Predictors of NFL Tackles

Isaac Kou (4502)
Isaac.Kou@colorado.edu
University of Colorado Boulder
Boulder, Colorado, USA

Sean Shi (4502)
Sean.Shi@colorado.edu
University of Colorado Boulder
Boulder, Colorado, USA

Timur Tripp (5502)
Timur.Tripp@colorado.edu
University of Colorado Boulder
Boulder, Colorado, USA

Abstract

This paper analyzes the 2024 National Football League (NFL) Big Data Bowl dataset which contains data from the 2022 season of the NFL. We show that the height, weight, and position of the player in the game have correlations, though ultimately non-significant, with whether or not the player was successfully tackled during that play. Additionally, strategies such as the formation used by defense and offense and the length of passes can increase or decrease the distance before tackle of a carrier or the likelihood of a tackle succeeding.

Keywords

Football, NFL, Data Mining, Tackles

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/XXXXXXX.XXXXXXX

ACM Reference Format:
Isaac Kou (4502), Sean Shi (4502), and Timur Tripp (5502). 2024. Predictors of NFL Tackles. In . ACM, New York, NY, USA, 10 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Every year, the Nation Football League (NFL) hosts the NFL Big Data Bowl, an annual sports data analytics competition in which teams analyze specific datasets provided by Next Gen Stats, the NFL's statistics division, in order to encourage data-backed improvements to the NFL. The data, which aligns with the competition's yearly theme, is

> *"the capture of real time location data, speed and acceleration for every player, every play on every inch of the field. Sensors throughout the stadium track tags placed on players' shoulder pads, charting individual movements within inches."* [1]

Specifically, the 2024 Big Data Bowl focuses on the theme of tackling in order to encourage the creation of metrics quantifying things like tackle probability, play impacts, expected points, and injury.

Tackling is imperative to the sport of American football, as the defending team's main goal is to tackle the opposing player with the football as soon as possible in order to help prevent the team from scoring. As such, the dataset used to analyze tackling includes categories on game data, play data, player data, tackles data, and tracking data. The goal of this project is to utilize such metrics in order to find correlations between known information such as weight, height, and time during a game and the outcome of plays (whether or not the tackle was completed, whether the tackle was successful in terms of gameplay, etc.).

2 Existing Work

A significant number of previous academic studies exist on the topic of tackling in American football. This section discusses three examples of those studies, one hobbyist analysis, and the foundation they lay for our proposed work.

The paper "Validating Tackle Mechanics in American Football: Improving Safety and Performance"[5] (Maerlender et al.) discusses the development of a program to reduce the risk of serious injury from tackling, specifically by identifying head-contact likelihoods and therefore alternative techniques to reduce such contact. Secondary research related to the impact of tackles on player performance was also conducted, finding a reduction in their number of yards run post-contact.

Similarly, "Effectiveness of the Heads Up Tackling (HUT) Program on Tackling Safety and Performance in American Football"[6] (Matsuo et al.) also identifies a program to promote tackle safety, though this study focuses on an existing program and its observed impacts. Results indicated that implementation of the studied HUT Program reduced rate and severity of player injury, with no detriment, or, in some cases, even benefit, to player performance.

Lastly, "Quantitative and qualitative analysis of head and body impacts in American 7v7 non-tackle football"[3] (Jadischke et al.) used video analysis to investigate the rate of head contact in the alternative non-tackle play style of American football. In a similar vein to the above studies, this research found that non-tackle football was associated with lower rates of head contact and therefore potentially lower rates of injury, though further research is needed.

Outside of academia, many football fans conduct hobbyist analyses of their own, often concerning points, player performance, game trends, and so on. An example of this is the article "Tackling Metrics with Big Data"[2] by Ezra Ford on Medium. In addition to being a subject of interest in and of themselves, these have applications in related sub-hobbies, such as sports betting and fantasy football, and can help fans maximize their success in those activities.

As shown, the existing work on tackles in American football often concerns the implications of tackling, instead of the ability to predict and understand when/where/how tackles happen. Thus, our work intends to investigate these trends instead.

3 Methodology

In this project, we seek to find the effects of different attributes on the success rate and time of tackles and subsequently, how tackles effects the outcome of a game. To this end, we will be employing a variety of statistical techniques covered in both CSCI 4502/5502 and from other classes.

3.1 Datasets

To analyze the effect of various traits of player, matches, and other factors on the timing and a success rate of tackle, we will primarily be using data from the Kaggle NFL Big Data Bowl [4] specifically the files "`tackles.csv`", "`games.csv`", "`plays.csv`", and "`players.csv`" which contain data on the play, players, success, teams, etc. involved in a tackle. We will also be examining data from 9 weeks of tracking data ("`tracking_week_[1-9].csv`") to determine the effects of factors such as position on the field or movement at the time of a tackle.

3.1.1 Games dataset. The "`games.csv`" dataset contains the game ID (a numerical identifier unique to each game), the home team, the visiting team, the resulting scores, the week of the season in which the game took place, the date, and the time of every game within the 9 week period.

3.1.2 Plays dataset. The "`plays.csv`" dataset includes the game ID, the play ID (a numerical identifier unique to each play performed during a game), the quarter of the game in which the play occurred, the down, the distance, the possessing team, the side of field, the offensive formation type, the defenders in the box, a text description of the play, the resulting yardage, the ball carrier player ID (a numerical identifier unique to each player, specifying, respective to the play, which player was carrying the ball during it), and penalties (if applicable) of every play within the 9 week period. Kicking plays (kickoffs, punts, and field goals) are excluded. Plays are identified by both the game ID and play ID, as play IDs are not unique across games.

3.1.3 Players dataset. The "`players.csv`" dataset contains a player ID (a numerical identifier unique to each player), player name, position, height, weight, date of birth, and the college from which they originated.

3.1.4 Tackles dataset. The "`tackles.csv`" dataset contains the game ID, play ID, player ID, whether it was a tackle or assisted tackle attempt, and whether the tackle was missed (a metric provided by the third-party vendor Pro Football Focus). The player ID is that of the ball-carrying player (usually the one being tackled).

3.1.5 Tracking dataset. Players on the field, as well as the football itself, are equipped with tracking devices which enable recording of their location, speed, acceleration, and altitude. These variables are included in the tracking dataset, along with the associated game ID, play ID, player ID (if applicable), and event description (such as `tackle`). This data is divided into 9 different datasets, one for each week of the 9-week period.

3.2 Exploratory analysis

Our primary goal for this project was to find correlations between previously known information such as weight, height, and time during a game and the tackles performed (location, likelihood of success and number of tackles). To this end, we aimed to find correlations between factors such as height and weight of player, and frequency and likelihood of success of tackles, as well as frequency analysis for examining which teams/players tend to engage in tackles. Analysis methods included both the naïve Bayes and decision tree classifiers.

Expanding beyond this, we also explored how strong of a correlation exists between defensive players making or missing tackles and their team winning games. We used data on the number of tackles performed during a game to see the effect on the game's outcome. This was intended to provide insight on how tackles can affect the course of the game.

Finally, we used compiled this information to determine potential strategies and implications on real-world applications, including the selection of offensive strategies and the use of certain formations.

3.3 Grouping tackles

When performing a tackle, the physical attributes of the players involved such as their relative speed, direction, location, etc. may have an effect on how tackles are performed and whether they are sucessful. By filtering tracking data from "`tracking_week_[1-9].csv`" on tackle events, we are able to determine each player's movement at the moment of a tackle. Since this data is largely numerical and are likely not independent, we have used cluster analysis using a k-means algorithm to determine any patterns in the dataset.

Using the Elbow Method, we choose 5 clusters for this dataset and applied the k-means methods to a number of combinations of attributes. By joining this dataset with the plays dataset, we were also able to determine which plays are more likely to have tackles and the relative motion of players with respect to each other at the time of tackle.

3.4 Offensive strategies

NFL offenses are always looking for ways to stay ahead of defenses. Putting ourselves in the shoes of an NFL offensive coordinator, we formulated three potential offensive strategies to investigate.

3.4.1 Optimizing the weight of running backs. Many play designs involve the quarterback handing the ball off to a player designated as the running back (RB), who attempts to run forward for as many yards as possible before being tackled or forced out of bounds by the defense.

For this offensive strategy, we will check if optimizing the weight of a running back is associated with a more favorable outcome for an increase in missed tackles, and if playing heavier running backs is associated with a positive impact on tackle performance.

3.4.2 Increasing passing yards. For a receiver, making catches is important, but just making the catch isn't all that matters. Yards after catch, the distance a receiver travels after making a catch but before a tackle, being forced out of bounds, or scoring, is also an important metric. Higher yards after catch will result in more yards gained during passing plays.

To increase yards after catch, offenses may need to make longer passes down the field to areas where there are fewer defenders nearby to make the tackle. For this offensive strategy, we analyze whether longer passes resulting in more yards after catch is supported by the dataset.

Note that this analysis will only include plays where a pass was made which requires filtering out observations without a value for `passLength`.

3.4.3 Choosing specific offensive formation types. An offensive coach must decide on a formation to use depending on what the play is trying to achieve, e.g. short run or long pass. The game permits several different formation types, with different formations placing players at different locations behind the line of scrimmage, or the line which separates offense and defense at the start of a play and marks the distance remaining for a first down.

A formation also directs the movement of players after a play begins, based on their positions in the game. For example, a wide receiver may be directed to run down the field to provide a target for a pass. Defenses are allowed to respond to offensive formations based on how the offensive players line up, but may not always know the offense's full intentions.

For this offensive strategy, we will analyze offensive formations to see if any formation types are associated with more favorable outcomes for offenses seeking to increase the amount of plays with missed tackles. We will also determine how offensive formation preferences change based on variables such as down and distance.

3.5 Tackles as a predictor of victory

"The best offense is a good defense" and "defense wins championships" are popular sayings, but is it really true? We will explore how strong of an association exists between defensive players making or missing tackles and their team winning games, and whether tackles can be used to predict game outcomes.

4 Evaluation

4.1 Exploratory analysis

By analyzing the frequency of `yardsToGo`, we were able to determine that the vast majority of plays started at 10 yards to go. This is likely because 10 yards is the starting point for each round. We also see more observations less than 10 yards to go than greater meaning moving forward is much more likely than backwards. The slight increase observed at 15 or 20 yards to go is likely due to penalties as -5 or -10 yards are common penalties.

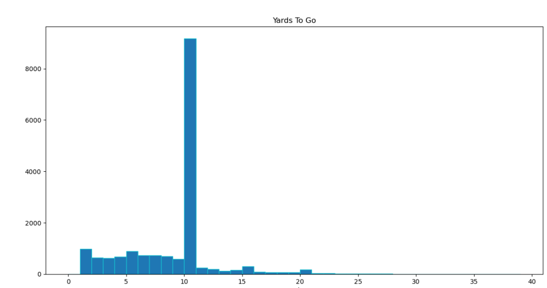

Figure 1: Yards to go

This also shows that most tackles occur fairly close to the starting line of the play.

4.2 Classifier predictions of tackle success

To determine if certain factors about the ball-carrying player (the player being tackled) made a tackle more likely to succeed, we attempted to use both a naïve Bayesian classifier and a decision tree classifier to predict the likelihood of a tackle succeeding based on the following attributes:

- Weight: The player's height, in feet and inches, sorted into 3 categories: ($< 5'10''$, $5'10'' - 6'2''$, $> 6'3''$, which were codenamed as "Short", "Average", and "Tall")
- Height: The player's weight, in pounds, sorted into 3 categories ('Under 200', '200 - 250', 'Over 250')
- Position: The player's position on the team (FB, QB, RB, TE, WR)

Two different classifiers were chosen in order to compare the results. Meanwhile, the attributes listed above were chosen because they were intrinsic to the player and likely meaningful to the effect of the tackle, or involved with their role on the field (reflecting possible influences such as positioning, play style, etc. and thus meaningful to the game), and thus could inform decisions about how to be more successful with future tackles.

We used 67% of the data for training the Bayesian and decision tree models, and the rest for testing each model's effectiveness in predicting success. Both resulting models were able to predict whether or not a tackle succeeded with accuracy of 88.6% and F1-score of 0.94. This shows that the factors chosen were, at least at face value, effective in predicting the likelihood of a tackle succeeding.

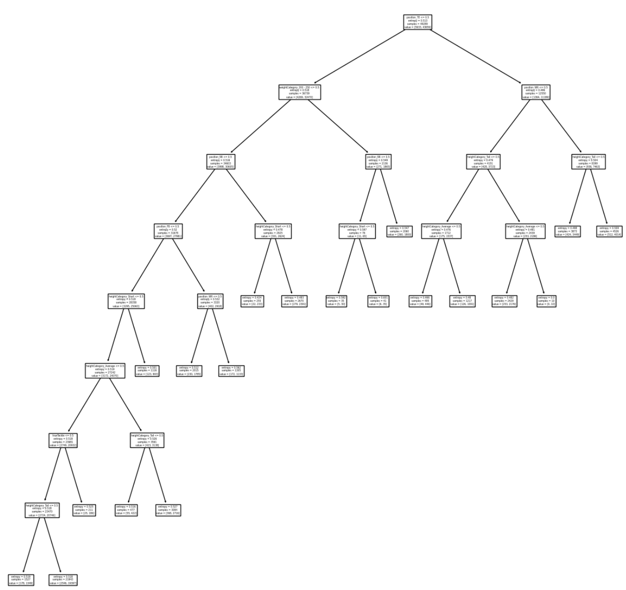

Figure 2: The resulting decision tree for tackle likelihood prediction

However, it is worth noting that the dataset of tackles provided by the NFL had relatively few records of "missed" tackles in comparison to "made" or "successful" tackles. This likely affected the accuracy of the classifiers and made them seem more predictive than they are.

When the dataset was adjusted using SMOTE (Synthetic Minority Oversampling Technique), the Bayesian model had an accuracy of 50.9% and an F1-score of 0.51, while the decision tree had an accuracy of 50.4% and an F1-score of 0.67. This supports the idea that the accuracy of the classifiers is inflated by the proportion of made to missed tackles, and once balanced, their predictive abilities are no longer as strong.

4.2.1 Random forest. To classify plays based on the presence of a missed tackle, a random forest classifier was constructed and trained using only variables known at the start of the play: quarter, down, side of field (own territory, midfield, or opponent territory), yards to go, offensive formation type, and defenders in the box. SMOTE was performed to account for the imbalance in the dataset, and the resulting dataset was split into 80% training 20% testing. The classifier achieved an accuracy of 84.0% and F1-score of 0.83 when predicting whether the play would result in at least one missed tackle.

4.3 Grouping tackles

Using the x and y attributes of the tracking dataset which correspond to the position of the football in yards at the time of the tackle, we attempted to determine if there any patterns in position at the time of tackle. We were unable to find significant clusters of positions where tackles occurred. As shown in the figure below, the locations of tackles do not appear to be in any statistically significant clusters. Notably, tackles tended to occur much more frequently close to the center of the field in both the x and y axes. One explanation for this is that players may be forced out of bounds rather than tackled, which also ends the play.

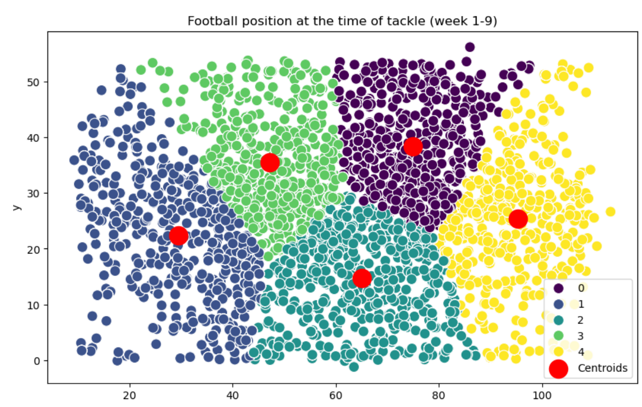

Figure 3: k-means analysis of football position

By using speed and acceleration of the football which acts as a proxy for the speed and acceleration of a player as they are being tackled, we are able to see that most tackles are clustered at fairly low speeds and accelerations. A second much smaller cluster of tackles also happen at low speeds but high accelerations. One explanation for this is that most players are seeking to avoid injury, so will not try to go faster if they are already likely to be tackled.

Additionally, by comparing the speed and direction of the "tackler" and the ball carrier, we can see that in most cases, the 2 are traveling in the same direction. However, the speed of the ball carrier tends to be higher than that of the defender likely because they

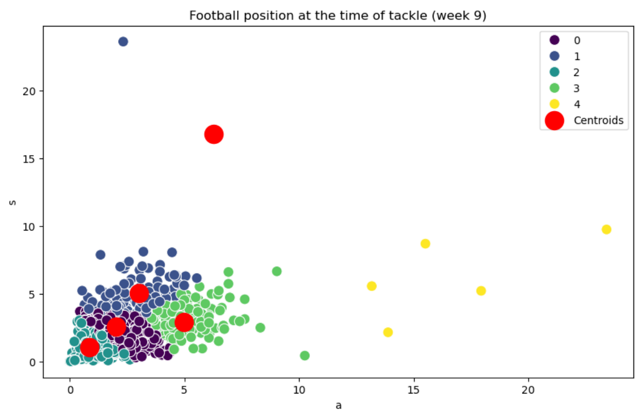

Figure 4: Speed and acceleration at time of tackle

must respond to what the ball carrier is doing and the ball carrier must outrun multiple defenders.

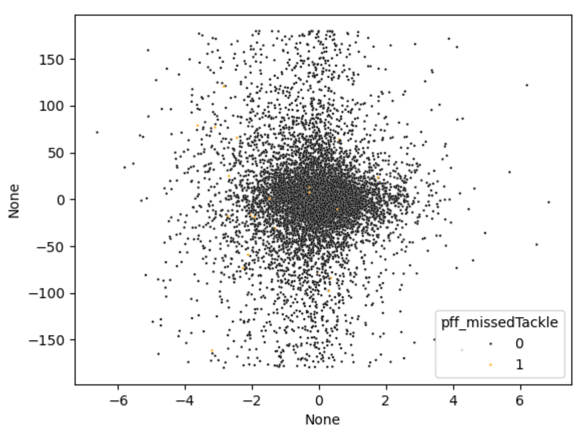

Figure 5: Difference in speed and direction between defender and ball carrier

4.4 Comparing prediction models

When comparing the predictive models of tackle likelihood based on player attributes, the differences between the naïve Bayes and decision tree classifiers was very small. They had identical accuracy and F1-scores under the original dataset, and very similar ones under the adjusted SMOTE dataset. Thus, in terms of predictive power, neither is necessarily better than the other. However, the decision tree offers a more easily interpretable visual representation, while the Bayes model is more flexible in terms of programming and individual predictions, so each has their own strengths and weaknesses in practical use.

4.5 Offensive strategies

4.5.1 Optimizing the weight of running backs. Running backs aren't significantly heavier than other offensive players. The heaviest players play at the key blocker positions of center (C), offensive guard (G), and offensive tackle (T), where the goal is to block defensive players from entering the offensive backfield and running long distances is rare. Players playing quarterback (QB), wide receiver (WR), and tight end (TE) don't need to be as heavy to do their jobs effectively. Players playing at wide receiver (WR) come in at the lightest weight, which makes sense for a player needing to be agile enough to make mid-air catches.

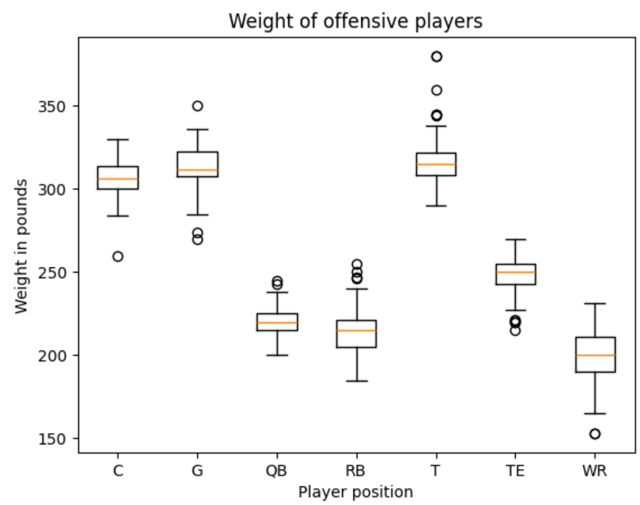

Figure 6: Weight of offensive players, grouped by position

For the rest of this analysis, we focused on played running backs and removed those who weren't involved in at least one play. The weight of played running backs is roughly normally distributed, with a mean of 213.82 pounds and standard deviation of 12.75 pounds. This was confirmed using a quartile-quartile (QQ) probability plot, which also revealed some outliers at the tails.

We binned running backs by weight into 5 bins to determine if running back weight had an affect on the number of plays with at least one missed tackle. Ultimately, it was inconclusive whether running back weight had a significant impact on tackle performance. The dataset lacked a sufficient number of missed tackle observations to meet our target number of at least 5 for each bin.

Table 1: Running back weight and play counts

Weight category	Missed tackle	Tackle or assist (not missed)
185 - 198	4	463
198 - 211	9	2007
211 - 224	17	2504
224 - 237	14	1078
237 - 250	2	596

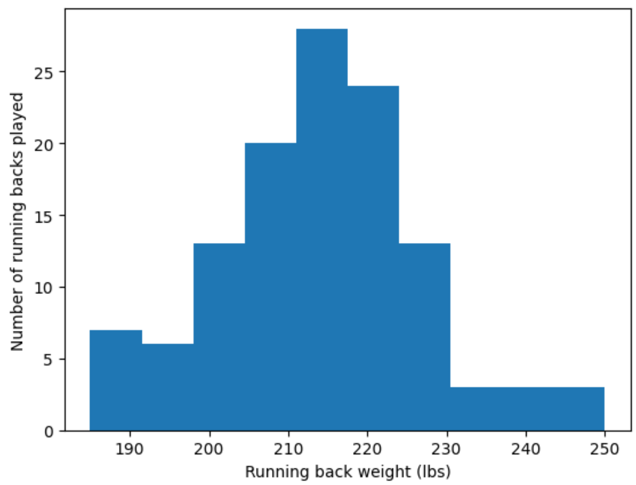

Figure 7: Histogram of played running back weight

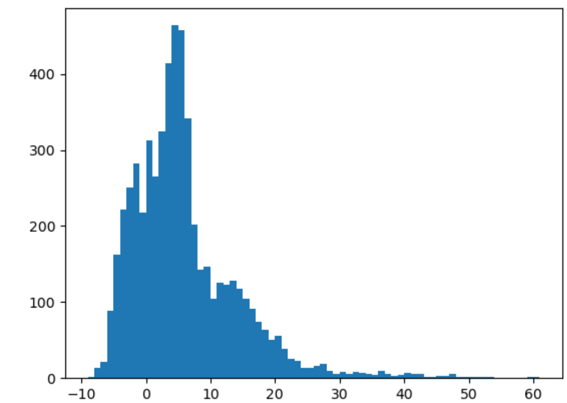

Figure 9: Length of passes

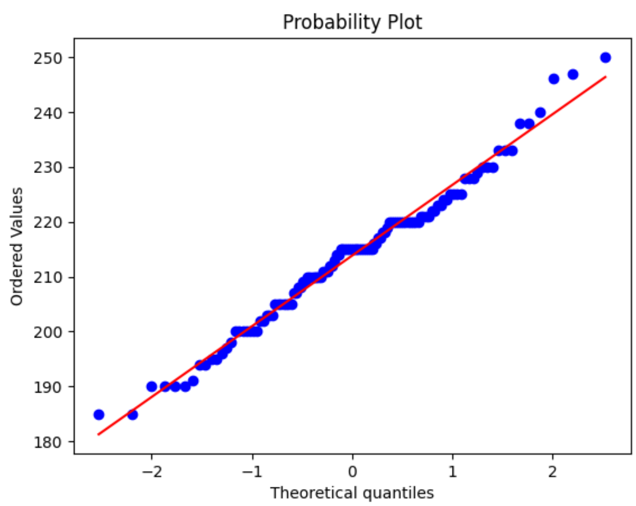

Figure 8: QQ plot of played running back weight

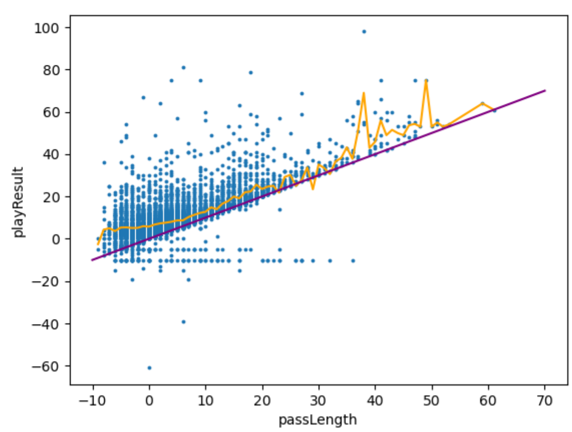

Figure 10: Pass length vs. result

4.5.2 *Increasing passing yards.* Examining the distance the quarterback passes the ball, we can see that distribution of passes is strongly skewed right with 1st and 3rd quartiles of 0,9. This indicates that most passes tend to be low with most falling between -5 and 10 as shown by the bar chart below

Comparing the pass result with passing yards, we can see that the most play results tend to be higher than the yards passed because players will try to move the ball forward. It is also noticeable that several outliers fall below this line which can be explained by penalty points of 5 or 10 or if the ball was intercepted by the opposing team. There are also some outliers which are very far above the average distance the player goes which indicates the carrier was able to dodge most of the defenders.

Examining the difference between play result and pass length (figure 11) as a proxy for the amount of distance an offensive player goes before being tackled, we can see that in the range of approximately [-5,8], the average distance before tackle decreases as the pass length increases probably because at these distances, other team members can help prevent the ball carrier from being tackled. The average distance covered changes by -0.76231569 yards for each extra yard the pass covers. Note that this linear regression has score of 0.027 which indicates the data is likely not linear.

We can also see that the play result seems to be the primary variable used by the NFL to predict the expected score for a game(represented by the color of a point) and that the prediction "penalizes" teams which have their ball intercepted presumably because this indicates a lack of skill.

4.5.3 *Choosing specific offensive formation types.* Examining the Plays data, we found six formation types played by offenses: Empty,

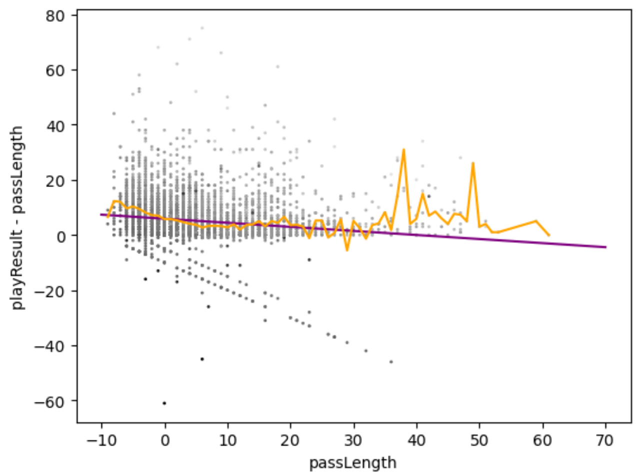

Figure 11: pass length vs distance before tackle

I-formation, Jumbo, Pistol, Shotgun, Singleback, and Wildcat. Overall, Shotgun is the most popular formation type, followed by Singleback.

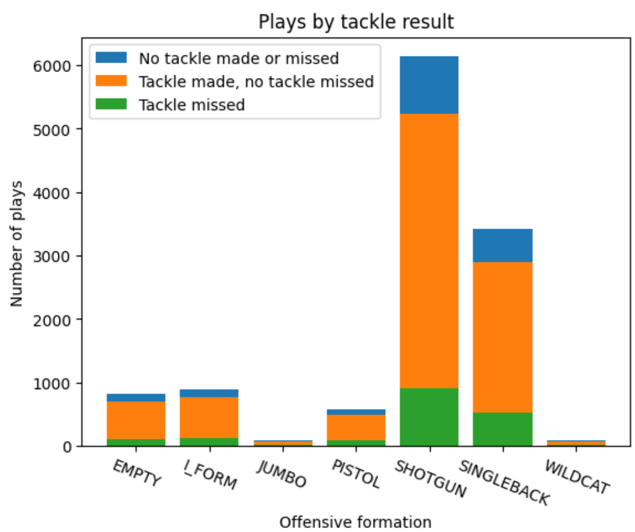

Figure 12: Plays by tackle result, grouped by offensive formation

However, formation distributions change based on down and distance. One special case is third and long situations (yards to go greater than 10, as a result of penalties, quarterback sacks, or other negative yardage plays). For these plays, picking up a large amount of yards is crucial to the offense maintaining possession of the ball.

For third and long situations we found that the popularity of the Singleback formation dropped considerably, and the popularity of the Empty formation rose considerably, compared to the overall distribution. This supports a conclusion that Empty formations are optimized for long passes down the field, which may have a lower chance of success overall but are seen as necessary to convert a third and long.

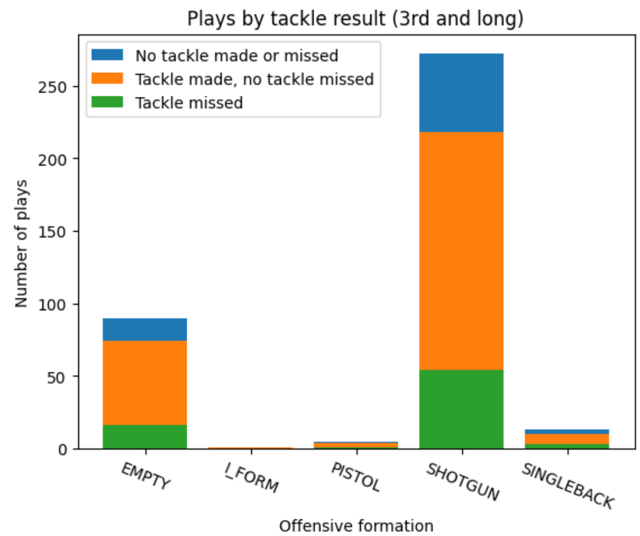

Figure 13: Plays by tackle result in third and long situations, grouped by offensive formation

Missed tackles are very detrimental for a defense, as the majority of offensive touchdown plays involved at least one.

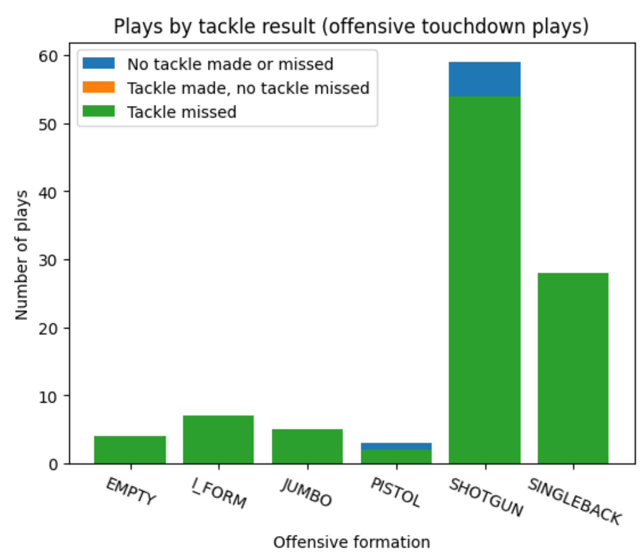

Figure 14: Plays by tackle result resulting in an offensive touchdown, grouped by offensive formation

We performed a chi-squared test to determine if offensive formation had an affect on the number of plays with at least one missed tackle. Ultimately, with $p = 0.888$, it was insufficient to disprove the null hypothesis and find that the offensive formation had a significant impact on tackle performance. The dataset lacked a sufficient number of missed tackle observations and was heavily imbalanced toward certain offensive formations.

4.6 Tackles as a predictor of victory

To determine if the quantity of tackles themselves contributed to the likelihood of winning a game, another naïve Bayesian classifier was constructed based on the number of successful tackles performed by each team during a game.

From the provided dataset of tackles, the number of successful tackles was aggregated by game and team. This was then combined with whether or not the team won the game in question. From there, the number of tackles made by a team during one game was bucketed into three categories: "High" (over 61 made tackles), "Average" (between 45 and 61 made tackles, inclusive), and "Low" (under 45 made tackles). This information was then used as the attribute for a Bayesian classifier to predict likelihood of victory based on the number of successful tackles.

The resulting classifier had an accuracy of 60% and an F1-score of 0.62. Broadly, it categorized a "High" number of tackles as a lost game, and an "Average" and "Low" number of tackles as a won game. Thus, it does not appear as if the number of successful tackles has a significant impact on the likelihood of victory. Additionally, when the model was slightly modified to predict victory based on the number of missed tackles, the results were comparable.

Again, however, the accuracy scores were likely affected by the low proportion of missed to made tackles in the dataset. When the dataset was adjusted with SMOTE, the accuracy was then about 57.8%, and the F1-score was 0.60. Accuracy was therefore lower, but the general implications with regard to the lack of applicability and predictive power of the number of tackles remained the same.

5 Discussion

5.1 Imbalance in missed tackle observations

The missed tackle class in the tackles data, provided by the third-party vendor Pro Football Focus as noted in the dataset description, contains a significantly smaller number of missed tackle observations (2,090) than non-missed (15,336). This imbalance made producing statistical significance when attempting to check for a correlation between missed tackles and another variable, or training classifiers to predict missed tackles, difficult. We were able to use Synthetic Minority Over-sampling Technique (SMOTE) to partially mitigate this for classifiers, but ideally we would have a much larger dataset to work with in order to produce statistical significance in all our forms of analysis.

5.2 Other sources of data

This dataset is limited to only nine weeks of gameplay in the 2022 season, which is a tiny sample of the decades of NFL gameplay data that exists. Websites such as espn.com and footballdb.com have compiled extensive amounts of data on NFL games, up to and including the current 2024 NFL season, and made it available to the public. The catch? These websites haven't made it available as an easily downloadable CSV file, requiring the use of a web scraper to collect it. They also tend to lack the detailed play-by-play tackling and tracking data of this NFL-provided dataset.

Nevertheless, many useful insights have been collected by these websites based on the massive database of gameplay data they contain. A more extensive analysis should also include data from these other sources when analyzing metrics such as player weight and yards gained.

5.3 Expanding project scope

This project struggled to produce many useful results because of its narrow focus on tackling which was believed to be well-suited to this dataset. Tackling is just one of dozens of metrics which can be used to predict play and game outcomes for NFL games. Perhaps the scope should've been expanded to embrace other metrics for analysis more thoroughly.

Despite the narrow focus, some useful insights not related to tackling were discovered throughout the analysis of this dataset. This included the weight distribution of running backs, the length of passes, offensive formation type preferences, and the yards to go distribution, across all the included NFL plays.

5.4 Utilization of tracking data

We had considered using temporal data particularly the tracking data to see if we can find any patterns over time. However, this was difficult due to the lack of historical data, short period of the tracking data, and its incompleteness i.e. only including a few frames around each notable event in the game. This was deemed to be outside the scope of the project though possible of interest for future projects.

5.5 Reflection

As with any form of data analysis, we aimed to find meaningful patterns and trends that can illuminate bigger truths about our area of study without unwarranted manipulation. Additionally, in order to conclude that these patterns and trends are in fact meaningful, we must be able to disprove the null hypothesis.

In essence, if we are mining the data for the effect of an example X variable, the null hypothesis would state that there is no statistical difference in tackle rate that can be attributed to X variable, and that any observed difference is due to nothing more than chance. Thus, in order to disprove it, the pattern would need to appear a sufficient amount of times within the dataset to surpass the significance level (typically set at 5%).

Notably, this is not the same as *proving* the hypothesis. Rather, we only had the ability to disprove that it does not hold. This is often the case with real-world data, and we are still be able to present our findings with some significance if that is the case.

In addition to these metrics, we evaluated our work based on the thoroughness of our investigation and our consideration of all factors. Real-world data is rarely straightforward, and contains many confounding variables that can all have impacts on the observed results. Accordingly, we didn't make make finding an "answer" our goal, but instead focused on exploring the data as best we can.

Even if we were unable to find conclusive or applicable implications for play strategies, we value what information our work does provide, and look forward to the future work that may result.

6 Conclusion

6.1 Predicting tackle success

We used three main methods of prediction of tackles' success: clustering tackles by position, speed, etc and finding clusters of success or failure, using classifiers and to predict whether a tackle will occur and if it succeeds, and analyzing potential offensive and defensive strategies to see if they have an effect.

6.1.1 Exploratory analysis. Most tackles occur fairly close to the starting line of a play (i.e. with 10 yards to go) generally ahead of it. Additionally, fouls by either team can move the starting point by 5 or 10 which can be identified as oddities in the data.

6.1.2 Grouping tackles. Attempting to group tackles using k-means clustering, we were unable to find any significant clusters of failure or success. We were able to determine that tackles tend to be more densely clustered close to the center of the field and players are likely to be going in the same directions when tackling and the ball carrier tends to be going at higher speed. Players also tend to have lower speed and acceleration at the moment of tackle regardless of success.

6.1.3 Classifying tackles. Initially, it appears as if player attributes have a fairly strong correlation with their likelihood of being tackled. However, this is likely due to the prevalance of made/successful tackles in the dataset, and therefore the relative scarcity of missed tackles.

When the dataset is adjusted for this using SMOTE, the accuracy and F1-scores of the models greatly decrease. As a result, we were not able to find a meaningful association between player characteristics of height, weight, and position, and their likelihood of being tackled.

A random forest classifier trained on variables known at the start of a play proved to be more promising, with an accuracy of 84.0% and F1-score of 0.83 predicting whether the play would result in at least one missed tackle.

6.2 Offensive strategies

We tested three offensive strategies for their usefulness: the weight of running backs, increasing passing yards, and offensive formations.

6.2.1 Optimizing the weight of running backs. The weight of played running backs is roughly normally distributed, $m = 213.82$ pounds, $sd = 12.75$ pounds. Running backs aren't significantly heaver than other offensive players. We were not able to find an association between running back weight and missed tackles.

6.2.2 Increasing passing yards. With shorter passes, the distance players can go before being tackled is decreasing possibly because they are more likely to get surrounded. This effect is on average overshadowed by the extra yards gained by longer passes. There is also very little data for long passes but there is some data to suggest that long passes can lead to longer runs without being tackled.

6.2.3 Choosing specific offensive formation types. The dataset contains six offensive formation types, with Shotgun and Singback being the most popular overall. Offensive formation preferences change based on down and distance, suggesting that some formations work better depending on what the offense's goals are. We were not able to find an association between offensive formation and missed tackles.

6.3 Tackles as a predictor of victory

As previously noted, the provided dataset is unfortunately limited in terms of the representation of missed tackles. Many more made tackles are available in the data, and thus predictive classifiers are skewed towards them.

Even so, the number of made (or missed) tackles does not have strong predictive ability with regards to a team's victory, whether or not the dataset was adjusted with SMOTE. We were not able to find an association between number of tackles and winning games.

6.4 Overall

Though our research is limited in its applications and significance, we still had interesting findings in a variety of game aspects, and we look forward to seeing what future research can be done in this field.

References

[1] NFL 2024. *Next Gen Stats.* NFL. Retrieved Sep 27, 2024 from https://nextgenstats.nfl.com/glossary
[2] Ezra Ford. 2023. *Tackling Metrics with Big Data.* Medium. Retrieved Sep 27, 2024 from https://medium.com/@ezra.ford/tackling-metrics-with-big-data-0812b5ab65f0/
[3] Ron Jadischke, Jessica Zendler, Erik Lovis, Andrew Elliott, and Grant C. Goulet. 2020. Quantitative and qualitative analysis of head and body impacts in American 7v7 non-tackle football. *BMJ Open Sport & Exercise Medicine* 6, 1 (2020).
[4] Michael Lopez, Thompson Bliss, Ally Blake, Andrew Patton, Jonathan McWilliams, Addison Howard, and Will Cukierski. 2023. *NFL Big Data Bowl.* Kaggle. Retrieved Sep 27, 2024 from https://www.kaggle.com/competitions/nfl-big-data-bowl-2024/
[5] Arthur Maerlender, Caitlin J. Masterson, Rex Norris, and Adam Hinthorne. 2020. Validating Tackle Mechanics in American Football: Improving Safety and Performance. *Annals of Biomedical Engineering* 48, 1 (2020), 2691–2700.
[6] Hirokazu Matsuo, Yukio Yamada, Katsuyuki Masuchi, and Tsuyoshi Matsumoto. 2018. Effectiveness of the Heads Up Tackling (HUT) Program on Tackling Safety and Performance in American Football. *International Journal of Sport and Health Science* 16, 1 (2018), 77–88.

A Honor Code

On my honor as a University of Colorado Boulder student I have neither given nor received unauthorized assistance.

B Individual contributions

- Isaac Kou
  - Classifying tackles: Naïve Bayesian
  - Classifying tackles: Decision tree
  - Tackles as a predictor of victory: Naïve Bayesian
  - Classifying tackles & Tackles as a predictor: SMOTE re-analysis
- Sean Shi
  - Exploratory analysis
  - Grouping tackles: k-means analysis of football speed and acceleration
  - Offensive strategy: Increasing passing yards
- Timur Tripp
  - Grouping tackles: k-means analysis of football position
  - Offensive strategy: Optimizing the weight of running backs

- Offensive strategy: Choosing specific offensive formation types
- Classifying tackles: Random forest

Received 15 December 2024