Computing Strategies of American Football via Counterfactual Regret Minimization

Yuki Shimano, [1] **Kenshi Abe,** [2] **Atsushi Iwasaki,** [3] **Kazunori Ohkawara,** [3]

[1] Independent
[2] CyberAgent, Inc
[3] University of Electro-Communications
yuki.shimano0401@gmail.com, abe_kenshi@cyberagent.co.jp,
atsushi.iwasaki@uec.ac.jp, ohkawara_kazunori@e-one.uec.ac.jp

Abstract

In this research, we analyze the strategies of American football in the game-theoretical framework. The past decade has seen a tremendous growth of interest in sports analytics, not only from a commercial perspective but also from a scientific one and so on. It is currently used as an indicator for tactical decisions in various sports such as baseball, soccer, and American football. By contrast, most of the previous efforts are limited to decision support based on statistics and do not take into account situations in which one's actions might change the actions of others. Even if such situations are taken into account, there is also the issue that the action space is not very large. We thus formulate American football as a two-player zero-sum extensive-form game with imperfect information and compute the Nash equilibrium of this game. To compute it, we use Counterfactual Regret Minimization which computes approximate Nash equilibrium in extensive-form games with imperfect information. Our simulation suggests that the equilibrium computed in our settings are history-based strategies that are observed in the National Football League. Moreover, empirical results show that the following the computed equilibrium is superior to following the probability of choosing a play computed from the game data and it is possible to compute an equilibrium by expanding the action space, which makes it possible to compute a strategy that is less likely to be exploited.

1 Introduction

1.1 Sports Analytics

In the sports, we have determined the plays via experiment or intuition for several years. However, the analysis based on statistics has spread rapidly since the movie "Money Ball" has made SABRmetrics famous. Strategy analysis like this has been called Sports Analytics.

Over the past few years, it has become mainstream to analyze strategy using any technologies in the sports with the development of information processing, the increasing sophistication and ease of acquisition of game data. American football (AF), in particular, is one of the sports where the use of data is most advanced. In the National Football League (NFL), the trajectories of players are captured from RFIDs attached to each player, and the plays are evaluated by using machine learning and other techniques [1]. In addition, they are soliciting applications on Kaggle, a data analysis platform, for tactical evaluation methods [2]. As can be seen, there is a lot of enthusiasm for analysis using computer and data utilization.

1.2 Previous Works

As above, AF goes well with data science because it has a better data infrastructure than other sports; therefore, there have been several types of researches about analyzing and computing strategies in AF. (Yurko, Ventura, and Horowitz 2019) provide a framework called nflWAR using multilevel models to isolate the contributions of individual offensive skill players, and providing estimates for their individual Wins Above Replacement (WAR). (Yurko, Matano, and Richardson 2020) introduced a general framework for continuous-time within-play valuation using player-tracking data and used a Long Short Term Memory (LSTM) recurrent neural network to construct a ball-carrier model to estimate how many yards the ball-carrier gets. These researches we introduced are how to evaluate and analyze offensive strategies. On the other hand, we explain researches on how to compute offensive optimal strategies. (Jordan, Melouk, and Perry 2009) quantified the decisions what to play using game theoretic techniques; updating optimal decision policy as new information becomes available during a game. In these papers (Yee, Rodríguez, and Alvarado 2014; Yee, Campirán, and Alvarado 2016), the analysis of American football strategies what offensive and defensive players should choice play for their role was by applying Nash equilibrium. Furthermore, they propose AF modeling by means of a context-free grammar. Besides these researches, there were some projects about computing strategies by applying Nash equilibrium and Minmax strategy (AdvancedFootballAnalytics 2008; Kovash and Levitt 2009; SBNATION 2015; Adams 2020).

AF is the highest strategic sport that offense and defense need to anticipate what the other think. We thus have to analyze and compute strategies in American football in a game-theoretical framework. Previous works formulated AF into the normal-form game but there was nevertheless no work

[1]NFL Next Gen Stats. https://nextgenstats.nfl.com
[2]https://www.kaggle.com/c/nfl-big-data-bowl-2021

that AF is formulated by extensive-form game. AF is the game that offense and defense repeatedly play for the point, and so we have to formulate AF as the extensive-form game. Additionally, the player's action was only selected as Run or Pass in most of the previous studies. Each plays has its own characteristics in AF, so it is unrealistic to abstract plays down to two choices for analysis and estimation.

In this work, we formulate AF into two-player zero-sum extensive-form games with imperfect information expanding the action space rather than previous works, and compute approximate Nash equilibrium of this game via Counterfactual Regret Minimization. Additionally, we examine these results that are computed in several situations. Through these experiment, we confirmed that the equilibrium is a history-based strategy of what plays each other has played in the past and what yards they have got. These equilibria are the strategies observed in the real games such as the NFL. Moreover, empirical results show that the following the computed equilibrium is superior to following the probability of choosing a play computed from the game data and it is possible to compute an equilibrium by expanding the action space, which makes it possible to compute a strategy that is less likely to be exploited. Our research shows that effective equilibrium can be computed in a more realistic setting than in previous efforts. Our research serves as a window to an understanding of the process how the player in AF determine the their strategies, and open up a new field in Game-AI research.

2 Notation and Background

2.1 Two-player Zero-Sum Extensive-form Game with Imperfect Information

A finite extensive-form game with imperfect information has the following components.

- A finite set $N = \{1, 2\}$ of players.
- A finite set H of sequences, the possible histories of actions, such that the empty sequence is in H and every prefix of a sequence in H is also in H.
- A finite set $Z \subset H$ of terminal histories.
- A player function P that assigns to each nonterminal history a member of $N \cup \{c\}$. If $P(h) = c$ then chance determines the action taken after history $h \in H$.
- $A(h)$ are the actions available after a nonterminal history $h \in H$.
- A function σ_c that associates with every history h for which $P(h) = c$ a probability measure $\sigma_c(\cdot \mid h)$ on $A(h)$, where each such probability measure is independent of every other such measure.
- For each player $i \in N$ a partition $\mathcal{I}_i$ of $\{h \in H \colon P(h) = i\}$ with the property that $A(h) = A(h')$ whenever h and h' are in the same member of the partition. For $I_i \in \mathcal{I}_i$ we denote by $A(I_i)$ the set $A(h)$ and by $P(I_i)$ the player $P(h)$ for any $h \in I_i$. $\mathcal{I}_i$ is the information partition of player i; a set $I_i \in \mathcal{I}_i$ is an information set of player i.
- For each player $i \in N$ a utility function u_i. Since the game in this paper is two-player zero-sum, $u_1 = -u_2$.

2.2 Strategies and Nash equilibrium

A strategy for player i, σ_i, in an extensive game is a function that assigns a distribution over $A(I_i)$ to each $I_i \in \mathcal{I}_i$, and σ_i is the set of strategies for player i. A strategy profile σ consists of a strategy for each player, σ_1, σ_2, with σ_{-i} referring to all the strategies in σ except σ_i. A best response is the optimal strategy for player i to use against opponent profile σ_{-i}. It is defined as

$$b_i(\sigma_{-i}) = \underset{\sigma_i' \in \Sigma_i}{\operatorname{argmax}} u_i(\sigma_i', \sigma_{-1}).$$

Two-player zero-sum games have a game value, V_i, that is the lower bound on the utility of an optimal player in position i. In this case, we use the term exploitability (Johanson et al. 2011) to refer to the difference

$$\epsilon_i(\sigma_i) = V_i - u_i(\sigma_i, b_{-i}(\sigma_i)).$$

If an agent plays according to σ then its exploitability is

$$\epsilon(\sigma) = u_1(b_1(\sigma_2), \sigma_2) + u_2(\sigma_1, b_2(\sigma_1)).$$

The traditional solution concept of a two-player extensive-form game is that of a Nash equilibrium. A Nash equilibrium is a strategy profile σ where

$$u_i(\sigma_i^*,\ \sigma_{-i}^*) \geq \max_{\sigma_i \in \Sigma_i} u_i(\sigma_i^{'},\ \sigma_{-i}^*) \quad \forall i \in N.$$

An approximation of a Nash equilibrium or ϵ-Nash equilibrium is a strategy profile σ where

$$u_i(\sigma_i^*,\ \sigma_{-i}^*) + \epsilon \geq \max_{\sigma_i \in \Sigma_i} u_i(\sigma_i^{'},\ \sigma_{-i}^*) \quad \forall i \in N.$$

In two-player zero-sum games, Nash equilibrium is unexploitable, $\epsilon(\sigma) = 0$. Moreover, it is important to compute Nash equilibrium properly because it has property minimizing utility loss. In particular, we could get utility more than $u_i(\sigma_i^*, \sigma_{-i}^*)$ in two-player zero-sum game no matter what opponent use strategies by using Nash equilibrium.

2.3 Counterfactual Regret Minimization

Counterfactual Regret Minimization (CFR) is an iterative algorithm that converges to a Nash equilibrium in any finite two-player zero-sum extensive-form game with a theoretical convergence bound of $O(\frac{1}{\sqrt{T}})$ (Zinkevich et al. 2007). Let σ^t be the strategy profile on iteration t. The counterfactual value $v_i(\sigma^t, I)$ of player $i = P(I)$ at I on t is the expected payoff to i when reaching I, weighted by the probability that i would reached I if she tried to do so that iteration. Formally,

$$v_i(\sigma^t, I) = \sum_{h \in I} \pi_{-i}^{\sigma^t}(h) \sum_{z \in Z_h} \pi^{\sigma^t}(h, z) u_i(z)$$

and $v_i(\sigma_{I \to a}^t, I)$ is the same except it assumes that player i plays action a at I on t with 100% probability. The instantaneous regret $r^t(I, a)$ is the difference between $P(I)$'s counterfactual value from playing a and playing σ on iteration t. The counterfactual regret for I action a on iteration T

is $R^T(I,a) = \sum_{t=1}^{T} r^t(I,a)$. Additionally, $R_i^{T,+}(I,a) = \max(R_i^T(I,a),\ 0)$. Total regret for i in the entire game is

$$R_i^T = \max_{\sigma_i'} \sum_{t=1}^{T} \left(u_i(\sigma_i', \sigma_{-p}^t) - u_i(\sigma_i^t, \sigma_{-p}^t)\right).$$

In regret matching (Hart and Mas-Colell 2000), a player picks a distribution over actions in an infoset in proportion to the positive regret on those actions. On each iteration $T+1$, i selects actions $a \in A(I)$ according to probabilities

$$\sigma_i^{T+1}(I,a) = \frac{R_i^{T,+}(I,a)}{\sum_{a \in A(I)} R_i^{T,+}(I,a)}.$$

Moreover, define $\overline{\sigma}_i^T(I,a)$ to be the average strategy for player i from iteration 1 to T. In particular, for each infoset I, each a, define

$$\overline{\sigma}_i^T(I,a) = \frac{\sum_{t=1}^{T} \pi_i^{\sigma^t}(I)\sigma_i^t(I,a)}{\sum_{t=1}^{T} \pi_i^{\sigma^t}(I)}.$$

In two-player zero-sum game at time T, if both player's average total regret satisfies $\frac{R_i^T}{T} \leq \epsilon$, then $\overline{\sigma}^T$ is a 2ϵ equilibrium (Zinkevich et al. 2007). Thus, CFR constitutes an anytime algorithm for finding an ϵ-Nash equilibrium in two-player zero-sum game.

3 Rules of American Football

3.1 What is American Football?

AF is a team sport played offense and defense teams. The offense have four downs (1st, 2nd, 3rd, and 4th Down) and must advance at least 10 yd in four downs. If they fail, they turn over the football to the defense, but if they succeed, they are given a new set of four downs to continue the drive. The goal of the offense is to get the point by advancing the ball into the opposing team's end zone. This is one of the main characteristics of AF, where the offense is trying to figure out how to move the ball forward, while the defense is trying to figure out how to stop the offense. Furthermore, unlike rugby or soccer, AF is characterized by a clear division of attacks and clear delimitation of plays.

3.2 Settings

In this research, we will construct a game restricted to the following settings, since formulating all of AF would be computationally prohibitive.

- Formulated as a two-player game with offense and defense. The individual roles are not taken into account in our case.
- We will cut out a particular 1st to 3rd Down in AF and analyze only that sequence. In other words, the offense and defense play at most three times.
- The game is over if the offense advances over 10 yd.
- The current location on the field and the distance to the touchdown are not taken into account.

As mentioned earlier, the offense has four downs, but in general, the offense often gives up the play on 4th Down to recover the position, so we defined the game until 3rd Down.

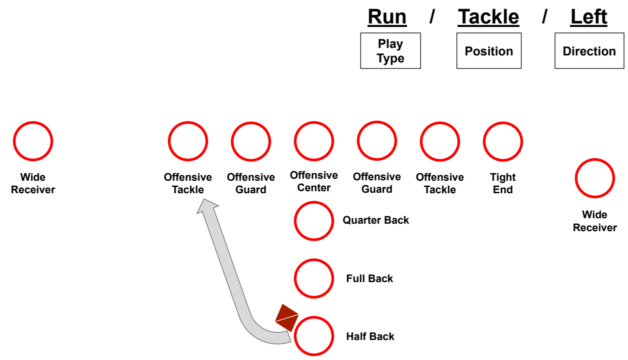

Figure 1: Offensive action example $(Run|Tackle|Left)$

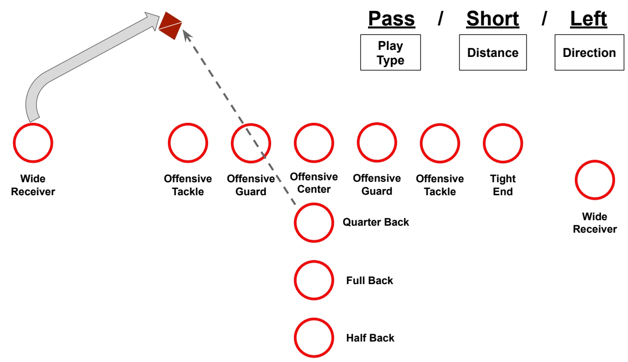

Figure 2: Offensive action example $(Pass|Short|Left)$

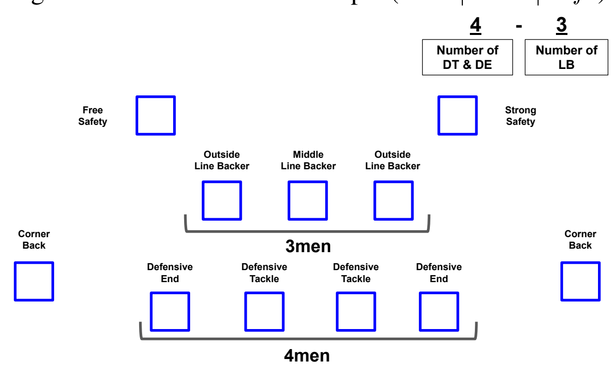

Figure 3: Defensive action example $(4-3)$

3.3 Player's Action

In our study, we extend the action space from the previous works and formulate it. Offensive actions are roughly classified into two, Run and $Pass$. Run is subdivided by $Position$ and $Direction$. On the other hand, $Pass$ is subdivided by $Distance$ and $Direction$. From the above, the offense can choose from the following 13 types of actions. Note that when the direction of the Run is $Middle$, there is no need to define the position to aim for, so N/A (Not applicable) is used.

$$\{Run\} \times \{N/A\} \times \{Middle\}$$
$$\{Run\} \times \{Guard,\ Tackle,\ End\} \times \{Right,\ Left\}$$
$$\{Pass\} \times \{Short,\ Deep\} \times \{Middle,\ Right,\ Left\}.$$

Table 1: Pay-off matrix (excerpt)

	$3-4$	
$Pass\|Short\|Left$	$(u_{\text{off}}(h'),\ u_{\text{def}}(h'))$ $(u_{\text{off}}(h''),\ u_{\text{def}}(h''))$ $(1,\ -1)$	if got 0 yd if got 6 yd if got 13 yd

Figure 1 and 2 show the example of offensive Run action and $Pass$ action. $Run|Tackle|Left$ means that offense play Run aiming for $Left$ $Tackle$. Meanwhile, $Pass|Short|Left$ means that offense play $Short$ range $Pass$ on the $Left$. Run doesn't basically have as many yards to gain, but it has a higher probability of success. By contrast, $Pass$ has more yards to gain, but the probability of success is lower than Run.

Defensive actions are classified into three formations,

$$\{4-2,\ 4-3,\ 3-4\}.$$

Figure 3 displays the example of defensive action $4-3$. The defensive formation are roughly composed of three layers: the forward, the middle and the rear. The formations are divided by how many players are placed in the forward and the middle. $4-3$ means that a formation with 4 $players$ in the forward and 3 $players$ in the middle. In general, the more people placed on the scrimmage line, the stronger against the Run. If the total number of forward and middle player is small, the number of rear player increase and thus the action will be stronger against $Pass$.

3.4 Game Structure and Utility

We show conceptual scheme of AF game in Figure 4. In this way, we can represent AF as a game tree. The game starts with a 1st Down, and the total yards required for the offense to get a positive utility is 10 yd. The pair of plays selected by each player determined the getting yards of offense in each Down. If the total number of yards got exceeds 10 yd, the offense gets utility 1 and the defense gets utility -1. By contrast, if the total number of yards got does not exceed 10 yd at the end of 3rd Down, the offense gets utility -1 and the defense gets utility 1. We then show detail scheme of AF game in Figure 5. Figure 5 displays that after the offense makes a decision, the defense makes a decision. However, since offense and defense select plays at the same time in each Down in the real game, the decision points of defense belong to the same information set. The getting yards is determined for a pair of plays, but since it is generally not uniquely determined, it is determined according to a probability distribution.

Table 1 displays the pay-off matrix when the offense chooses $Pass|Short|Left$ and the defense chooses $3-4$ on 1st Down. In this example, the offense is assumed to have some probability of gaining either 0 yd, 6 yd, or 13 yd. If the offense gets 13 yd, the game is over, and then the offense gets utility 1 and the defense gets utility -1. Otherwise, the game continues as the offense doesn't advance 10 yd. No utility can be gained since the game is not over. Therefore, this pay-off matrix assigns the expected pay-off u_{off}, u_{def} that

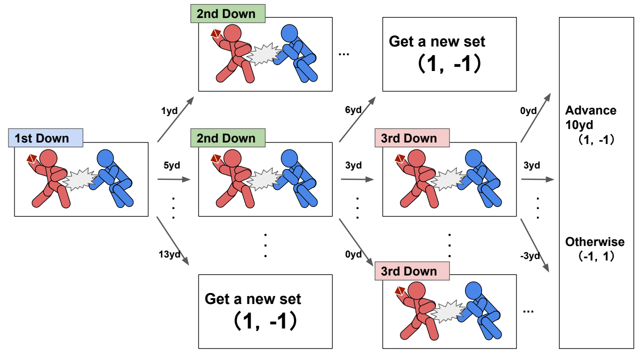

Figure 4: Conceptual scheme of AF

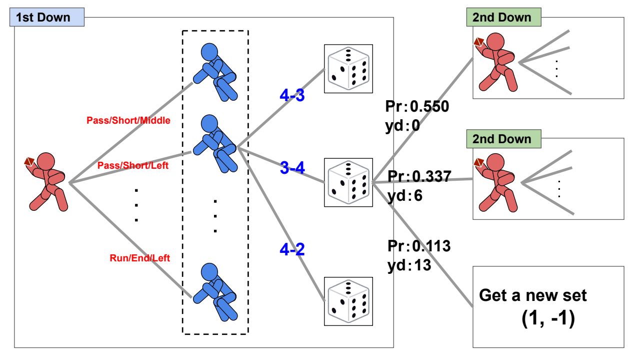

Figure 5: Detail scheme of AF

will be obtained after the 2nd Down. Note that h' and h'' denote the histories after selecting plays from each other from the history h at the start of Down and gaining yd from the result. 1st Down 10 yd is the starting point of the game, so $h = \varnothing$, and each history can be described by

$$h' = \{Pass|Short|Left,\ 3-4,\ 0\ yd\}$$
$$h'' = \{Pass|Short|Left,\ 3-4,\ 6\ yd\}.$$

Note that $\{Pass|Short|Left,\ 3-4,\ 13\ yd\} \in Z$.

3.5 Getting Yards via NFL Game Data

In general, the getting yards for a given pair of plays is not uniquely determined, and there is variability. Therefore, the getting yards for a given pair of plays is determined according to σ_c computed from the 2017 NFL game data. Figure 6 shows the tally of yards gained against the defensive play from the game data when the offense play $Pass|Short|Left$. In this case, getting yard x corresponding to the following are treated as outliers. Note that Q_1 is first quartile, Q_3 is third quartile, and IQR is Interquartile Range.

$$x < Q_1 - 1.5(IQR), \quad Q_3 + 1.5(IQR) < x.$$

Then, create a histogram using the outlier-processed data, the mode in each bin is the getting yards, and the area of each bin is the probability of selecting that yard. We show the histogram at $bin = 5$ when the offense plays $Pass|Short|Left$ and the defense plays $3-4$ in Figure 7.

According to Figure 7, the offense would have a probability of 0 yd at 0.432, 6 yd at 0.260, 8 yd at 0.136, 12 yd at 0.124, and 18 yd at 0.047. The larger the size of bin, the closer σ_c is to the actual distribution of the getting yard. Note that σ_c for the case where the getting yard is 6 yd and the selection probability is 0.260 can be written as follows

$$\sigma_c(6\ yd) = 0.260.$$

4 Experimental Setup and Result

In this section, we compute the equilibrium via CFR with our settings and examine the results. In this calculation, iteration $T = 10,000$, $bin \in \{3, 4, 5\}$. Table 4 displays some of the offensive strategies that were computed as a result of the calculations [3]. These strategies are shown to be those that we would use in a real game, and are history-based strategies. For instance, the equilibrium at 2nd Down 2 yd with history $\{Pass|Short|Middle,\ 4-3,\ 8\ yd\}$ is to choose $Run|Tackle|Left$ at 0.570, $Run|Tackle|Right$ at 0.419, and $Run|Guard|Right$ at 0.011. In general, when there are only a few yards left to get a new set, the offense tends to choose Run, which has a high probability of getting a few yards, and rarely choose Pass. The above confirms that it is possible to compute the equilibrium strategies in our setting that is similar to that observed in the real games such as the National Football League.

Moreover, since the equilibrium we have computed is an approximate Nash equilibrium, we calculate the exploitability to evaluate how close the equilibrium is to the Nash equilibrium. Figure 8 shows transition of exploitability. $bin = 3$ setting reaches the exploitability of 2.03×10^{-3}, $bin = 4$ one reaches the exploitability of 1.98×10^{-3}, and $bin = 5$ one reaches the exploitability of 1.42×10^{-3}. We could denote them the approximate Nash equilibrium from these previous works (Lanctot et al. 2009; Brown and Sandholm 2019a; Brown et al. 2019).

In this experiment, the exploitability in the final iteration is smaller when the size of bin is large than when the size of bin is small. We expect that this is because the larger size of bin, the more detailed the plays can be described, and thus the fewer on-path nodes. Figure 9 displays the transition of average strategy in $bin = 3$, and Figure 10 displays the transition of average strategy in $bin = 5$. Both histories are $h = \{Pass|Short|Left,\ 4-3,\ 8\,yd,\ Run|Tackle|Right,\ 4-2,\ 0\,yd\}$. At the $t = 100$, there are 6 types of average strategies greater than 0 in $bin = 3$, while there are only 3 in $bin = 5$. Thus, it can be seen that the larger size of bin is, the fewer on-path nodes there are.

5 Discussion

This section walks through various results and analysis of equilibrium computed with our model.

5.1 Stats vs Equilibrium

We denote how effective playing the strategies computed are. We test this by pitting players who choose tactics according to equilibrium strategies against players who choose them according to statistical data. Players who follow the statistical data shall select plays according to the probability calculated for each $Current\ Down$ and $Remaining\ yards\ for\ getting\ new\ set$ from the 2017 NFL game data. The experimental setup is shown below.

- Both player play 1st to 3rd Down.
- If the offense gets 10 yards or more in three plays, the offense wins. On the other hand, if the offense fails to get 10 yards, the defense wins.
- The histogram of getting yards at $bin = 3$.
- One simulation is run until one side wins.
- This simulation is performed 10,000 times, and they are considered as one set.
- Do this for 10 sets, and compute the average number of offensive and defensive wins.

Table 2 displays the number of offensive and defensive wins. For instance, the Table 2 shows that the offense win on average 5,226 and the defense won on average 4,774 when both offense and defense chose plays according to the statistical data. If the offense follows the statistical data, the winning rate is around 50%. By contrast, if the offense follows the Nash equilibrium, the winning rate is in the upper 70%. For the defense, following the Nash equilibrium also

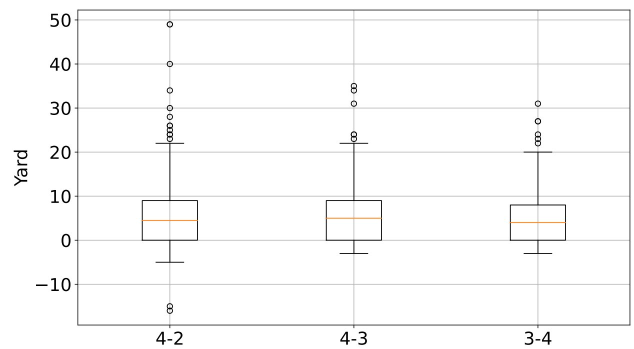

Figure 6: The tally of yards gained against the defensive play when the offense plays $Pass|Short|Left$

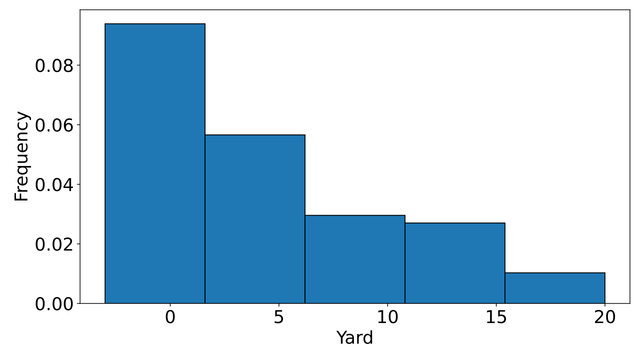

Figure 7: The histogram of getting yards at $bin = 5$ when the offense plays $Pass|Short|Left$ and the defense plays $3-4$

[3]For full results, see the following URL. https://www.dropbox.com/sh/wk0o3nmcbtkgv9m/AADEkoK_5tYjCnhpc-XNSkCVa?dl=0

Table 2: Number of offensive and defensive wins

		Defense	
		Stats	Equilibrium
Offense	Stats	(5,226, 4,774)	(4,778, 5,222)
	Equilibrium	(7,716, 2,284)	(7,676, 2,324)

Table 3: Exploitability average and SD

	exploitability	SD
large space: $bin = 3$	0.194	0.032
large space: $bin = 4$	0.195	0.035
large space: $bin = 5$	0.192	0.022
small space: $bin = 3$	0.286	0.030
small space: $bin = 4$	0.398	0.030
small space: $bin = 5$	0.285	0.031
small space: $bin = 20$	0.277	0.030

increases the winning rate. It shows that following the Nash equilibrium can increase the winning rate. In summary, we have shown the superiority of following the Nash equilibrium in determining plays.

5.2 Effect of the Action Space

The previous works compute equilibrium by limiting the choice to Run and $Pass$. On the other hand, in our study, we extend the number of actions by considering directions, formations and so on. Assuming that the setting of the previous work is a *small space* and the setting of the our work is a *large space*, we examine the effect of the size of the action space on equilibrium using exploitability.

In order to make a comparison, we first compute equilibrium in each action space based on the setting of our study. Note that $bin \in \{3, 4, 5, 20\}$ in the *small space* and $bin \in \{3, 4, 5\}$ in the *large space*. Since the games in the *large space* and *small space* are different, we apply the equilibrium computed in each to the common game (called *real game*) based on deterministic reverse mapping (Gilpin, Sandholm, and Sorensen 2008) and calculate the exploitability. Specifically, the equilibrium of the same or close to the history in the *real game* is directly applied. For example, the equilibrium in the history $h = [Pass|Short|Middle,\ 4-2,\ 8\ yd]$ in the *large space* implies mapping as equilibrium in the following history in the *real game*.

$$h' = [Pass|Short|Middle,\ 4-2,\ 8\ yd]$$
$$h'' = [Pass|Short|Middle,\ 4-2,\ 6\ yd]$$

The mapping in small game is almost the same, but we distribute the probabilities uniformly. For example, consider the case where the probability of choosing Run in a history is 0.7. The probability of choosing $Run|Tackle|Left$ in the corresponding *real game* history is $0.7/7 = 0.1$ because the number of tactics in the Run is 7. Note that, in both settings, if there is no corresponding equilibrium, the probability of choosing each action is $1/|A|$.

Table 3 displays the average and standard deviation (SD) of exploitability. For example, the results for "*large space*: $bin = 3$" in the table show that the offense is an equilibrium with "*large space*: $bin = 3$", and when the exploitability is computed for each of the defensive equilibrium computed in each setting, the mean is 0.194 and the standard deviation is 0.032. Since the exploitability in the *large space* is smaller than that in the *small space*, this experiment shows that it is possible to compute a strategy closer to Nash equilibrium by expanding the action space. From the previous experiment, it was found that determining plays according to Nash equilibrium increases the winning rate. At least in the setting of this experiment, we could state the advantage of computing equilibrium by expanding the action space.

On the other hand, if we focus on the size of bin, we can see that a large bin does not necessarily mean a small ex-

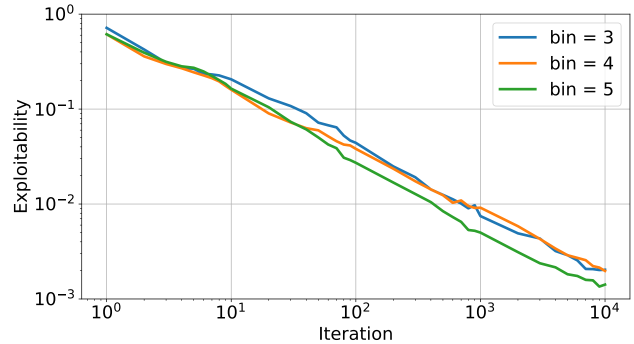

Figure 8: Transition of exploitability

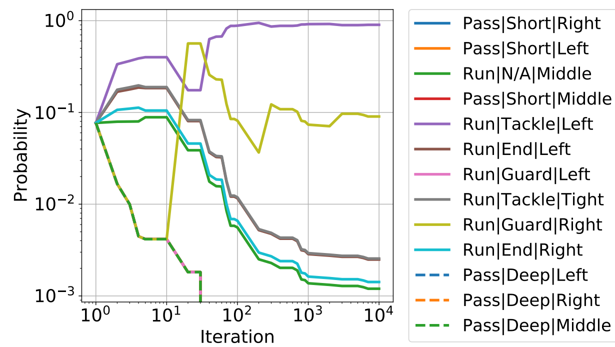

Figure 9: Transition of average strategy in $bin = 3$

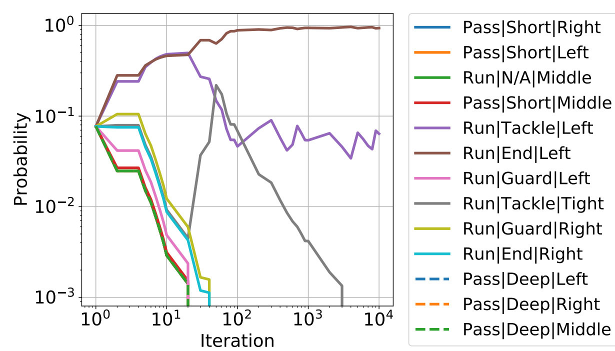

Figure 10: Transition of average strategy in $bin = 5$

Table 4: Offensive equilibrium (excerpt)

<table>
<tr><th>bin</th><th>Down</th><th>yd</th><th>History h</th><th>Equilibrium</th></tr>
<tr><td rowspan="5">3</td><td>1</td><td>10</td><td>$\varnothing$</td><td>[Pass|Short|Left, Pass|Short|Middle, Pass|Deep|Right]
= [0.001, 0.510, 0.489]</td></tr>
<tr><td rowspan="4">2</td><td rowspan="2">10</td><td>{Pass|Short|Left, 4-3, 0 yd}</td><td>[Pass|Short|Middle, Pass|Deep|Left, Pass|Deep|Right, Pass|Deep|Middle]
= [0.474, 0.002, 0.414, 0.110]</td></tr>
<tr><td>{Pass|Short|Middle, 4-2, 0 yd}</td><td>[Pass|Short|Middle, Pass|Deep|Right, Pass|Deep|Middle]
= [0.057, 0.298, 0.645]</td></tr>
<tr><td rowspan="2">2</td><td>{Pass|Short|Left, 4-3, 8 yd}</td><td>[Run|Tackle|Left, Run|End|Left, Run|Tackle|Right, Run|Guard|Right]
= [0.288, 0.005, 0.568, 0.139]</td></tr>
<tr><td>{Pass|Short|Middle, 4-3, 8 yd}</td><td>[Run|Tackle|Left, Run|Tackle|Right, Run|Guard|Right]
= [0.570, 0.419, 0.011]</td></tr>
<tr><td rowspan="5">4</td><td>1</td><td>10</td><td>$\varnothing$</td><td>[Pass|Short|Right, Pass|Short|Left, Pass|Short|Middle, Pass|Deep|Right]
= [0.024, 0.410, 0.413, 0.153]</td></tr>
<tr><td rowspan="4">2</td><td rowspan="2">10</td><td>{Pass|Short|Left, 4-3, 0 yd}</td><td>[Pass|Short|Middle, Pass|Short|Middle, Pass|Deep|Right]
= [0.117, 0.413, 0.470]</td></tr>
<tr><td>{Pass|Short|Middle, 4-2, 0 yd}</td><td>[Pass|Short|Right, Pass|Short|Left, Pass|Short|Middle, Pass|Deep|Right, Pass|Deep|Middle]
= [0.122, 0.004, 0.494, 0.370, 0.010]</td></tr>
<tr><td rowspan="2">2</td><td>{Pass|Short|Left, 4-3, 8 yd}</td><td>[Run/Tackle/Left, Run/Tackle/Right]
= [0.368, 0.632]</td></tr>
<tr><td>{Pass|Short|Middle, 4-2, 8 yd}</td><td>[Pass|Short|Right, Run|Tackle|Left, Run|Tackle|Right, Run|Guard|Right, Run|End|Right, Pass|Deep|Middle]
= [0.004, 0.396, 0.513, 0.037, 0.040, 0.010]</td></tr>
<tr><td rowspan="5">5</td><td>1</td><td>10</td><td>$\varnothing$</td><td>[Pass|Short|Left, Pass|Short|Middle]
= [0.607, 0393]</td></tr>
<tr><td rowspan="4">2</td><td rowspan="2">10</td><td>{Pass|Short|Left, 4-3, 0 yd}</td><td>[Pass|Short|Middle, Pass|Deep|Right, Pass|Deep|Middle]
= [0.401, 0.453, 0.146]</td></tr>
<tr><td>{Pass|Short|Middle, 4-2, 0 yd}</td><td>[Pass|Short|Middle, Pass|Deep|Right]
= [0.558, 0.442]</td></tr>
<tr><td rowspan="2">2</td><td>{Pass|Short|Left, 4-3, 8 yd}</td><td>[Run|Tackle|Left, Run|Tackle|Right, Run|Guard|Right]
= [0.803, 0.195, 0.002]</td></tr>
<tr><td>{Pass|Short|Middle, 4-2, 8 yd}</td><td>[Run|Tackle|Left, Run|End|Left, Run|Tackle|Right, Run|Guard|Right]
= [0.414, 0.351, 0.005, 0.230]</td></tr>
</table>

ploitability and SD. In general, the more elaborate the game structure, the more desirable it is for accurate prediction. A larger bin means that the distribution of got yards for a pair of plays is closer to the distribution in reality, and thus the game structure is more elaborate. Therefore, we expected both Exploitability and SD to decrease as bin increased, but this was not the case. The relationship between size of bin and equilibrium will be the subject of future work.

The mapping method in this experiment uses a simple technique. There is a possibility that the exploitability in the $small\ space$ can be further improved by devising a mapping method (Schnizlein, Bowling, and Szafron 2009). In other words, the present experiment does not allow us to state that it is $always$ superior to compute equilibrium in $action\ large$. It would be an ongoing task to dig deeper into the relationship between the action space and equilibrium.

6 Conclusion

In this paper, we formulate American football as a two-player zero-sum extensive-form game with imperfect information, and compute the approximate Nash equilibrium of this game via Counterfactual Regret Minimization. Empirical results show that the following the computed equilibrium is superior to following the probability of choosing a play computed from the game data. We also showed that it is possible to compute the strategy that is less likely to be exploited by expanding the action space, although this is a limited approach. This research shows that effective equilibrium can be computed in a more realistic setting than in previous efforts. Future work includes generalizing the problem setting and improving the efficiency of equilibrium computation by abstracting the game (Brown and Sandholm 2018, 2019b).

References

Adams, C. B. 2020. *Learning Microeconometrics with R*. Chapman and Hall/CRC.

AdvancedFootballAnalytics. 2008. Game Theory and Run/Pass Balance. http://archive.advancedfootballanalytics.com/2008/06/game-theory-and-runpass-balance.html. Accessed: 2020-03-14.

Brown, N.; Lerer, A.; Gross, S.; and Sandholm, T. 2019. Deep Counterfactual Regret Minimization. In *Proceedings of the 36th International Conference on Machine Learning*, 793–802.

Brown, N.; and Sandholm, T. 2018. Superhuman AI for Heads-up No-limit Poker: Libratus beats top professionals. *Science*, 359(6374): 418–424.

Brown, N.; and Sandholm, T. 2019a. Solving Imperfect-Information Games via Discounted Regret Minimization. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, 1829–1836.

Brown, N.; and Sandholm, T. 2019b. Superhuman AI for Multiplayer Poker. *Science*, 365(6456): 885–890.

Gilpin, A.; Sandholm, T.; and Sorensen, T. B. 2008. A heads-up no-limit Texas Hold'em poker player: Discretized betting models and automatically generated equilibrium-finding programs. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems (AAMAS)*, 911–918.

Hart, S.; and Mas-Colell, A. 2000. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica*, 68: 1127–1150.

Johanson, M.; Waugh, K.; Bowling, M.; and Zinkevich, M. 2011. Accelerating Best Response Calculation in Large Extensive Games. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 258–265.

Jordan, J. D.; Melouk, S. H.; and Perry, M. B. 2009. Optimizing Football Game Play Calling. *Journal of Quantitative Analysis in Sports*, 5: 1–34.

Kovash, K.; and Levitt, S. D. 2009. Professionals Do Not Play Minimax: Evidence from Major League Baseball and the National Football League. NBER Working Papers 15347, National Bureau of Economic Research, Inc.

Lanctot, M.; Waugh, K.; Zinkevich, M.; and Bowling, M. 2009. Monte Carlo Sampling for Regret Minimization in Extensive Games. In *Proceedings of the 22nd International Conference on Neural Information Processing Systems*, 1078–1086.

SBNATION. 2015. Understanding balanced, effective play-calling with Game Theory and Nash's Equilibrium. https://www.fieldgulls.com/seahawks-analysis/2015/1/9/7517551/understanding-effective-play-calling-with-game-theory. Accessed: 2020-03-14.

Schnizlein, D.; Bowling, M.; and Szafron, D. 2009. Probabilistic State Translation in Extensive Games with Large Action Sets. In *Proceedings of the 21st international jont conference on Artifical intelligence*, 278–284.

Yee, A.; Campirán, E.; and Alvarado, M. 2016. Gaming and Strategic Choices to American Football. *International Journal of Mathematical and Computational Methods*, 1: 355–371.

Yee, A.; Rodríguez, R.; and Alvarado, M. 2014. Analysis of strategies in American Football using Nash Equilibrium. In *Proceeding of AIMSA*, 286–294.

Yurko, R.; Matano, F.; and Richardson, L. F. 2020. Going Deep: Models for Continuous-Time Within-Play Valuation of Fame Outcomes in American Football with Tracking Data. *Journal of Quantitative Analysis in Sports*, 16: 163–182.

Yurko, R.; Ventura, S.; and Horowitz, M. 2019. nflWAR: A Reproducible Method for Offensive Player Evaluation in Football. *Journal of Quantitative Analysis in Sports*, 15: 163–183.

Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2007. Regret Minimization in Games with Incomplete Information. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, 1729–1736.