. Optimizing the winnings is an interesting exploration of probability and decision making. This paper delves into the world of intelligent gambling agents.

In particular, this paper lays out a method of creating an intelligent gambling agent capable of playing blackjack using learning, utility theory and decision-making to maximise the expected probability of winning the game. This paper will explain how this learning and decision-making is applied and how the probability of winning the game is increased, ultimately increasing the expected utility of playing blackjack.

The intelligent agent described in this paper is capable of perfectly determining the game probabilities based on the perception of its environment. Based on the input, the agent is able to make optimal decisions in both play and betting.

1 Introduction

Casinos earn their money by playing the odds. Having a set of games with probabilities tipped slightly in favour of the casinos multiplied by the number of games played makes the casino game an excellent example of applied probabilities. While there is clearly a significant factor of "luck" involved in any of the common casino games, there is also an element of "skill."

Maximizing the winnings from a casino is every gambler's ideal. If there is such a thing, an intelligent gambler would be one that gains all reasonable understanding of the game instance being played and makes decisions based on that understanding to maximise the expected payout.

It seems reasonable, then, to say that an intelligent agent would be able to process the game environment and determine an action that would maximize the expected utility of the game instance. What are the factors involved in such an agent? How does one create such an agent and how specifically does the agent achieve its goal?

This paper provides a sufficient background to understand the game of blackjack even for those unfamiliar with the game. This background focuses on the mathematical underpinnings of the game. The paper then describes a rational agent that combines a number of techniques in attempt to provide a near-optimal blackjack agent.

2 Prior Work

Casinos are a significant economic driver in modern civilization. Because casinos are economically significant, there is great incentive to study them. This incentive has lead to a wealth of studies and papers on the topic. The focus of these papers has developed a number of techniques and statistical background for players of casino games.

The surveyed studies and papers have a significant focus on either the mathematical basis for the rules and probabilities of the game based on random deals or are focused on practical applications which allow humans to more easily succeed at the games. The former category of work proves very useful to the work of creating a set of rules driving the actions of a rational agent. The latter is often focused on making the storage of information easier (card counting) which is not of much use in the design of a rational agent which can have perfect recollection. In other cases, the central theme is some helpful tips that reflect underlying statistical principles or the human nature of other gamblers, both of which are invaluable in the design of a gambling rational agent.

3 On Blackjack

Before putting forth an intelligent agent that is capable of playing blackjack, it is imperative to have a solid understanding of exactly what is this blackjack. Blackjack is a popular casino gambling game that uses standard decks of 52 cards in four suits of even size. Suits are not relevant to the play of blackjack, so it is safe to say that each deck has four each of the thirteen values for these purposes. The values are as follows: Ace; face values 2 through ten; face cards jack, queen, and king. Because there are many variants of the blackjack rules, the exact rules considered in this paper are presented here in some detail.

At the beginning of each round, two cards are dealt to each player and to the dealer. Each player then has their turn, determining if they would like another card added to their hand or if they are content with their current hand. The goal of each hand is to approach or achieve the summation 21 with the value of their hand. Going over the value 21 causes a player to "bust". If the player wishes another card, she "hits" and if content, she "stands".

There is one complication to the card play of blackjack. If a player is dealt two cards of the same value, excepting the values 4, 5, and 10, the player has the option to "split". If the player chooses to "split", the player then divides the two cards into two hands and receives one new card on each resultant hand. Each resultant hand is played as if it is a fresh hand. A player may only split four times within one round.

After all players have played out their hands, the dealer then plays out her hand until meeting or exceeding the value 17. If the dealer exceeds 21, the dealer "busts". Each player's hand is compared against the dealer's hand. If a player "busted", that player loses. If not, but the dealer "busted", then that player wins. If neither "bust", then the three cases remaining, dealer's hand is higher, hands are the same value, and player's hand is higher result in player loses, "push" where neither side profits, and player wins. If the player wins and had a value of 21 consisting of only 2 cards, the player has a "blackjack".

Betting is what drives the players and represents a utility function for the game. Before each round, the player determines the amount the player wishes to bet on that round. At the end of each round, a player with a losing hand forfeits their bet. A player with a winning hand wins the amount of their bet again unless the player has a blackjack for which the player wins 1.5 times the amount of their bet. If the player "split", then the original bet was risked on each resultant hand.

There is one additional complication to the betting of blackjack. For any hand where a player is dealt a sum of 9, 10, or 11 in their first two cards, the player has the option to double their bet for the hand, an act known as "doubling down". A player may "double down" on any hand in which the first two cards are of the correct sum including the resultant hands of "splitting."

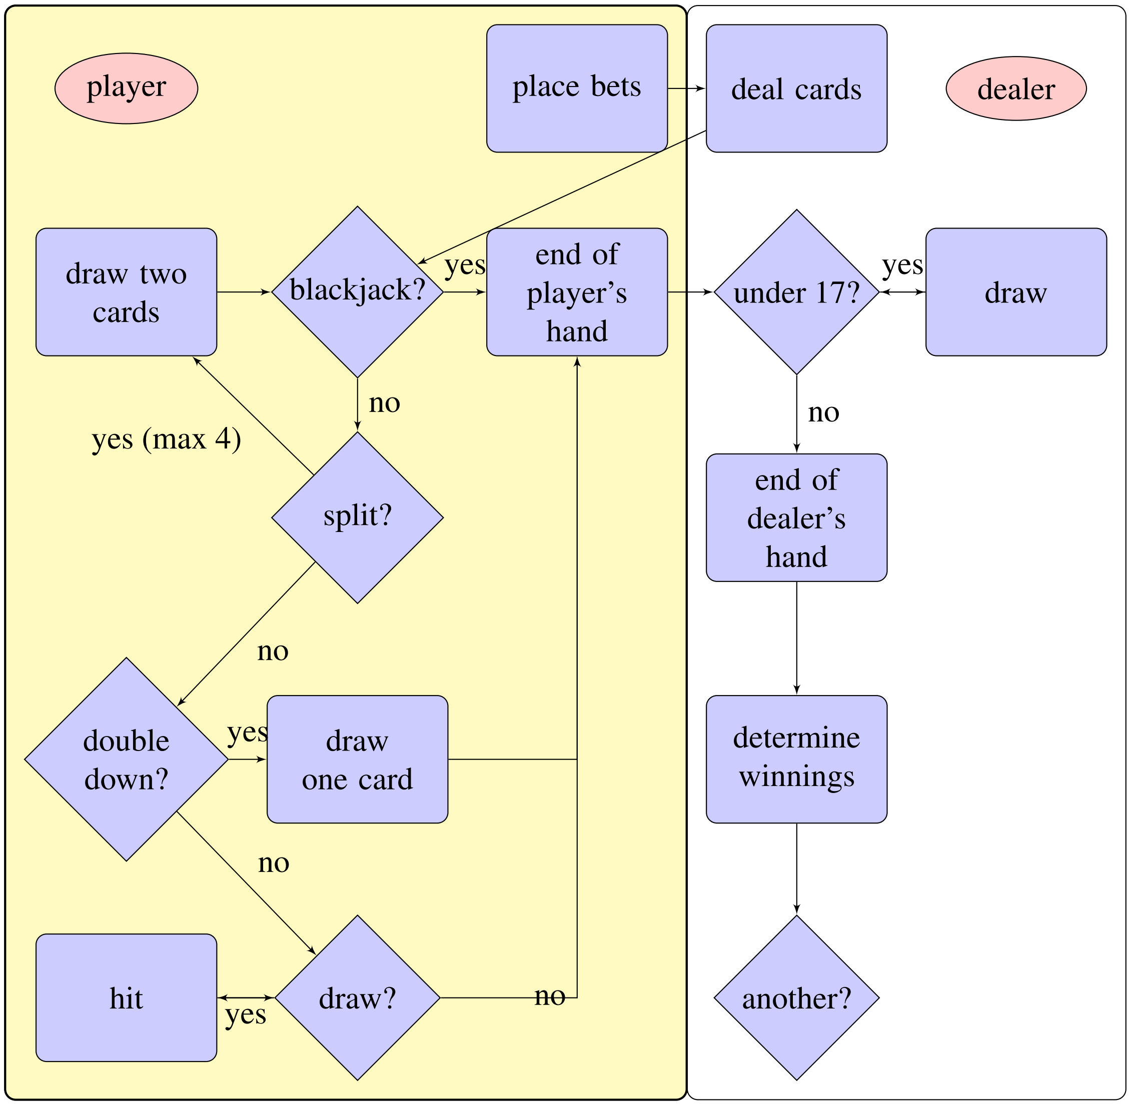

Figure 1: A flow chart describing the rules of blackjack.

4 On Utility

An ideal blackjack player would maximise the payout of each hand. The approach to doing this in a rational agent is to make choices that maximise the expected utility of that hand. The utility of each hand considering a bet b, is $-b$, 0, b, or $\frac{3b}{2}$ with no splitting or doubling down taken into account.

The casino would not offer blackjack if it were not profitable for the casino to do so. The odds, therefore, must be tipped in favour of the dealer. Some amount of work has been done to determine exactly what this advantage is [2, 4]. The values in Table 1 represent a survey [4] on this very topic.

Blackjack Strategy	Player's Advantage (in per cent)
Thorpe	-0.64
Acting like the Dealer	-5.7
Never bust	-6.0
Typical Casino Player	-2.0 to -15.0

Table 1: Players' advantage across surveyed blackjack strategies

Table 1 shows that with none of the proposed theories give the player a positive advantage over the casino. The particular advantage, however, can be raised significantly based on the particular technique used. This very observation indicates that a good enough player with a sufficient technique can perhaps tip the odds towards the player and away from the casino.

5 On Hitting and Standing

Implementing an intelligent agent for playing blackjack requires at least an agent that can make intelligent decisions during the game play. As described above in the description of blackjack rules, the agent has during the play, the choices of "hit", "stand", "split", and "double." This yields the first few requirements of an agent beyond its basic understanding of rules. The agent will need a decision network that takes in the current game state and output which of those four possible actions should be taken. Previous works on the subject have focused primarily on this subset of what an intelligent blackjack agent must do.

Deciding which action to take is based on the game state as it is known by the agent. The game state was traditionally thought to be the set of cards on the table and much work has been done to determine based solely on the cards showing what is the best possible action. Many more modern theories including work by Thorpe and others have added substantial work in the area of card-counting [7, 8].

Card-counting is the creation of a database consisting of the cards that have been exposed. Using this information, the probabilities of getting a particular card can be determined with much more confidence when compared to only looking at the current set of cards that have been dealt. The nearly perfect nature of agent memory makes this one area in which an intelligent agent can truly excel.

6 On Betting

In addition to determining how to best play the game, there is one other thing that a blackjack player has under her control. The first action a blackjack player takes in a game is to bet. The utility of the hand is a function of the bet. If a player were confident that she would win, then the player would bet everything. Similarly, if the player were confident that she would lose, the player would bet nothing. In practise, the confidence of the player is incomplete and thus the bets tend nearer nothing than the player's entire assets. Deciding the exact amount to bet on each hand has great impact on the amount a player will win or lose over time.

6.1 The Kelly Criterion

Perhaps the most prominent formula in the area of betting utility is the Kelly Criterion [1, 3, 7]. This formula provides a relationship between the probability of winning and the amount that can be "safely" bet on the game. That is to say, following the Kelly Criterion will yield an optimal rate of return over time for players with a logarithmic utility function on money as the one described in [5]. The equation,

$$f^* = \frac{bp - q}{b} \tag{1}$$

shows an optimal fraction of the bankroll, f^*, in terms of the odds of winning, p, and losing, q, taking into account b, the payout percentage for wins. This works for bets with a known probability and simple payout, but it gets a bit more complex for blackjack. With blackjack, there are two distinct types of payout, $b \in \{1, 1.5\}$. Each type of payout corresponding to its own probability. To express this complication, a new equation,

$$f^* = \frac{1p_{sw} - q_{sw}}{1} + \frac{1.5p_{bj} - q_{bj}}{1.5} \tag{2}$$

is created. This equation sums the fraction of the bankroll that should be wagered for a straight win, p_{sw}, and what should be wagered on a blackjack win, p_{bj}. For each probability of win, there is a corresponding probability of loss, q_{sw} and q_{bj}.

One may wonder, "What ever happened to the push and how does it fit in?" This has not been forgotten about, but has been omitted since the value of the payout of push is 0 and it adds nothing to the evaluation of the Kelly Criterion.

7 Agent Implementation

The blackjack playing agent is a probabilistic utility-based reasoning system. As such, it fundamentally senses its environment by observing the cards that are dealt and performing some action to maximise the expected utility of the environment. This simplistic view hides the underlying complexities and finer points of the agent. While the sensing of the environment is left to the realms of an API and the interaction with the environment comes in the form of printing information to a screen, the determining of the action provide plenty of room for interesting insight into the creation of a rational agent. Figure 2 provides an overview of the complete agent.

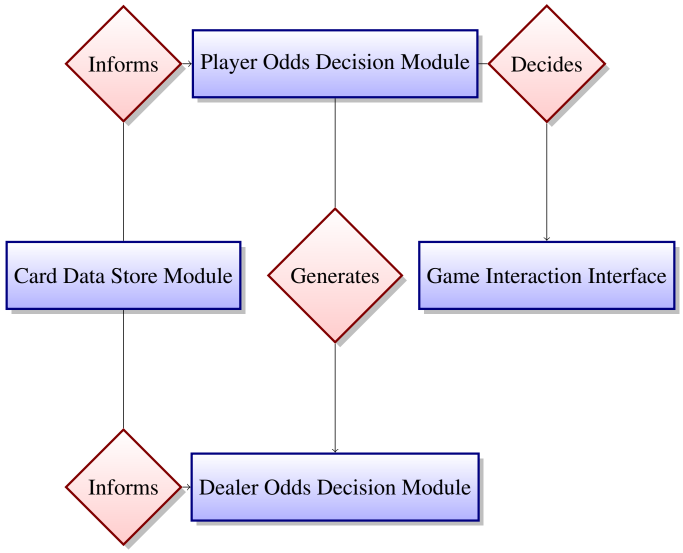

Figure 2: Entity Relationship Diagram of the Agent

A number of components are involved in this model and an overview of these components will help ease the understanding of the overall agent design. The agent is driven by a number of dynamically generated decision trees fuelled by a dynamically updating data store that provides the perceived odds of the available cards.

7.1 Probability Calculations

As time passes in the game, it becomes increasingly advantageous for a black jack player to keep track of probability of drawing of each card present in the deck. There are a number of published techniques [6, 10] that exist to aide a human player in approximating an accurate card count. However an artificial agent has an advantage over a human player in both speed and accuracy of retaining a count of the cards. The method used by the agent is to maintain an array of probability values of the denominations used in the game of blackjack.

The agent initialises its card count based on the number of decks being used at the table as described in function 1.

These figures can be used to calculate probability values across the denomination list after every card is dealt. Now these calculated and stored probabilities of each denomination can

function 1 CalculateInitial(N)

Input: N {Number of decks in play}

Output: C, the array with initial count of each card

```
for i ← 1 to 9 do
    C[i] ← 4 * N
end for
C[10] ← 4 * 4 * N
```

be used to make relatively complex calculations of important events in the game, such as $P(dealerbust)$, $P(playerbust)$, and so forth.

$$P(dealerbust) = P(dealertotal > 21) = P(10, 10, 2) + P(10, 10, 3) + \ldots$$

Figure 3: Mathematical expression of probability of dealer having a bust hand

In Figure 3, $P(x, y, z)$ is the probability of dealer's hand with cards of denominations x,y, and z. However $P(dealerbust)$ is a more important value than $P(x, y, z)$ as far as our agent is concerned. It will directly assist the agent in the calculation of $P(win)$ at every time step in the game. This $P(win)$ value in turn will be used in the Kelly Criterion equations to determine the amount of bet placed at every step in the game where money is required to be involved. Figure 4 shows the formula that is used to calculate $P(win)$.

$$\begin{aligned} P(win) &= P(dealerbust)P(\neg playerbust) \\ &+ P(\neg dealerbust)P(\neg playerbust)(P(playertotal > dealertotal) \\ &+ P(playerblackjack)P(\neg dealerblackjack)) \end{aligned}$$

Figure 4: Mathematical expression of probability of player winning

The exact probabilities are calculated using a decision tree approach. The following sections explain the procedures used in some depth.

7.2 Computing Player Odds

At any point in the game, the player has a large but finite set of possible hands, each hand having a certain probability of occurring. Each of these hands is represented by a decision-tree similar to that described in [5]. The decision tree has one decision node for each possible card value that could be dealt at a given time.

With the exception of being dealt a blackjack, each hand also represents a set of decisions. Only a subset of the possible hands are hands that the agent would choose to play. It

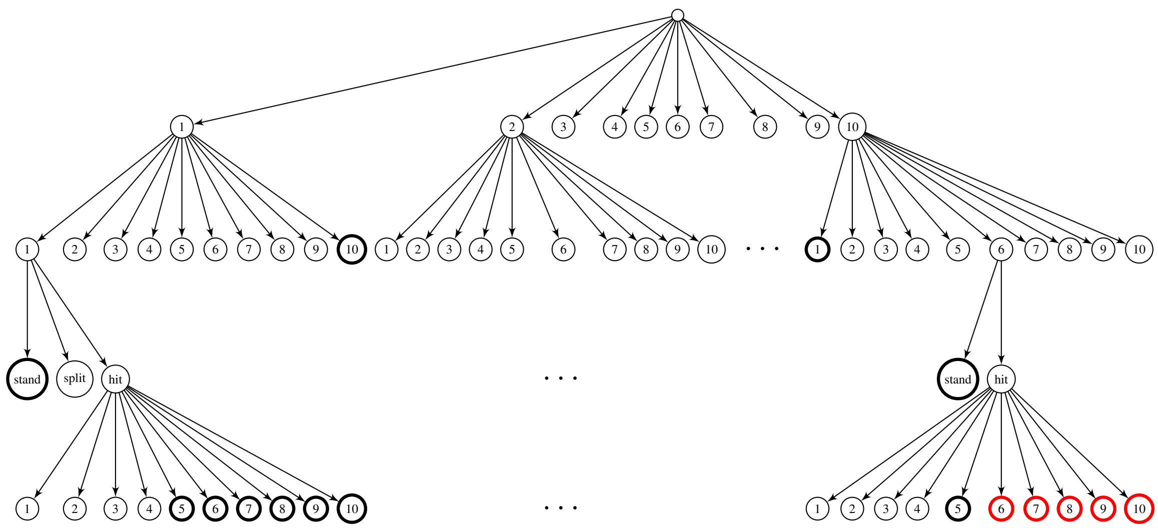

Figure 5: The Computation of Player Odds

is only the hands that the agent would choose to play that should be counted towards the probability of winning. Computing the actual probability of winning is therefore dependent on how the player plays the game. Knowing what action to take at any particular point in the game depends on the complete game state including the dealer's "up" card. How can the agent determine the odds of winning before any cards are dealt?

7.3 Computing Dealer Odds

When determining the odds of blackjack for the agent, there are only two players of note. Those two players are the agent itself and the dealer. The dealer as described follows a very simple set of rules. To assess the overall odds of the game before the play starts, it is important to understand the odds of both the dealer and the agent. Computing the odds of the dealer can be done quickly by enumerating the set of possible hands. This enumeration is done in a fixed-policy decision tree based on the house rules for the dealer.

7.4 Combining Player and Dealer Odds

Combining the two decision trees into one very large decision tree yields a method of determining the probability of winning given any amount of known state, from a fresh shoe with no cards dealt to a table where the agent has already made any number of decisions. For each possible player hand in the player's decision tree, the dealer could have had any possible card up, each with a certain trivially computed probability. For each possible card the dealer could have up, the dealer could have any possible second card and so forth, taking

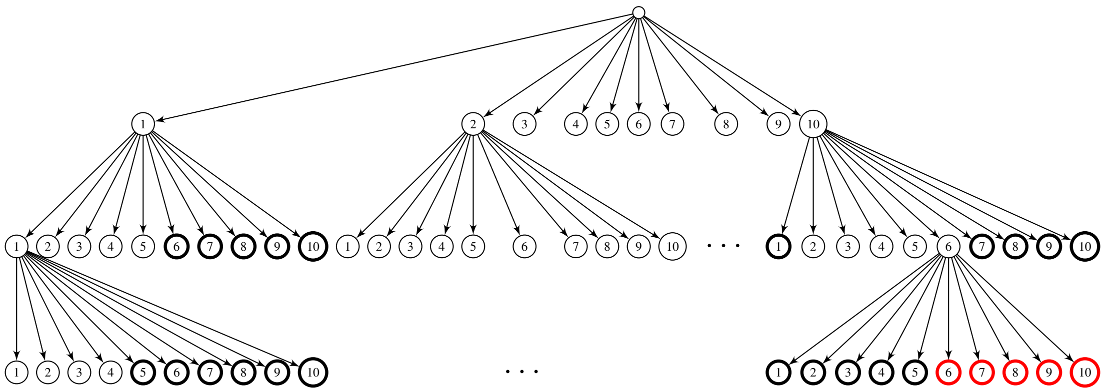

Figure 6: The Computation of Dealer Odds

into account the fixed policy of the dealer, the entire set of hands with their probabilities can be computed as described above.

With this "double tree" computed, each combination of hands has a particular result: win, lose, push, or blackjack. Finally, armed with this information the probability of winning for each user action is computed. For each decision node, the decision that corresponds to the highest expected utility is chosen over the decisions with lower expected utility. The decisions that are not chosen are discounted when computing the overall odds of winning.

7.5 Placing a Bet

Play begins with a player placing a bet. Similarly, the first action that the agent must perform is to place its optimal bet. As discussed above, the Kelly Criterion provides a formula for determining the optimal bet given known odds. The decision trees used by the agent for determining the probability of winning. Using this new-found knowledge allows the agent to compute the portion of the bankroll that should be bet for the round.

7.6 Card Counting

In order to make the optimal decision, the agent must store the data necessary to determine the most accurate probabilities possible. One element in doing so is storing which cards the agent believes are still available. This is done in an array with one element per card value with each element storing the number of cards remaining of that value. The agent also stores the total number of cards remaining. From these, the agent determines the probability of drawing any particular card.

In the initial state of the agent, coming newly to a freshly shuffled table, determining these probabilities is simply a matter of knowing the number of decks in play. Each particular card is assumed to be equally likely to be in any particular place within the shoe. As the play progresses, the agent decreases the count remaining for each card as the agent "sees" them – a technique commonly knows as card counting.

7.7 Playing a Hand

At the beginning of a hand, the dealer deals two cards to each player and, effectively, one to herself. The same tree concept that drives determining the winning odds before the hand is reused to determine which action to take given a current game state. The first two levels of the tree are reduced from broad branches including all possible cards at each level to simply the two cards that are displayed. The dealer's card is both removed from the array of available cards and used as the root node for all dealer sub-trees, again reducing the overall computation substantially while giving a much greater degree of accuracy between the expected utility and the actual utility of the hand.

As during the phase where the agent thinned out the possible hands into hands that the agent would choose, the agent computes the probabilities for each decision given the current state and discounts decisions with a lower expected utility than the optimal decision. Note that the actual decision may or may not reflect a decision that remained in the original tree when estimating the odds of winning the hand.

While the above makes clear how the agent chooses to "hit" or "stand", there are occasionally other decisions available to the agent. When presented with the opportunities, the agent must also decide whether to "split" or "double down". Each of these requires an additional bet to be laid down. Determining whether this bet is worthwhile, and thus whether the action is advantageous is a matter of evaluating the Kelly Criterion given the current context and seeing if the amount determined is greater than or equal to the current bet.

The play of each hand takes the form of a Bayesian network, but it is the combination of hands over time in a dynamic Bayesian network that vastly increases the agent's knowledge of the card state thus enabling the agent to make informed and intelligent decisions. In Figure 7 the relationship amongst the bankroll, B, over time, τ. Using the Kelly Criterion, the agent will determine a bet, b, for each hand based on the (knowledge of) the actual card ratios, c, and the bankroll. Once bets are placed, the agent's hand, h and dealers hand, d drive the agent's actions, a. These actions result in a total for each of the agent, H, and the dealer, D. The comparison of the two totals determines the winner, w. The outcome of the game is a combination of the bet and the winner.

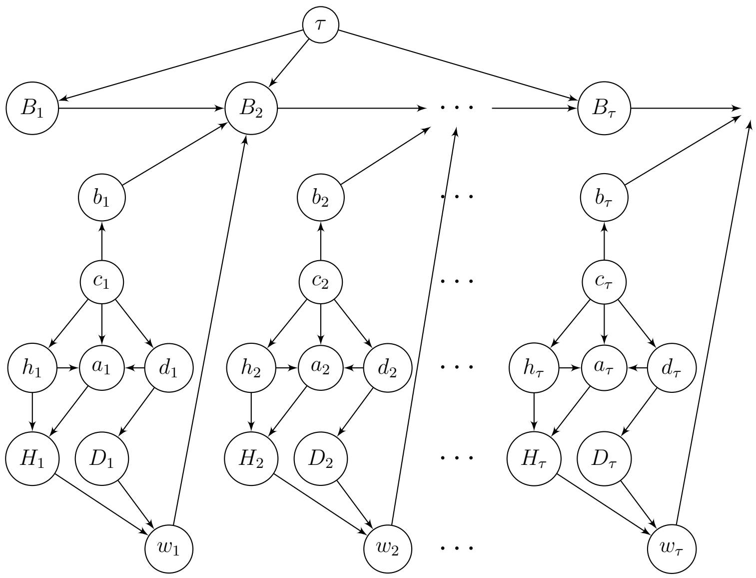

Figure 7: Dynamic Bayesian Network of Agent

8 Inherent Inefficiency of a Blackjack Agent Solution

The game of black jack is inherently based on the Subset Sum problem from the well-known NP Complete class of problems of the computational complexity theory. The subset sum problem is a decision problem which tries to answer, "given a set of integers and an integer s, does any non-empty subset sum to s?"

A blackjack agent typically faces a bigger problem based on the subset sum at the beginning of every card deal in a blackjack game which can be stated as follows.

1. Given the total of the agent's hand and the cards in dealer's hand including the 2 hidden cards, is there any set of cards within the remaining deck that will make my hand total greater than 21? If such set of cards are there, what is the probability of getting one such set in addition to the hand in possession.

2. While taking care of point 1, are there any set of cards in the remaining deck which will make my total $= 21$ or at least total $>$ Dealer's Stand Point. (Dealer's Stand Point $\in \{17, 18, 19\}$ depending on the casino rules and practices).

A game of blackjack revolves around this dilemma faced by a player. Based on the solution to the question posed above, an agent can derive the probability of winning a game and thus control its moves and the flow of money in the game. However due to the NP complete nature of the problem, there is no known efficient solution to this problem.

Number of Decks	Our Design		Literature	
	Dealer Busts	Dealer Blackjack	Dealer Busts	Dealer Blackjack
1	0.2835824	0.048265465	—-	—-
2	0.28257743	0.047796864	0.28263	0.4803
3	0.28225237	0.04764268	—-	—-
4	0.28208637	0.047565963	0.28209	0.4757
5	0.2819857	0.04752005	—-	—-
6	0.28192252	0.04748949	0.28192	0.4749
7	0.2818731	0.047467686	—-	—-
8	0.281835	0.047451347	—-	—-
9	0.28181112	0.047438648	—-	—-
10	0.2817892	0.04742849	—-	—-

Table 2: Agent calculated dealer odds vs. literature

9 System Testing and Experimentation

The decision calculates the odds in favour of the dealer at any time step in the game. The validity was verified using a number of simulations to find the probabilities like the probability of the dealer busting and the probability of the dealer getting a blackjack. The implementation predicts the odds of these events eventually happening in game, given the 1st card is being drawn out on the table. The probabilities thus obtained are listed in the Table 2 side-by-side with computed values discovered in a literature on blackjack probabilities.

The accuracy in the results of the simulation increases with increasing number of decks being shuffled in the shoe. This likely implies that the agent should play better on the tables with a larger shoe size. However, one more important inference that that can be drawn from the result obtained from these simulations is that as the number of decks used in the game increases, the probabilities of dealer getting a blackjack and the dealer getting busted decrease too. This clearly indicates to following two deductions.

1. The losses of the house due to dealer busting would decrease with the number of decks involved in the game.
2. The importance of an agent's skill increases with the number of decks involved, since he busts more but also starts getting more blackjacks on the other hand.

Just as the probability of the dealer getting a blackjack decreases, so decreases the probability of the player getting a blackjack. The probability of the player busting, however, is not so easily determined as it is dependent upon the particular actions of the player for each possible hand.

10 Summary

This paper has described the problem of blackjack and the design of an intelligent agent that is capable of playing blackjack in an optimal way. This optimal way combines the concepts of optimal blackjack playing as well as optimal betting. The agent uses decision tree learning to determine the optimal way to play a hand. It uses the same method of determining optimal actions to also calculate aggregate probabilities of winning and uses this information in the placing of a bet based on its utility function.

11 Future Work

The game of blackjack has many variants, each with a unique set of rules. This paper considers only one such set of rules. Therefore the agent design could be modified to include or exclude features based on the rules employed by a specific group of casinos.

This paper presents a very general design approach towards an intelligent agent which aims for an overall monetary profit in a long streak of games employing the Kelly criterion, instead of simply taking a game by game approach. However the Kelly criterion in itself is imperfect as a utility function for money due to the assumptions on which it is based. A wide area of improvement in the design can be in improving the Kelly criterion for the anomalies or in finding a more perfect utility function.

A more efficient implementation for the agent design presented in this paper is also an area of possible future work. The current implementation develops very large decision trees in order to consider all possible probabilities at each stage of play.

References

[1] S. Browne and W. Whitt, *Portfolio Choice and the Bayesian Kelly Criterion*. Advances in Applied Probability, Vol. 28, No. 4, 1996.

[2] G. Gottlieb, *An Analytic Derivation of Blackjack Win Rates*. Operations Research, Vol. 33, No. 5, 1985.

[3] J. L. Kelly, Jr., *A New Interpretation of Information Rate*. 1956.

[4] G. Kendall and C. Smith, *The Evolution of Blackjack Strategies*. 2003.

[5] Russell and Norvig,*Artificial Intelligence: A Modern Approach, 2nd ed.*. Prentice-Hall, 2003.

[6] E. O. Thorp, *Beat the Dealer: A Winning Strategy for the Game of Twenty-One*. New York: Vintage Books, 1962.

[7] E. O. Thorp, *The Kelly Criterion in Blackjack Sports Betting, and the Stock Market*. 10th International Conference on Gambling and Risk Taking, Montreal, June 1997.

[8] E. O. Thorp, *The Mathematics of Gambling*. 1984.

[9] E. O. Thorp, *Optimal Gambling Systems for Favorable Games*. Review of the International Statistical Institute, Vol. 37, No. 3, 1969.

[10] S. Wong, *Professional Blackjack*. Pi Yee Press, 1994.

[11] K. Yu and H. Wang, *Economic Analysis of Blackjack: An Application of Prospect Theory*. Canadian Economic Association 43rd Annual Conference, May 2931, 2009.