A Game-Theoretic Intelligent Agent for the Board Game Football Strategy

Sean McCulloch
Ohio Wesleyan University
stmccull@owu.edu

Abstract

While much work has been done in the development of intelligent agents for "classic" board games (e.g., Chess, Checkers, Othello), there have been many games developed in recent years that have not received as much attention. These games usually include large amounts of randomness or non-public information and often have interesting underlying theoretical structure. In this paper, we investigate a game called *Football Strategy*, which can be viewed as having similar structure to a normal-form game. We discuss the game-theoretic techniques used to arrive at a mixed strategy to play the game. We also discuss the methods used to incorporate information about the game which are not easily represented in game-theoretic ways into our agent. Our agent was evaluated by playing against one of the world's top players, and gave him a competitive game.

.

Introduction

For decades, the Artificial Intelligence community has worked on the development of intelligent agents to play "classic" games like Chess, Checkers, and Othello. In recent years, many new games have been published that have gained widespread popularity. Often, these games have two common factors that make them interesting to study from an intelligent agent perspective:

- They often have a large state space, or branching factor. Sometimes this is due to randomness (cards or dice), sometimes this is due to non-public information (different players each having secret victory conditions, for example), and sometimes this is due to a large variety of potential moves available to the player. Whatever the cause, traditional tree-based searching is often impractical.
- They often are based on interesting underlying theoretical or mathematical ideas which can potentially be exploited by the agent.

There have been some efforts to design agents to play these games. The most common approach is to manage the large state space created by these games, either by using clever variants of classical search methods (Heyden 2009) (Schadd et al 2007) (Schadd and Winands 2009) or by applying multiple agents to the problem (Johansson 2006) or by keeping a simpler model of the game to reduce the complexity (Thomas 2003). Another approach is to develop a rule-based system derived from specific rules and strategies of a given game. These rules can then either be applied directly to choose a move, or serve as the basis of an evaluation function in a search (Hall et al 2004) (Heyden 2009).

These approaches obviously have some inherent limitations. Often, attempts to reduce a game's complexity result in removing components that were necessary to play the game well. Additionally, a rule-based system is only as good as the rules it contains. If there is no rule to apply to a specific situation (for example, to react to a new strategy by an opponent), then the quality of the agent will suffer.

Our approach is to design an agent that exploits the underlying mathematical nature of a game. In this way, the agent's performance can be based on well-known rules of math, and the existence of mathematical theorems and tools can be used to aid the agent. In this way, our agent's play can approach "optimal" (in a mathematical sense).

In this paper we concentrate on a game called *Football Strategy*, originally published by Strategy Game Company in 1958, but more widely published by Avalon Hill in several editions starting in 1962. (For details on the game, see BGG 2015a.) The game represents an American football game between two generic teams who each call plays that are resolved in a two-dimensional matrix. The offensive player chooses a number representing a column, and the defensive player selects a letter representing a row. A small version of the matrix is included in Figure 1. In the actual chart, there are 20 offensive plays and 10 defensive plays, for a total of 200 entries.

	1 Power Up Middle	2 Power Off Tackle	3 QB Keeper	4 Slant Run
A	-2	-1	0	0
B	-1	FUMBLE	-1	+2
C	+10	-2	+15	-3
D	+1	+4	+2	+3

Figure 1: A subsection of a Football Strategy chart

In each play of the game, the offensive player and defensive player each secretly and simultaneously choose a letter or number. The intersection of those choices in the chart determines the result of the play. In most cases, the chart gives the gain (or loss) of yardage as the result of the play. For example, the combination of offense play “1” and defense play “D” (we will abbreviate such a combination as “1D”) is a gain of one yard. In some cases, other results can happen-- penalties, or as in the case of play 2B, a turnover, giving the other team the ball. The normal rules of American football apply (for scoring, first downs, et cetera.)

While some combinations of offensive and defensive plays lead to the same outcome (for example, plays 4B and 3D above both result in a gain of two yards), the chart has over 60 unique outcomes, leading to a branching factor that is too large to be addressed with traditional methods. By way of comparison, Chess has an average branching factor of about 35.

Game Theoretic Approach

In general, the chart given above can be seen as a zero-sum normal form game, in the game-theoretic sense. In these games, a *mixed strategy* is a probability that each player will choose one of each available options. A *Nash Equilibrium* (Nash 1950) in this case is a mixed strategy that cannot improve any player’s utility, given the other player’s strategy. Von Neuman has shown (von Neumann and Morgenstern 1944) that a mixed strategy Nash Equilibrium exists for all zero-sum normal form games. The mixed strategy for each player can be computed using linear programming (for example, see Mendelson 2004.)

We have developed an agent that uses these techniques to play the game of *Football Strategy*. The game chart gives the basic utility function measured in yards gained by the offensive player. The agent then modifies the values in the chart by adding (or subtracting) values to each table entry based on how the play affects the overall game. The new weighted yardage table is then solved using linear programming to generate a mixed strategy.

The resulting mixed strategy will give the agent a probability of choosing each offensive (or defensive) play that will maximize (or minimize) the expected performance of the offensive player. Table 1 shows a sample mixed strategy for an offensive player, and Table 2 a sample mixed strategy for the defensive player. Both tables show the same situation (first and ten, on the offense’s twenty yard line.) Low-numbered offensive plays tend to be running plays, and high numbers tend to be passing plays. Low-lettered defensive plays tend to be defenses against runs, and high-lettered defensive plays tend to be defenses against passes. In this case, the intersection of the highest probability offensive play call and the highest probability defensive call is 4E, which would result in a gain of five yards for the offense. However, most of the other defensive play calls that have positive probability (A,B,C,D,and H) give lower gains against the offensive play call of 4, sometimes even resulting in losses of yardage for the offense.

Offensive play number	**Probability of Play Call**
1	0%
2	.29%
3	12.9%
4	53.92%
5	12.3%
6	0%
7	0%
8	0%
9	2.04%
10	0%
11	0%
12	0%
13	10.21%
14	0%
15	0%
16	0%
17	5.98%
18	2.23%
19	0%
20	0%

Table 1: An example of an offensive mixed strategy

It is interesting to note that there are several play calls that have zero probability-- they will never be called. This result indicates that the solver has determined that for the given game situation (defined by a combination of yard line, down, distance to a first down, score, and time remaining in the game), those play calls are dominated by others. It is the case that each offensive and defensive play has a game situation in which it will be called with a nonzero probability.

Defensive play letter	Probability of Play Call
A	.24%
B	3.41%
C	5.94%
D	7.72%
E	43.9%
F	10.64%
G	12.72%
H	15.89%
I	0%
J	0%

Table 2: An example of a defensive mixed strategy

The fact that the mixed strategy generated is a Nash Equilibrium means that this probability gives the optimal expected gain if we assume the opposing player is also using the Nash Equilibrium mixed strategy for their side. If the opponent deviates from the equilibrium strategy, the expected gain may improve.

Adapting the Game Matrix

Simply defining the utility of the game matrix in terms of yards gained from each play is not sufficient to play *Football Strategy*—or any simulation of a football game— intelligently. The agent must manage several additional aspects that relate to the rules of football:

- While on offense, the offensive player has four attempts to gain a total of ten yards (a "first down",) which will reset the counter of four attempts. Often, the offensive player just uses three of those attempts, using the final attempt to give up possession to the opponent further downfield ("punting".)
- As shown in play 2B above, some combinations of plays result in the offense immediately losing possession of the ball.
- The ultimate goal of the offensive player (and what the defensive player is trying to prevent) is not merely to gain yardage-- points are scored either by crossing the opponents goal line, or by attempting to kick a field goal (which in this game is resolved randomly, with increasing probability given to attempts made closer to the opponent's goal line)

As a result, the utility matrix in the game must be modified to encapsulate a more accurate utility value based on these situations. To do this, we added five constants to the utility value in several situations.

The first constant value is added when the offensive player makes a first down. To encourage the agent to choose plays that gain some yardage, rather than merely trying a low-probability attempt to get the entire ten yards in one play, we add a prorated portion of the first down bonus for each yard gained (and subtract a similar amount if yardage is lost). To further encourage shorter yardage gains, we give an additional bonus (defined as a fixed percentage of the overall first down bonus) for a play that results in a favorable second or third down situation.

The second constant gives a large addition to the utility score when the play results in a touchdown. Since scoring a touchdown is the main goal of the offense (and preventing one is the main goal of the defense), this bonus should be large enough to have a major impact on play calling, but not so large that it causes the agent to become imbalanced when only a few risky plays give any chance of scoring a touchdown on a single play.

Additionally, it is possible for the offense to kick a field goal. In practice, this is usually done when the offense is on fourth down, and relatively close to the other team's end zone. Field goal attempts in the game are resolved by a random die roll based on the location on the field that the attempt is made from. Scoring a touchdown gives the offense six or seven (usually seven, depending on a random die roll) points, while kicking a field goal gives the offense three points. To simulate this, the extra utility added to the chart on plays that result in a successful field goal is set to 3/7 the value gained by a touchdown.

The third constant gives a large penalty to the utility score when the play results in a fumble or interception. While balancing the values of all of the constants has been challenging, it has been especially challenging to balance the interplay between the bonus utility added by scoring a touchdown versus the negative utility added by creating a turnover. Since a play that is neither a turnover nor a touchdown is usually better for the offensive player (since they will typically have at least one more chance to make another offensive play), we have found that setting the turnover penalty higher than the touchdown bonus works well.

The final two constants are more rarely relevant. First, if the offense loses enough yardage to end the play in their own end zone, this results in a "safety," giving the other team two points and the ball. This can only arise in rare situations. Finally, we added a special constant for turning the ball over on fourth down. We found that if we applied the normal turnover penalty described above to the fourth down situation, the agent would never find punting to be a viable option, since even a remote chance at a first down is better than the guaranteed turnover that punting represents. As a result, giving a smaller penalty to the utility score on fourth down turnovers led to less radical decision-making by the agent.

Strategic Considerations

The mixed strategy, with the additional constants described above, works well in a tactical sense to find the optimal move in a single play of the game in isolation. However, to win the actual game, other strategic factors are necessary to add to the complexity of the decision. The fact that the offense has several attempts to reach a first down, the time remaining in the game, and the current score all impact the choices made by players, and an agent who does not adapt to the changing nature of the game will not be successful.

The two major strategic aspects we have added to the game involve the special case of fourth down, and adjusting the play call in reference to the current time remaining and score of the game. In keeping with the ideals of designing an agent based on mathematical principles, we have made an effort to allow the agent to create its own strategy by implementing these strategic factors as modifications to the game matrix.

Fourth Down

On fourth down, the offensive player has three choices: punt the ball to the other team (losing possession but forcing the other player to gain more yards to score), try to gain a first down (losing possession at the current yard line if the attempt fails), or try to kick a field goal (represented in this game as a random chance to score points based on how close the offense is to the end zone). There are two punting plays, one that can be called at any time, and one that can only be called in "punt formation," which the offense must declare before both players choose their plays. The offense is allowed to call a non-punt play in punt formation (a "fake punt,") but suffers a penalty on any yardage gained.

The agent determines the best option between these three choices by finding the expected utility of all three possible play call situations (calling a normal play, calling a play from punt formation, and attempting a field goal). The agent will choose the mixed strategy that has the highest utility. An interesting consequence of this method is that the agent will occasionally decide that going into punt formation and fake punting is the correct move, just like real human players would.

Clock Management

The game of football is timed and scored, and the ultimate goal is to have the most points when time runs out. As a result, teams that are losing towards the end of the game often need to take more risks to gain more yards quickly, and teams that are winning can play more conservatively and choose slower plays that consume more time. Teams get three time-outs each half which are used to cause the previous play to consume less time, allowing both teams the ability to conserve the clock, if desired.

To simulate this, once the time remaining gets below a certain threshold, the agent estimates the number of offensive plays the player currently losing the game will have, based on the time left in the game and the number of timeouts each player has. If the offensive player is currently losing, this is the number of plays left in the game. If the defensive player is currently losing, there is a minimum number of plays that must elapse before they can get possession of the ball and try to score.

Next, the agent calculates the number of yards that need to be gained by the losing player to have a chance at winning the game. If a field goal is sufficient to win, we calculate the yardage to the next position on the field that changes the odds of making a successful field goal; otherwise we calculate the yardage needed to score a touchdown.

From these values, we can generate an average "yards per play" that the player needs to make to reach the goal of winning the game. The elapsed time and yardage of each play considered by the agent will have an effect on this average. The weighted change to this average yards per play will be added to the play's utility value.

Evaluation

Since Football Strategy is not as widely played as other games, and to our knowledge there are no computer programs that play it, we evaluated the performance of the program against several humans of varying ability, ranging from novices to experts. After seeing encouraging results against these players, we also tested the agent in several games against Bruce Reiff, the world's top player of *Football Strategy*, as rated by the Boardgame Player's Association (BPA 2014) (BPA 2015).

Mr. Reiff played against our agent in a series of three games. He won all three games, but each was within a touchdown, and in one game, his victory required a successful field goal attempt on the last play of the game. Despite losing all three games, we feel that our agent can play at a level competitive with the world's top players. After our official experiment, Mr. Reiff played a few "recreational" games against our agent, and our agent did manage to get some wins against him.

As an additional experiment, we set up some games where Mr. Reiff was allowed to see the agent's mixed strategy probabilities before choosing a play, though not the agent's actual play call. Mathematically, if the utility values are computed correctly, this should not give the human player any advantage, since it is an equilibrium strategy. However, the extra strategic elements added to the game, as well as the effects of the constant values added to the utilities, made the game slightly exploitable, and the human player won by a wide margin.

Future Work

There are several possible avenues of future improvement that can be added to the agent.

We would like to generate the values of the five constants used to adjust the values placed in the game matrix in a better way. Right now we have just used trial-and-error in playing against humans. We would instead like to make these values more rooted in the underlying mathematical principles of the game.

We would like to integrate calling time-outs into the utility function, especially towards the end of each half of the game. Currently, time-outs are called outside of the game matrix decision making, using a set of pre-defined rules. Incorporating them into the utility matrix would allow for more strategic decision making by the agent.

We would like to explore ways in which the agent can plan a sequence of downs to realize short-term goals. For example, a “second and short” situation- where it is second down and very few yards are needed to obtain a first down- is often a good time to attempt a risky play like a long pass. If it is successful, there is a large gain of yardage. If it is unsuccessful, it is likely that the short yardage required for a first down can be gained on the next two plays. Currently, the agent as constructed merely sees a high probability for obtaining a first down and chooses that over the higher-potential play.

Humans tend to play *Football Strategy* by predicting tendencies of play calls made by the opponent. It would be interesting to explore whether a learning element could be added to the game to allow the agent to make similar decisions.

A similar game to *Football Strategy* exists, called *Paydirt* (BGG 2015b). This game uses a chart of play calls similar to the one found in *Football Strategy*, but the charts represent actual historical teams. The intersection of an offensive and defensive play call is resolved probabilistically based on the historical team’s abilities. The underlying ideas behind our current agent should translate well to a program playing this similar game.

We would like to extend this idea of exploiting the underlying mathematical structure of a game to other games. While there are not many games that are so directly related to normal-form games, many other games use concepts from probability, graph theory, or auctions that can be exploited in similar ways.

Acknowledgements

We would like to thank Julide Yazar for an invaluable education into the basics of Game Theory. We would also like to thank Bruce Reiff for being our resident *Football Strategy* expert and experimental subject.

References

BGG. 2015a. Football Strategy. Available online at http://boardgamegeek.com/boardgame/951/football-strategy

BGG 2015b. Paydirt. Available online at http://boardgamegeek.com/boardgame/1498/paydirt

BPA. 2014. Boardgame Players Association Football Strategy Rankings. Available online at http://boardgamers.org/yearbook14/fbspge.htm

BPA. 2015. Boardgame Players Association. Available online at http://www.boardgamers.org.

Hall L.; Gordon A.; James R.; and Newall L. 2004. A Lightweight Rule-Based AI Engine for Mobile Games. In Proceedings of the 2004 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology, 284-289. New York, NY: Association of Computing Machinery.

Heyden C. 2009. Implementing a Computer Player in Carcassone. Master’s Thesis. Department of Knowledge Engineering, Maastricht University, Maastricht, Netherlands.

Johnansson, S. 2006. On Using Multi-agent Systems in Playing Board Games. In Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, 569-576. New York, NY: Association of Computing Machinery

Mendelson, E. 2004. Introducing Game Theory and its Applications. Boca Raton FL:CRC Press.

Nash , J. 1950. Equilibrium Pointes in n-person Games. Proceedings of the National Academy of Sciences, 36(1):48-49.

Schadd, M,.; Winands M.; Bergsma M; and Uiterwijk J. 2007. Solving Fanorona. ICGA Journal, 30(1):43-45.

Schadd, M; and Winands M. 2009. Quiescence Search for Stratego. In Proceedings of the 21st Benelux Conference on Artificial Intelligence. Eindhoven, the Netherlands.

Thomas, R. 2003. Real-time Decision Making for Adversarial Environments Using a Plan-based Heuristic. PhD diss., Department of Computer Science, Northwestern University, Evanston IL.

von Neumann, J., and Morgenstern, O. 1944. Theory of Games and Economic Behavior. Princeton University Press.