Measuring Grit in NFL Cornerbacks using Statistical Analysis

by

Cole Kingston

B.S. Computer Science and Engineering,
Massachusetts Institute of Technology, 2023

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Authored by: Cole Kingston
Department of Electrical Engineering and Computer Science
May 12, 2023

Certified by: Anette Hosoi
Associate Dean, MIT School of Engineering
Neil and Jane Pappalardo Professor of Mechanical Engineering
Thesis Supervisor

Accepted by: Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Measuring Grit in NFL Cornerbacks

using Statistical Analysis

by

Cole Kingston

Submitted to the Department of Electrical Engineering and Computer Science
on May 12, 2023, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Using the pass play tracking data from the 2018 National Football League (NFL) season, I compiled a Grit Score that measured cornerback responses to an adverse result to a play. I calculated this Grit Score using the results of whether a cornerback allowed their opposing receiver to catch the ball to measure change in performance. When comparing performance, I used the difference in average distance between the cornerback and opposing receiver to compile one score for each player in the NFL. I validated my calculations with Pro Football Focus Coverage Ratings and was able to classify players into 6 different categories based on talent and Grit Score. Overall, I found that most NFL players have high grit, or play consistently through adversity, which explains why they have made it to the highest level of football. NFL coaches and general managers prefer players who have increased performance following a bad event as those players tend to stay in the NFL for longer than those with decreased performance.

Thesis Supervisor: Anette Hosoi
Title: Associate Dean, MIT School of Engineering
Neil and Jane Pappalardo Professor of Mechanical Engineering

Acknowledgments

I must first thank my advisor, Peko Hosoi. This thesis would not have been possible without her invaluable advice and commitment. Her incredible insights and cheery personality made this process a far more enjoyable experience. I cannot say enough about her incredible leadership ability and passion for helping her students.

I would also like to thank my parents, Sarah and Jeff, and my sisters, Claire and Jane, for their constant love and support as I would not be where I am today without them. In addition, I also thank my girlfriend Gigi for her help with editing this thesis and providing encouragement and compassion throughout this process. My dog Quinn deserves acknowledgement as well for being a consistent source of enjoyment and unwavering loyalty.

There are too many other people to name that helped me get to this point, but I would like to recognize a few other groups of people that helped me. I am grateful to my friends in E Calcs who never fail to make me laugh and have stood behind me since early childhood. I would also like to thank the gentlemen of DKE for their mentorship and for making MIT a fun place to learn. Lastly, I must thank MIT Football for helping me get into MIT and providing the inspiration for this thesis.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

In today's National Football League (NFL) games, the level of physical talent among players is becoming increasingly balanced across the league, causing a need for new evaluation techniques. Measurable attributes such as speed, strength, and playmaking ability are no longer enough to understand whether a particular player will succeed. As the talent gap narrows in football due to increasing athleticism, teams are now placing a particular emphasis on the mental aspect of the game. Specifically, NFL coaches and general managers want players that can carry out the game plans to near perfection throughout the entirety of the game. While most players are able to do this when the game is going as expected, the variance of football almost guarantees that each player will have to go through adversity at some point during each game.

One of the most desirable traits of NFL players is grit. While physical grit is often closely associated with toughness, such as playing through an injury, mental grit is far more difficult to quantify. Trainers can detect the degree of a player's injury through various physical tests. Therefore, observing the effects on performance following these injuries will give an accurate measure of physical grit. This aspect of grit is seemingly precise, but the mental part of grit poses more ambiguity.

Mental grit in football is a largely unexplored field due to limited testing strategies which create inaccurate results. In current sports analytics and psychology, grit is only measured by mental fortitude tests or subjective opinions from football experts such as coaches, scouts, or analysts. Therefore, we do not have concrete evidence as to whether the perception of each player's grit is actually true in the context of their performance on the field. Current perceptions of a player's grit are generally controlled by anecdotes from teammates or coaches, as well as the media that reports this to the rest of the league. In addition, players are incentivized to give answers that will result in a more desirable type on personality tests. This could cause a team to offer them more money and thus lead to extreme bias in the results of these tests.

1.1.1 What is Grit?

Angela Duckworth, a current professor of psychology at the University of Pennsylvania, is one of the leaders in the studying of grit and resilience. Her work has explored grit as a predictor of success in a variety of contexts such as education, sports, and business. Through compiling her research, she describes grit as "perseverance and passion for a long-term goal." [2] To translate this into a football context, having perseverance and passion means performing at one's best throughout the long-term goal of winning the game. When a player experiences a play where they have bad luck or are outperformed by another player, the performance on the following plays can be heavily influenced by that one event.

How players respond to a negative event is a strong indicator of their grit and mental toughness. There are three ways that this event can affect a player's performance. In one case, if someone on defense with high grit allowed their opponent to score a touchdown, then we would expect them to have a similar performance on the ensuing plays as they had before the bad play. Players with the highest grit should be able to continue performing to their potentital regardless of previous plays. We could also have the case where a player with low grit gives less effort or worse execution of the game plan, so their performance following a negative play should see a significant decrease. Lastly, we could see the case where a player's performance benefits from a bad play. While this shows some grit and perseverance, it leads to other questions as to why they need a negative play to achieve full potential.

1.1.2 How Grit Affects Team Performance

Knowing which players will be able to respond well to adversity is essential to NFL management as this could be the difference between winning and losing a game. In fact, the 2022 NFL season saw its lowest average margin of victory in 90 years at just 9.7 points per game, a very low margin considering that teams generally score in increments of 3 and 7 points. Thus, mental mistakes must be minimized, especially towards the end of games when players are more fatigued and mentally exhausted.

Throughout the duration of the game, mental grit will keep a player level-headed which is ultimately better for team morale. Large mood swings in one player will distract and affect the tempers of other teammates. In a sport where each person must do their own job to near perfection for the team to win, one player can have rippling effects throughout a game or season that causes significant deviations from achieving full potential.

1.2 The Cornerback Position

Since football has many positions with vastly different responsibilities, this thesis focuses particularly on cornerbacks due to the significant mental challenge that comes with playing the position. Cornerbacks have one of the most difficult jobs on the field and therefore commonly experience failure.

On each play, the offense can either pass or run the ball. Run plays usually happen towards the middle of the field and cornerbacks typically play on the edges of the field, so their performance heavily relies on passing plays. For any passing play, they must try to keep the distance between them and a wide receiver to a minimum, then attempt to prevent the wide receiver from catching a pass by knocking the ball away. Not only does the job itself put the cornerbacks at a disadvantage, but NFL wide receivers are also some of the fastest and most athletic people on the planet. They can run anywhere on the field to get open with the quarterback trying to throw the ball in a spot where only the receiver can catch it. Thus, cornerbacks must be in a peak mental and physical state to effectively execute on each play.

Because the position involves a considerable amount of failure, cornerbacks tend to be more intense, outspoken, and outwardly emotional. While it may seem as if they are having breakdowns on the field from a subjective view, this thesis hopes to understand whether the actual performance follows suit. In short, we want to see whether NFL cornerbacks have the ability to put their emotions aside before each play, and continue playing at peak performance on a consistent basis, no matter the result of the previous play.

1.2.1 Coverages: Man vs. Zone

An important part of the cornerback's job is executing a particular plan on each play. This plan includes responsibilities for both the cases where the offense passes or runs the ball. In the case that they pass the ball, the plan is called a coverage. A coverage is meant to strategically prevent the offense from completing a pass on that play. In the time before each play, the defensive coach will signal the specific coverage to the captain of the defense who will communicate it to his teammates. This way everyone is on the same page for defending the pass on the ensuing play.

Coverages can be broken up into two main categories: man and zone. Man coverage is when each defensive player chooses one player to follow from the start of the play until the end of the play. A single cornerback's performance for that play hinges on staying as close as possible to their assigned receiver. This deters the quarterback from throwing to that receiver, or puts them in the best position to knock the ball away if the ball is thrown their way.

Figure 1-1 shows a team, with the defense represented as X's in a specific type of man coverage, cover 0. We can see that each defender is assigned to one offensive player, or they are trying to tackle the quarterback in the case that there is no assigned receiver for them. Cornerbacks almost always have an assigned receiver, as their specialty is matching up with them.

Zone coverage, as seen in Figure 1-2, is when each defender covers a particular area on the field and defends any receiver that comes into that zone. This differs from man coverage as the cornerback will not know which receiver they are responsible for

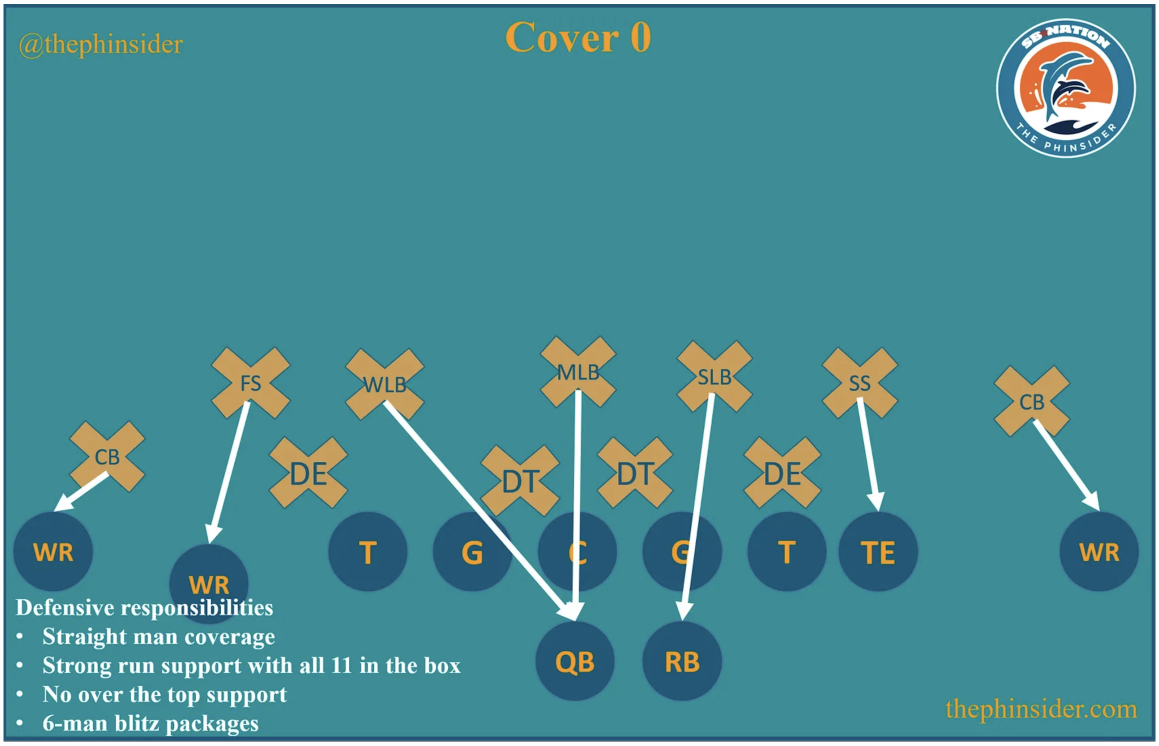

Figure 1-1: Example of Man Coverage with Each Defensive Player and an Arrow Pointing to Their Corresponding Receiver Responsibility [10]

until after the play begins. Having this late knowledge and changing responsibility typically means that the cornerback spends a portion of the play determining who they should guard, then moves towards them to defend a pass. Since they are only responsible for a receiver in their zone, they also may not have a receiver responsibility on a given play if there are more zones than recievers. Overall, zone coverage can have a lower amount of time or no time where a cornerback has a responsibility, but high performance during those times is essential to team success.

For both man and zone, the cornerback could start 1 to 7 yards away from the closest receiver depending on the coverage type and situation. A larger gap typically means that there is a lower chance of giving up a high yardage play, although, this gap is typically closed within the first seconds of the play.

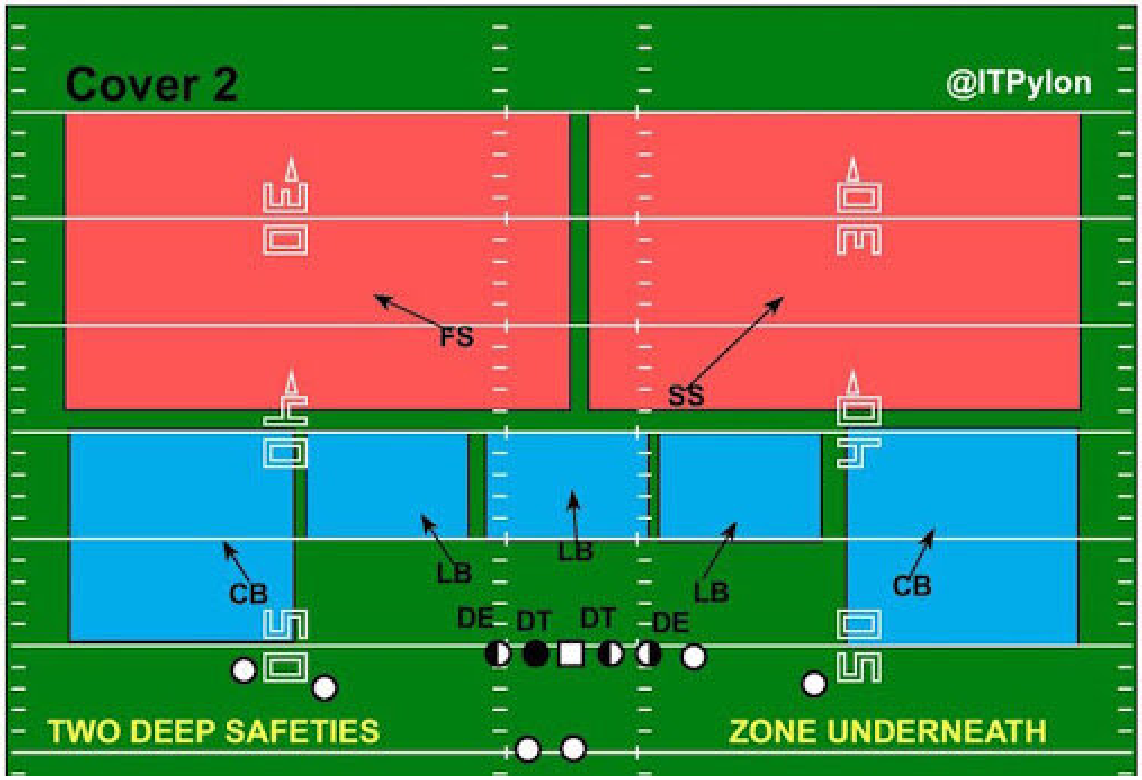

Figure 1-2: Example of Zone Coverage with Each Defensive Player and an Arrow Pointing to Their Corresponding Zone [7]

1.3 Impact

Through this thesis, I created a concrete way to measure the mental fortitude of NFL cornerbacks when responding to an adverse event. Using data from the 2018 NFL regular season, I extracted a metric for how each corner in the league responds to adversity and compared it to the current perceptions of those players. Using the information from this metric, NFL teams will be able to better evaluate potential free agents, and determine which players to sign to their team. Having players with more grit directly increases a team's chance to win, as they are less susceptible to big momentum swings during a game. In professional sports, winning games means more TV deals, media attention, and fan attendance to games, all of which increase the profits for the owners.

In addition to evaluating the current NFL players whom general managers have already seen play in the NFL environment, this could also be used to evaluate new players entering the league. In fact, the largest impact of this study would come when evaluating college players for the NFL Draft. The NFL Draft is an event where each team gets to claim rights to signing a college player for the following season. Historically, the player rankings before the draft do not closely correlate to success in an NFL career. The transition from the level of play of college football to the NFL is a considerable challenge. In addition to the increase in skill level, learning how to live on one's own creates further obstacles. These new experiences contribute to the high variance between the pre-draft rankings of college players and their actual performance. Due to the high variability of player success, having accurate evaluation techniques can create a substantial advantage for those teams who pick the best players in the draft. With access to the positional data from a college football player's season, these teams could use this research to understand a certain player's mental toughness.

Chapter 2

Related Work

Although the topic of quantitatively evaluating grit in the NFL has not been explored yet, there has been research to measure the impact of grit when attempting to achieve a long-term goal, as well as analysis of NFL players' cognitive abilities. In this section, I will review the findings of these projects, as this thesis will be an intersection of measuring cognitive ability and defining grit in terms of success. Understanding how each of these papers apply to this thesis is imperative to achieving the ultimate goal of creating a metric for grit in NFL players.

The most prominent publication regarding grit comes from Duckworth et. al. (2007) [2], where Ivy League students, West Point Cadets, and National Spelling Bee competitors were first evaluated on the Grit Scale which comes from a questionnaire that asks questions focusing on personality traits associated with grit. After understanding the grittiness of each participant, the study found the correlations between intelligence and eventual success using the Short Grit Scale [1] (GPA in the case of students and ranking in the case of Spelling Bee participants). There was a correlation between those with high grit and those who had a higher success rate in relation to their intelligence. In addition, grit accounted for about 4% of the variance in the outcomes of success for these high achieving academics. Although IQ is widely known to be the most direct predictor of success [6], other aspects of personality play a strong role as well. In environments such as high academic schools and the workforce, testing, prior experience, and applications have funneled people of similar intelligence together, making grit a strong differentiator. In fact, grit was proven to have a positive correlation with retention in the military, the workforce, school, and marriage. [3]

Grit is closely related to sports as sport participation levels have high correlation with work ethic in the real world. [9] Further, a study found that when comparing elite athletes and a comparison sample of non-athletes, the elite athletes had a significantly higher correlation to grit. [5]

In football, it is widely known that physical ability is the largest predictor of success, specifically speed. [13] With players at similar positions running at similar speeds in today's game, analytics have turned to other indicators of success on the field such as the mental aspect of the game. Each player entering into the NFL draft is asked to take the Wonderlic test, measuring their cognitive ability and problem-solving skills. However, recent studies have shown that there is no correlation between a player's scores on the Wonderlic test, and their eventual success in the NFL [8]. This is likely due to the fact that football intelligence is more based on preparation and instantaneous decision-making rather than understanding math or vocabulary questions. Current written tests for aptitude prove to be ineffective, likely due to the fact that they cannot simulate the speed and emotion of an NFL game. Instead, we can take a different approach of evaluating the cognitive ability, specifically grit in our case. The new model will use a results based statistical approach with real game data, allowing a direct computational analysis of grit rather than trying to measure cognition using a written test.

Chapter 3

Data Creation

3.1 Raw Data

To explore this new metric, I used the positional and informational data from the 2018 season through the 2021 NFL Big Data Bowl. The dataset includes all passing plays from each game from weeks 1 through 17 of the season. While the dataset provides a wide variety of information, this research focused on the positional tracking of each player over the course of a play. Using a chip in the player's equipment, the x and y positions of each player were sampled at a rate of 10 Hz. In addition, the plays were labeled with contextual information, events, and the results of the play. For this specific task we used player x positions, player y positions, player names, player teams, player type, team with possession, time, ball x position, ball y position, and play events.

3.2 Results vs. Performance

To understand which players are maintaining consistent performance throughout the game, we must first come to a conclusion on what success means for a cornerback. On any single passing play, a cornerback will be responsible for guarding one receiver or zero receivers in the case that they are playing zone defense and no receiver goes into their zone. The primary responsibility of a cornerback is to defend against passes and prevent the opposition from completing them down the field. A successful result would be an incomplete pass to a corner's receiver and a failure would be a complete pass to that same receiver. Because quarterbacks do not always throw to the most open receiver or can have inaccurate passes, the result of each play is not always representative of a particular cornerback's performance. Therefore, we must make a distinction between a result and a performance on a specific play. For the purposes of this research, the result of the play refers to whether the receiver that a certain cornerback was guarding caught the ball. Performance on the other hand, is how well the corner minimizes the average distance between them and their receiver. For this research, we want to see how each cornerback reacts following a negative event. Since humans tend to weigh outcomes heavier than the actions that they took leading up to it, the result is important for recognizing a failure. After finding this time of adversity, we want to measure performance of the cornerback before and after the play without outside influence of the quarterback.

3.3 Determining Receiver Responsibility

To understand performance and find plays where a cornerback faced adversity, we must first have a way to determine which receiver they were guarding on each play. Since a cornerback can often switch the receiver that they are guarding with another defensive player towards the beginning of a play, we want to see which receiver is closest to them at the time of pass release. Thus, we can take the x and y positions of all receivers on the other team and find the receiver with the minimum distance to the specified cornerback using the distance formula.

$$d_{rc} = \sqrt{(x_r - x_c)^2 + (y_r - y_c)^2}$$

We have x_r and y_r represent the x and y positions of some receiver, while x_c and y_c represent the x and y positions of the cornerback. d_{rc} is the distance between that receiver and our current cornerback who is being analyzed. We do this for all k receivers on the current play, then match the corner with his receiver based on which receiver has the minimum distance at pass release.

$$closest = min\{d_{rc}[1], d_{rc}[2], ..., d_{rc}[k]\}$$

If the cornerback's receiver were to catch the ball, we can also deduce that they are at fault and have failed their task.

To find the receiver who caught the ball, I calculated the shortest distance between all the receivers and the football at the time of pass completion. Again, I used the distance formula and took the minimum distance between any receiver and the ball. The receiver who had the minimum distance to the ball was then labled as such. Now that I had both the corner's receiver responsibilities and the pass catcher, I compared these to determine if any of the corners on that play were at fault.

3.4 Performance Metric Creation

In any NFL passing play, we can separate it into two parts: before and after the pass is released. Before pass release, a cornerback's job is to keep the receiver as close as possible, so they can be in the best position to prevent the receiver from catching the pass. Typically, the quarterback will throw to the receiver that has the most space between them and the cornerback, and the rest of the corners' jobs are essentially done. Because of this, I simplified the performance metric to only include the time before pass release. As mentioned previously, cornerbacks typically play a variable distance, or cushion, from their closest receiver based on the play call. To avoid penalizing cornerbacks who have larger average cushions or those who have nobody come into their zone in zone coverage, we start measuring performance of the cornerback once their receiver has come within a threshold of three yards on the play. While this choice of three yards may seem arbitrary thus far, it was strategically chosen through an analysis of options which will be discussed in section 4-2. If a cornerback's closest receiver never comes within three yards, then we omit the play because the ball was either thrown before their cushion was closed or a receiver never came into their zone. Now that we are looking at play times from when the cushion is reduced to three yards to pass release, we can simply define a cornerback's success as minimizing the distance between him and the receiver that he is guarding during that play.

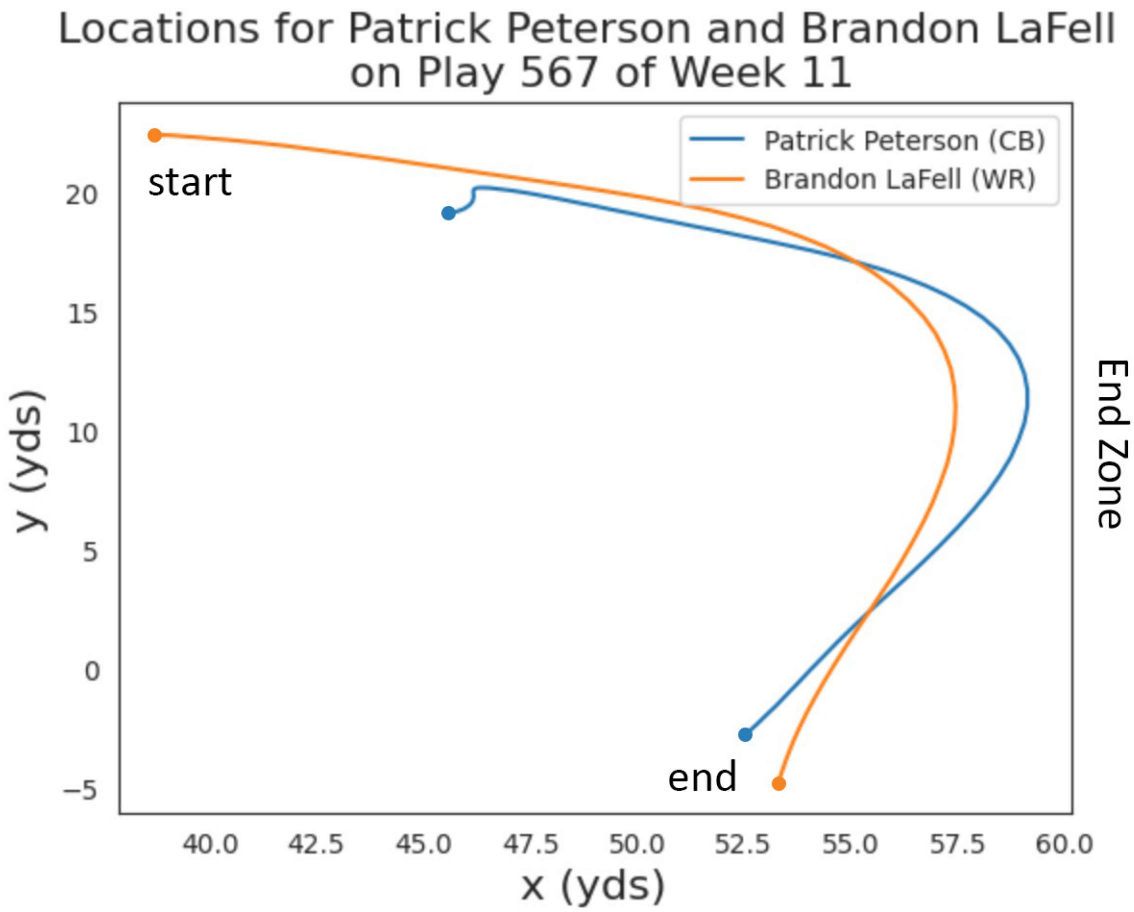

Figure 3-1: X and Y Locations of One Play for Patrick Peterson and Brandon LaFell on Play 567 of Week 11.

Figure 3-1 shows the locations of one of the best cornerbacks in the NFL, Patrick Peterson (blue), and the receiver that he is guarding, Brandon LaFell (orange), on a represenation of the football field. On this play, the offense is going from the left side of the field and trying to score by getting the ball to the right side of the field. Peterson starts with a 7.5 yard cushion, and stays close to LaFell as he approaches and closes the gap. We can also see that there is a divergence between the two players towards the end of the play. When seeing a situation such as this, we can infer that

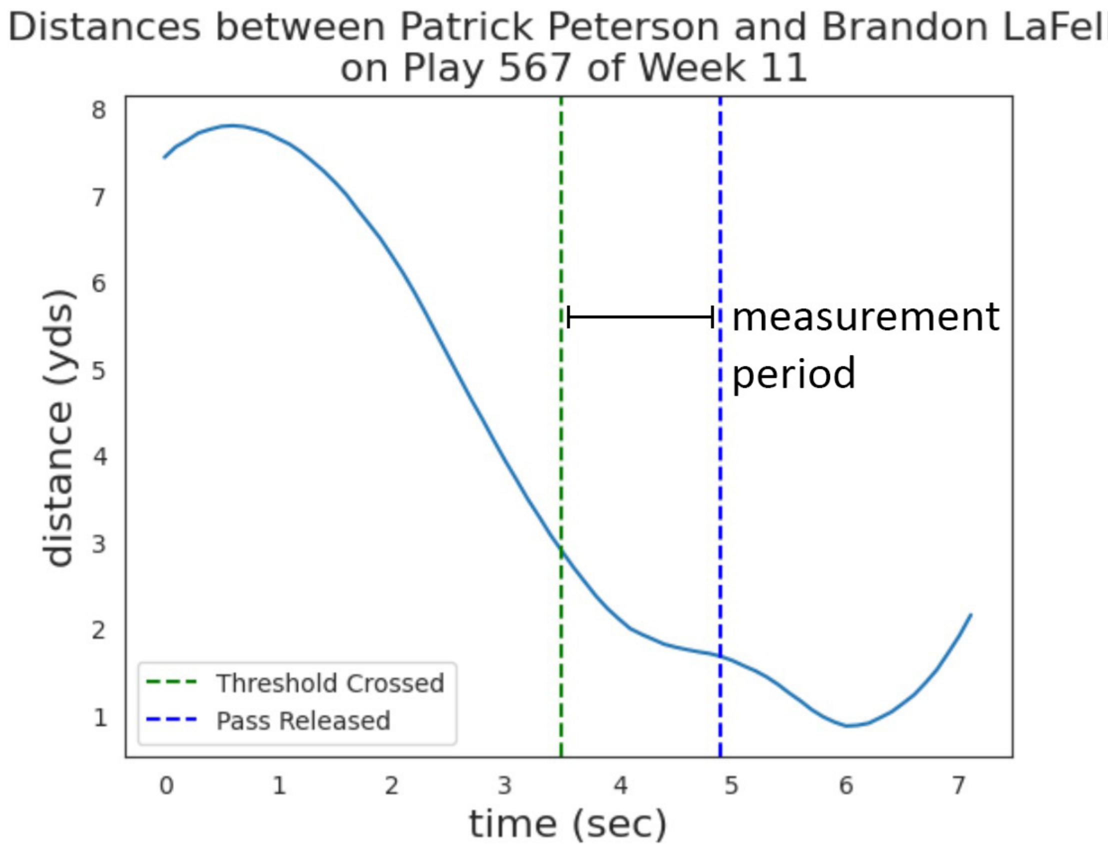

Figure 3-2: Distances Between Patrick Peterson and Brandon LaFell Over the Course of Play 567 of Week 11.

the pass was thrown before that separation was created, so Peterson stopped running or dove for the ball. In rare cases, this could also be the result of the cornerback falling before the ball was thrown. Either way, both of these cases are covered by our measurement choices.

To visualize our measurement choices, we can see the distances between a cornerback and their receiver as displayed in Figure 3-2 over the same play as Figure 3-1. At the start of the play, we again confirm the 7.5 yard cushion between the players. In addition to the distance over the course of the play, there is a green line and a blue line representing the time at which the three yard threshold is crossed and the time of pass release respectively. Using the threshold and timing of pass release prevents both penalizing for having a cushion given by the coach, and accounts for the distance created after the ball is thrown. In Figure 3-2, the distance between the two players rapidly increases at the end of the play. Had I measured the entire play, this data would be heavily skewed towards the time after pass release when the cornerback has already done his job. Through this visualization, only the essential times of the play are measured and therefore the performance measurement of the cornerback is more accurate.

3.4.1 Performance on Man vs. Zone Coverage

Due to the nature of each coverage, the tracking data of a cornerback and their receiver looks vastly different. We compare these again for Patrick Peterson in Figure 3-3 with man coverage and Josh Norman in Figure 3-4 with zone coverage. When comparing the differences between the cornerbacks' behavior in man coverage vs. zone coverage, we can see that in Figure 3-3a, Peterson follows the receiver as closely as possible throughout the course of the play and thus the shape of their locations are very similar. In Figure 3-4a however, Norman starts further away from his eventual receiver and moves closer to him as the receiver approaches his zone.

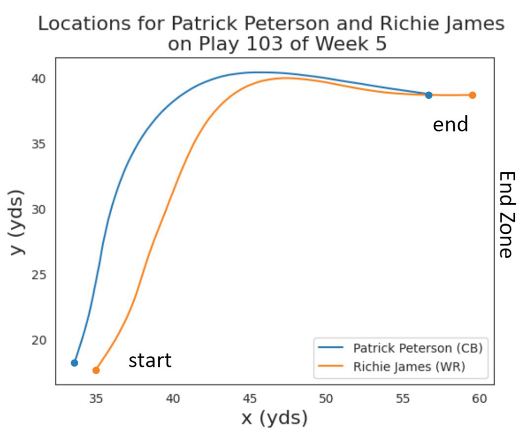

(a) Locations of a Receiver and Cornerback with the Defense in Man Coverage.

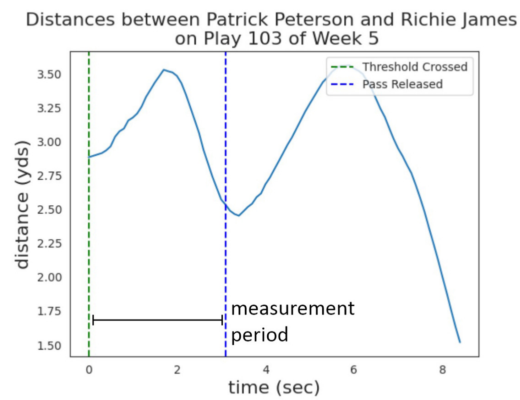

(b) Distances Between a Receiver and Cornerback with the Defense in Man Coverage.

Figure 3-3: Man Coverage

Now if we compare the distances over time in Figures 3-3b and 3-4b, we first notice that the measurement period for man coverage is far greater than the measurement period for zone coverage. Although it might not happen for all cases, this would be typical as a cornerback in man coverage knows his receiver responsibility at the start of the play, so he would immediately close the distance between the two players. For example, we can see that Norman does his job well as he quickly closes the distance between him and the receiver to under two yards at the time the pass is thrown. After the pass is thrown, their paths diverge because the defender must run towards the player with the ball. In zone coverage, the cornerback has to wait for the play to develop before determining who comes into his zone and becomes his responsibility. Ultimately, both of these graphs show a confirmation that we are not penalizing players for being in certain coverages while also maximizing the measurement period of their performance.

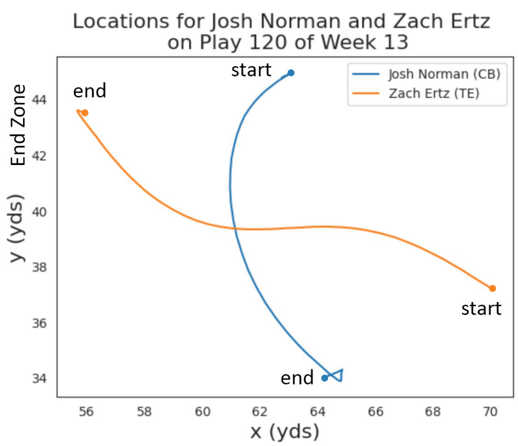

(a) Locations of a Receiver and Cornerback with the Defense in Zone Coverage.

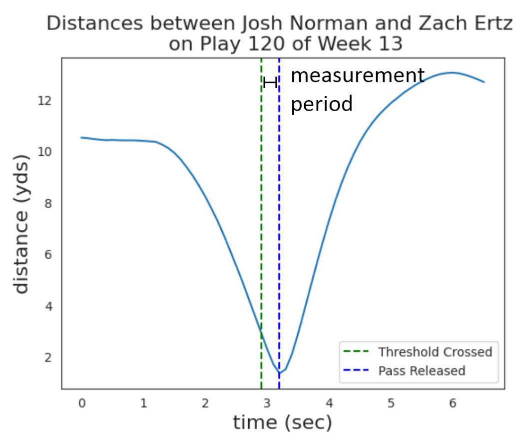

(b) Distances between a Receiver and Cornerback with the Defense in Zone Coverage.

Figure 3-4: Zone Coverage

3.5 Performance and Result Dataset Creation

Once I could identify the matchups for each cornerback and define an accurate metric to measure performance, the next step was to create a new dataset that encapsulated all cornerbacks' performance and the results of each of those plays. Using this new data set, we can find all adverse events for a specific cornerback and compare performance before and after those events.

For the creation of this dataset, I looked at each passing play and identified the cornerbacks on that play. For each of the cornerbacks, their performance was initially measured by both the mean distance and maximum distance from their receiver over the course of the measurement period. After creating the dataset however, I came to the realization that cornerbacks are very good at keeping their receivers close, so they rarely allow the maximum measurement distance to go above the three yard threshold.

	game_id	play_id	displayName	meanDistance	maxDistance	GotBeat
0	2018090600	146	Sidney Jones	1.7313354149697400	3.0	0
1	2018090600	190	Sidney Jones	3.3926836614073600	4.996048438516190	0
2	2018090600	320	Brian Poole	2.7232718800413700	3.0	0
3	2018090600	344	Robert Alford	1.690695628553290	3.0	0
4	2018090600	521	Ronald Darby	1.927766818266660	3.0	0
5	2018090600	521	Jalen Mills	2.253124822052680	3.0	0
6	2018090600	545	Ronald Darby	0.8627626884336130	3.0	0
7	2018090600	545	Jalen Mills	1.698864977758310	3.0	0
8	2018090600	545	Sidney Jones	0.9715923670505640	3.0	0
9	2018090600	636	Ronald Darby	1.503921945504770	3.0	0
10	2018090600	636	Jalen Mills	1.4615395761549000	3.0	0
11	2018090600	752	Brian Poole	1.5898607368701700	3.0	0
12	2018090600	776	Robert Alford	1.2354101766366800	3.0	0
13	2018090600	839	Robert Alford	1.2327159528362400	3.0	0
14	2018090600	889	Ronald Darby	1.695656011763330	3.0	0
15	2018090600	949	Sidney Jones	2.5580402815312900	3.0	0
16	2018090600	1037	Desmond Trufant	1.3540606299003400	3.0	0
17	2018090600	1037	Robert Alford	1.2272226164604100	3.0	1
18	2018090600	1061	Robert Alford	1.8446348510721900	3.0	1

Figure 3-5: Snippet of the Resulting Dataset with Results and Performance on Each Play.

In Figure 3-5, we can see a small piece of the larger data set. Due to the high performance of NFL cornerbacks, once the measurement threshold was crossed, the distance between the cornerback and receiver infrequently increased to a value higher than the threshold. In this particular snippet, we can see that row 1 is the only time where the distance goes below 3 yards and later reaches about 5 yards during the measurement period. Thus, I realized that having a maximum distance would not measure performance as well as using the average distance over the measurement period because there was almost no variability within the maximum distances. Since each player has the same measurement threshold, we will still have a fair evaluation across all players. We keep track of the specific game and play to help find the plays before and after each adverse event using the game_id and play_id labels. In addition, one game must be analyzed at a time because the players will likely have mentally recovered from any event in the week between two games. Each cornerback's mean distance and maximum distance on that play is also recorded in the row. The last column, labeled "GotBeat", represents whether the cornerback's receiver caught the ball and thus faced adversity. After creating this new dataset from all games in the 2018 NFL season, I could see whether a cornerback's performance changed following a poor result.

3.6 Validating the New Dataset

To confirm the accuracy of the "GotBeat" classification, I compared the percentage of time that each player was deemed to be at fault for letting the other team get a catch and their Pro Football Focus Coverage Score.[4] Pro Football Focus (PFF) is a widely accepted player performance scoring system and in this case was used to validate the dataset. Because some players only registered a small number of pass plays, I validated my results on players who registered at least 100 plays in the dataset. This way small sample sizes won't skew the percentages of the time players got beat.

In Figure 3-6, we can see that there is a negative correlation between the percentage of plays where a player gets beat and their PFF coverage grade. We would expect that players with lower PFF coverage grades would have a higher percentage of plays where their receiver catches the ball. Therefore, a negative correlation between our "GotBeat" calculation and the PFF coverage grade provides confirmation that the calculation was reasonably accurate. As mentioned previously, Figure 3-6 does not include low sample sizes. We exclude them because if some player has a very low number of eligible plays to measure, and happens to get beat on a few of them, we will see a disproportionate value for their percentage beat. There are similar excluded outliers with the opposite effect. Those who got beat zero times can attribute this to a low sample size for which they happened to not let their receiver catch the ball on any of those plays.

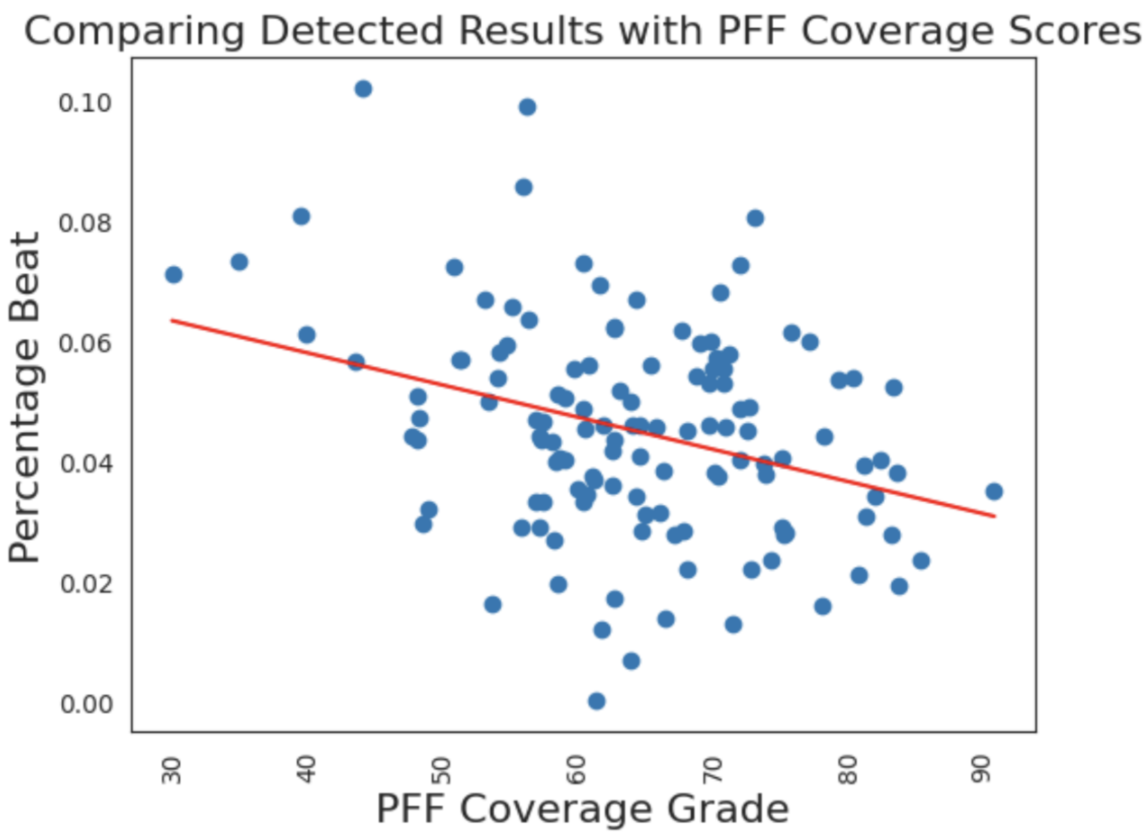

Figure 3-6: Scatter Plot of the Percentage of the Time that a Player was Detected as Getting Beat Compared to their PFF Coverage Grade.

Chapter 4

Analysis

4.1 Building a Grit Score

Now that we have our new dataset with calculated performance and results on each play, we can begin to build a Grit Score for each player. Using the new dataset, all times where any cornerback had their receiver catch the ball can be found by filtering for when the "GotBeat" parameter is 1 or true. We can then go through each of these and see if the cornerback's play changed or stayed the same following the event. To detect differences in play, we first take the average distance values on the 5 plays before and the 5 plays after the cornerback got beat. We call these P_Before and P_After respectively.

$$P\_Before = [pb_1, pb_2, pb_3, pb_4, pb_5]$$

$$P\_After = [pa_1, pa_2, pa_3, pa_4, pa_5]$$

In Figure 4-1, we can see a visualization of the timeline for the P_Before and the P_After arrays in the context of a real game. In this particular case, the cornerback, Stephon Gilmore, allowed his receiver to catch the ball at 6:13pm. The 5 eligible plays before he got beat are shown on the timeline in red, while the 5 eligible plays after he got beat are shown in blue. We can see that some neighboring plays are closer

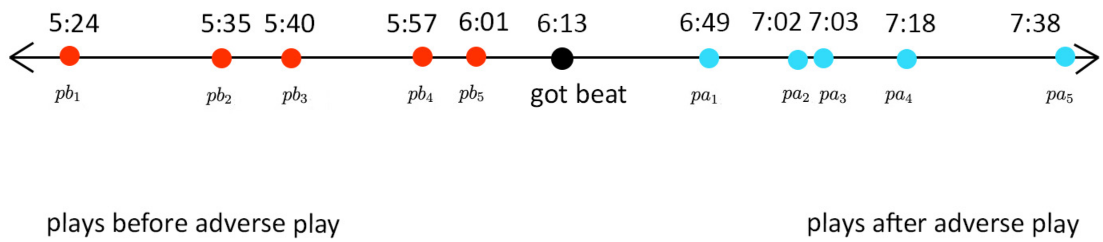

Figure 4-1: An Example Timeline of Plays Before and After a Cornerback (Stephon Gilmore) Gets Beat.

together while other neighbors can be more spread out. Because of events such as halftime, commercials, run plays, and having the team on offense create these gaps, there can be some variability in the timing between plays. Although these could have an effect on the cornerbacks's mental state, I continued using the play numbers because the next defensive passing play is the earliest time he can be impactful.

I also reversed the order of P_Before prior to taking the difference between the two, so we could see the change in performance for 1 play away from the bad outcome, 2 plays away, and so on. For example, since we use the 5 plays before the adverse event and the 5 plays after, if we were comparing the two arrays directly, then we would be taking the difference between the performance on a play 5 plays before and 1 play after. Instead by reversing the two arrays we can compare the same distance before and after one another. Now to find the differences in performance, we have:

$$\Delta_P = P\_After - rev(P\_Before)$$

$$\Delta_P = [\delta_1, \delta_2, \delta_3, \delta_4, \delta_5]$$

Subtracting the before performance array from the after performance array will give us an array of delta values for each timing where a positive value represents an increase in average distance. Therefore, the performance would decrease and we could say that the player has less grit. Players with the most grit would have unchanged performance and would have delta value of close to zero. Lastly, players with a negative Grit Score are inspired by the adverse event and are playing better after it happens.

After calculating the Δ_P values for each play, we still want a holistic view of each player rather than an anecdotal one. Thus, we take each of these values and add them to their own array according to their timestep. This gives us a distribution for differences in performance some number of plays after getting beat. By visualizing this distribution over the course of the ensuing plays, we can better understand how timing affects each player. That is, how quickly can they return to normal performance and whether or not there is a gradual return. We can see the visualizations for Jalen Mills, Jalen Ramsey, and Grant Haley below.

If we look at Figure 4-2a, we can see the distribution of the differences in performance before and after getting beat for Jalen Mills over the 2018 NFL season. We can see that these distributions stray towards the negative end of the spectrum with all averages being below 0. Since there is still variance in the data based on how many plays away from the event we are, we can take the average of all distance deltas to get a Grit Score. For this case, Mills clearly has a negative Grit Score which implies that he is inspired by an adverse event and thus plays better folllowing that.

Another interesting player to observe is Jalen Ramsey. Though he is widely regarded to be one of the most talented cornerbacks in the NFL, he is also known to be a more emotional and outspoken player. While this would lead many to believe that he would be heavily influenced by positive or negative events, we can see from Figure 4-2b that he is actually one of the most level headed players when it comes to keeping his performance steady. This would be a situation where the anecdotal evidence and perceived antics on the sideline do not correlate to his play. Based on this analysis, Ramsey had a considerable amount of grit throughout the entire season.

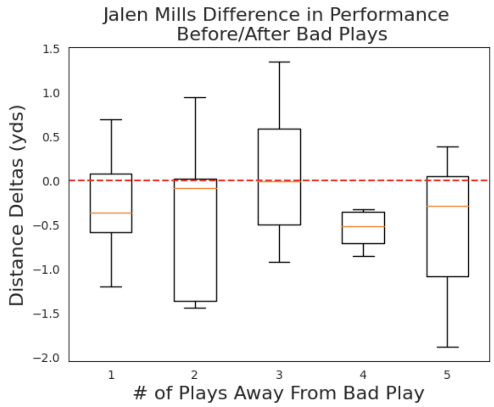

(a) Jalen Mills' differences in performance plays before and after getting beat.

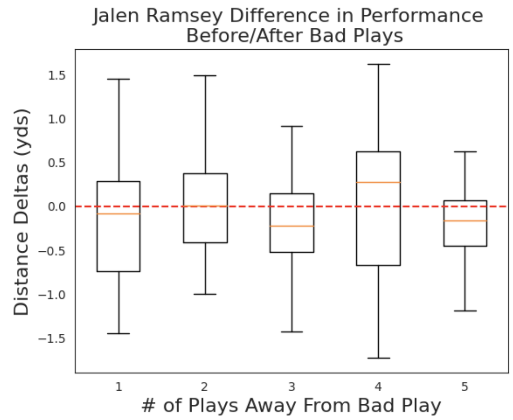

(b) Jalen Ramsey's differences in performance plays before and after getting beat.

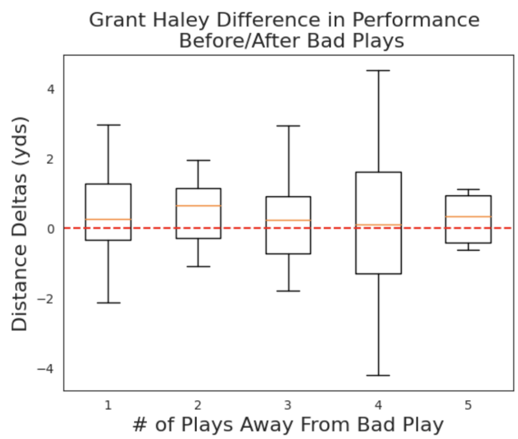

(c) Grant Haley's differences in performance plays before and after getting beat.

Figure 4-2: Player Distance Deltas

Lastly, we look at a lesser known player such as Grant Haley seen in Figure 4-2c. The distance between him and the receiver that he is guarding generally increases after he gets beat. This is a unfortunate outcome for the player as it means that they have less perseverance through a bad result. Because of this, we would classify Haley as having low grit which would be a less desirable trait for NFL teams.

To create one Grit Score for each player, we can take average deltas for each player's distribution and take the average of that to see the average change in distance before and after an adverse play. As mentioned above with each example, we will classify the players with a Grit Score of near 0 as those with the most grit.

Cornerbacks with more positive Grit Scores are those that perform worse following the adverse plays and thus have the lowest grit. Those with negative Grit Scores are inspired by adverse plays, however it does raise a question as to why they aren't reaching their full potential for the entire game.

4.2 Choosing a Measurement Start Time Threshold

Once the process for creating a Grit Score is well defined, I compared options for measurement threshold to most accurately capture performance. Following the guidelines that we created for the measurement period, a play is only eligible to be in the dataset if the distance threshold was crossed before the ball was thrown or when the measurement period would be positive.

$$Measurement\ Period = Time\ of\ Pass\ Release - Time\ Where\ Threshold\ Crossed$$

Choosing a distance threshold has many tradeoffs including the number of plays in the dataset that will be eligible, as well as having the most accurate possible measurements on each play. For this research, I examined measurement thresholds of 2 yards, 3 yards, and 5 yards.

2 Yards	3 Yards	5 Yards
48%	61%	79%

Table 4.1: Percent of Overall Plays Measured per Threshold

In Table 4.1 we can see that increasing the threshold for which we begin measurement yields a higher number of plays that can be included in our performance and result dataset. A lower number of plays included can lead to higher emphasis being placed on particular plays and thus a less complete evaluation. However, increasing this threshold also can include plays where the cornerback is never actually guarding the closest receiver to them, so their mean distance can be inflated.

Figure 4-3, shows the differences in average grit score when changing between particular thresholds. We can see that the changes in grit value from a threshold of 2

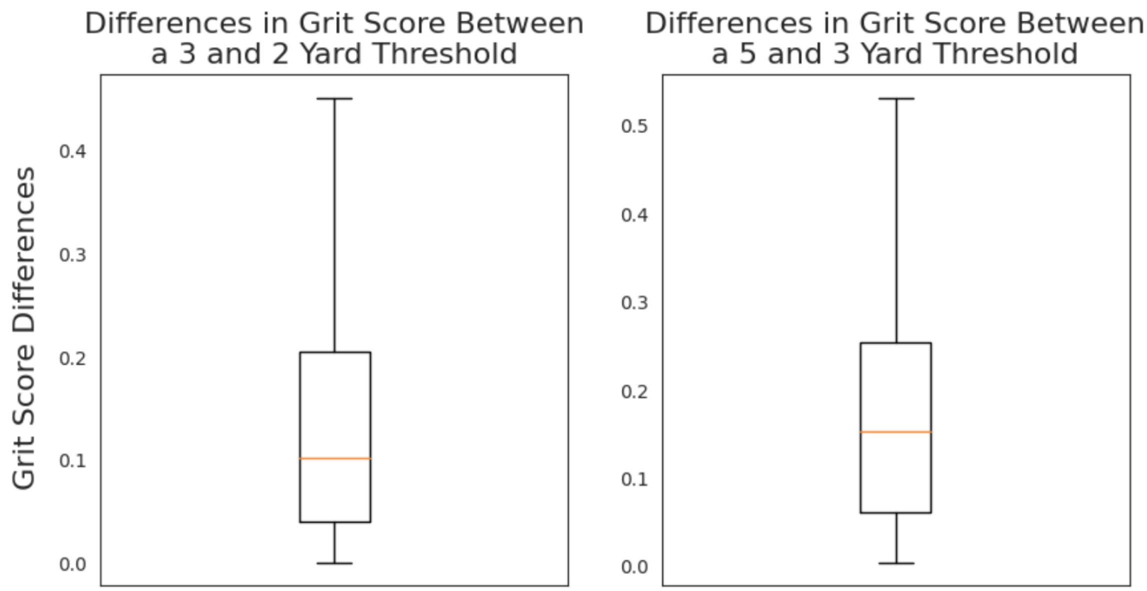

Figure 4-3: Effects on Changing the Measurement Thresholds on the Grit Metric

yards to a threshold of 3 yards is smaller than those from 3 yards to 5 yards, showing that using a threshold of 5 yards is inflating the distance values and thus the grit values. From these two figures, we can see that choosing a threshold of 3 yards makes the most sense for this calculation.

Chapter 5

Results

Now that the Grit Score is solidified, we can begin to understand which players are performing with more grit than others and whether the preconceived notions of them are correct. As mentioned previously, a Grit Score closest to zero is an ideal score because this represents the most consistent play. Most cornerbacks in the NFL display a significant amount of grit which is likely why they have made it all the way to the pro level.

Since all players in the dataset were not able to play a significant amount of time, we will zoom in on the top 100 most utilized (played in the most plays) players for the 2018 season. Figure 5-1 shows the outliers from these top 100 players, specifically the highest 10 and lowest 10 from the dataset. Immediately, Damarious Randall's score of 0.591 pops out as the most positive Grit Score. He and many of the other cornerbacks with positive Grit Scores are not the most immediately recognizable players and 7 out of 10 are not currently on an NFL roster.[11] Conversely, only 3 out of the 10 are no longer playing in the NFL. The negative Grit Score players are generally still in the NFL with bigger names such as Patrick Peterson, Jalen Mills, Marlon Humphrey, and James Bradberry. Peterson had the most negative Grit Score and is a 3x All-Pro cornerback. Due to the difference in longevity between the two groups, we can see that NFL coaches and general managers prefer the cornerbacks that respond well to adverse events.

Figure 5-1 also includes the 10 players with Grit Scores closest to zero. This group of players are less talented as the players with the most negative Grit Scores, however most of them have had long careers in the NFL or are still currently on a roster. Eight out of these 10 players are currently on an NFL roster and the two players who are not, Johnathan Joseph and Desmond Trufant, had 16-year and 9-year careers respectively. These are far longer than the average cornerback career length of 2.94 years.[12]

In addition, none of these cornerbacks are particularly known for having elite talent, so it is likely that their grit has been a significant factor in keeping them in the league. An interesting example of this would be a cornerback such as Eli Apple. While he is known for being a mediocre at best cornerback, he also talks a lot of smack to opposing teams through interviews and social media. No matter the result of the game, he is always very outspoken about his performance and usually speaks negatively of others. He would seem to have low grit based on public perception of his personality, however his performance-based Grit Score says otherwise. A situation such as this might be a time where our Grit Score can disprove preconceived notions about a player's mental perseverance.

5.1 Grit Scores vs. PFF Coverage Grades

Now that each of the cornerbacks have a Grit Score, we can compare this score to their overall PFF coverage grade in Figure 5-2. Again, since the overall dataset includes a few outliers who have not played as much, we can zoom in on the top 100 players in the dataset. Here we see two red dotted lines for reference. The horizontal line is along the line of a Grit Score of zero while the vertical line represents the average PFF coverage grade from 2018. Using these as reference, we can classify the cornerbacks into six different player types.

In the top right corner are the higher overall performing cornerbacks with low grit. This likely means that these players rarely allow their assigned receiver to catch the ball, however when that does happen, their play suffers. Stephon Gilmore, the highest PFF rated cornerback, is a prime example of this type of player. Cornerbacks

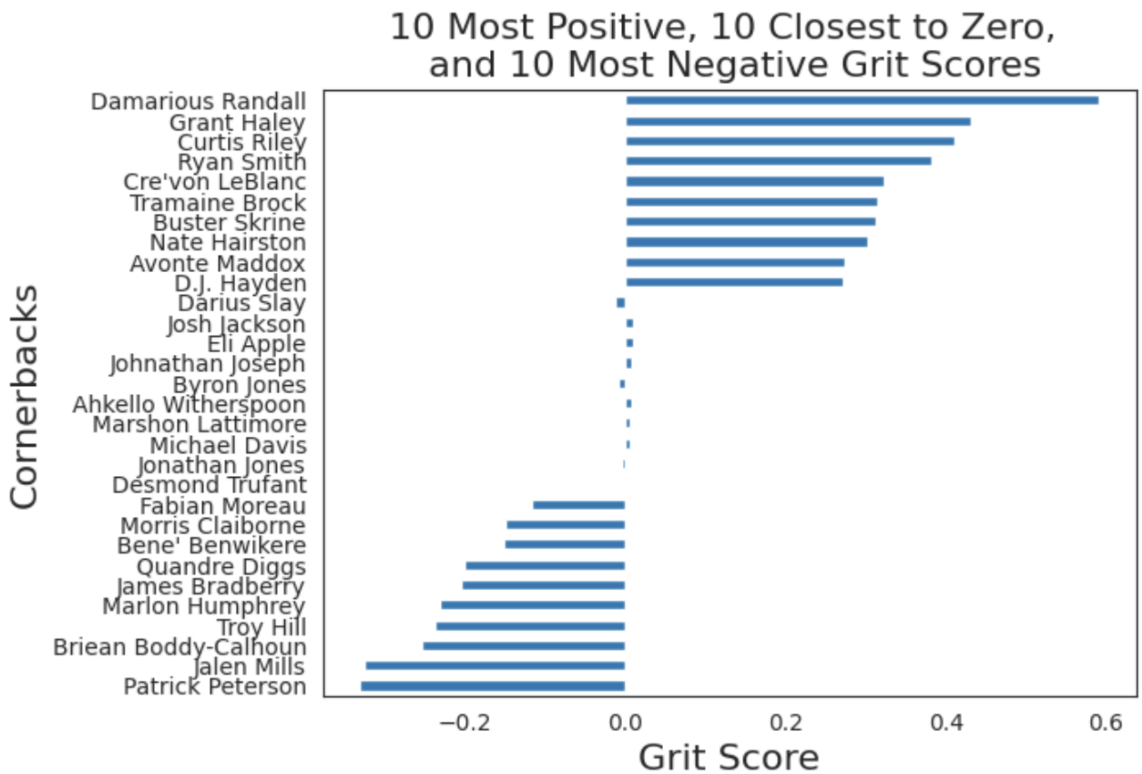

Figure 5-1: The 10 Most Positive, 10 Closest to 10, and 10 Most Negative Grit Scores for the Top 100 Most Utilized Players.

with high PFF coverage grades and Grit Scores near zero, such as Jason McCourty, are considered as the ideal cornerbacks for a team. They will always perform at a consistently high level, no matter what happens in a game. For the last of the cornerbacks with high PFF grades, those with negative Grit Scores typically play well during the course of a game, but need a bad circumstance to achieve their full potential. Patrick Peterson stands out the most in this group with a very negative Grit Score and very high PFF coverage score. While all of these players mentioned are considered star players, their consistency in the face of adversity is far different from one another.

We can do a similar analysis with the cornerbacks of lower PFF grades as we did with the higher performers. The lower left quadrant has the lower performers who can play better after getting beat. These are less talented players who might not come with the right mindset at the beginning of a game, but get back on track as soon as they get beat. While there are not many players that fit this category well, Jalen Mills had a relatively low coverage score and a very negative Grit Score. Lower PFF grades with near zero Grit Scores, such as Ahkello Witherspoon, are likely those who have strong mental fortitude, but don't quite have the talent to succeed in the NFL. Lastly, cornerbacks with low PFF grades and positive grit scores would be least desirable for NFL coaches and general managers. Nate Hairston is an example of this type of player, who has bounced around between 4 different teams in his 5-year career. These are players that see both a considerable talent gap, and cannot persevere as much as the rest.

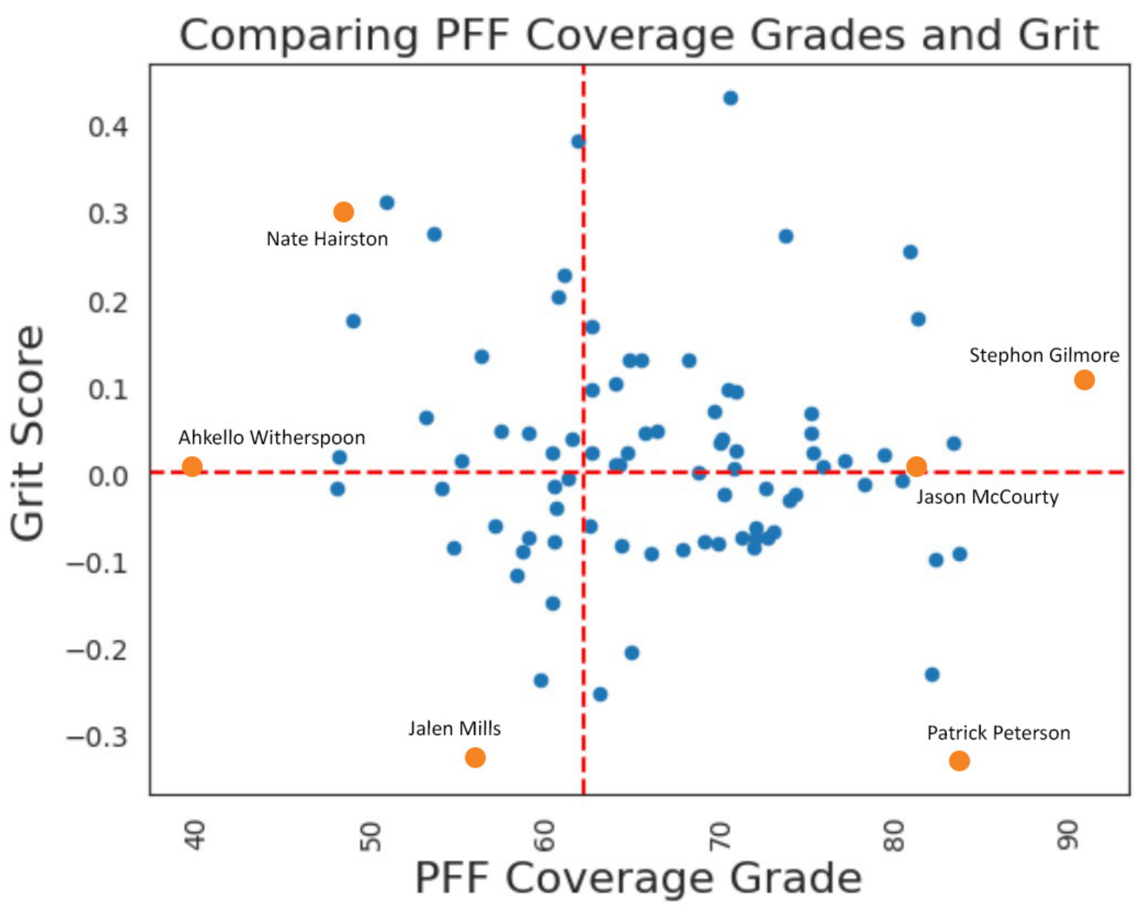

Figure 5-2: Comparing the Top 100 Most Utilized Players Grit Score vs. Their PFF Coverage Score.

Chapter 6

Future Work

Since finding a performance-based grit metric is relatively untapped, there are numerous topics that can be explored in this area. One could start by trying different approaches to the same project such as using neural networks, or a different performance metric to measure a player's success in the league. While I believe that this statistical and play-by-play approach is likely the most accurate, it is always important to have comparisions between methods.

In addition, one could perform this analysis to understand the state of NFL cornerbacks and how they are changing over time. An interesting project could include exploring the most recent seasons and comparing them to the 2018 season to see if players are getting more gritty over time. This could be a way to understand how management in the league is weighting overall talent against mental fortitude. It could even be taken a step further to compare salaries to grit as well. Could there be a correlation between the players of high grit and those who are making the highest salaires? While it might be assumed that this is the case, other circumstances such as an expiring contract, overall talent, and player demands could change these results.

Cornerbacks are typically viewed as the most emotional and variable in their play, but how would Grit Scores vary accross positions? One could provide an interesting argument for how to measure results-based adversity while also determining what play-to-play performance looks like. For example, a quarterback's adverse result might be throwing an interception and play-to-play performance could be based on their pass accuracy and receiver selection. Once a Grit Score is found for each position, one could explore the importance of having grit for each position and rank these.

Lastly, this new Grit Score could be a source of validation for the new cognitive tests that are constantly being created for college player analysis. With access to college football tracking data, one could perform the same statistical analysis and compile Grit Scores for each of the NFL Draft candidates. A high correlation between the scores on the cognitive tests as well as closer to zero Grit Scores would likely mean that these tests are highly effective. No correlation on the other hand, might imply that there is an issue with one of the two metrics. In total, there is a multitude of future projects to explore grit in football as well as any other sport.

Chapter 7

Conclusion

The NFL is a multi-billion dollar league where money and team success hinges on player performance. Therefore, understanding the value of each player that a team puts on the field and how much they will contribute to that team's success is vital. NFL teams already have accurate data and analysis for physical attributes on the field, but have yet to find a way to measure most of the mental aspects of the game, specifically grit. Through this Grit Score calculation, we can better understand the mental toughness of these NFL players during the game itself rather than basing it off of perceived reactions to unfortunate events.

In the process of creating the Grit Score, I verified my identification of when a cornerback was at fault for an opposing receiver catching the ball and isolated the measure of a cornerback's performance on a play to be independent from the opposing quarterback. The distribution of differences in play over time allowed us to see a more specific profile of each player, giving an insight as to how an impactful event of allowing a receiver to catch the ball will affect them over time. Since the play-to-play variability of football caused the differences before and after the play to also be variable, we aggregated these values for our final Grit Score.

Players with highly positive Grit Scores generally do not last as long in the NFL as those with a near zero Grit Score or negative Grit Score. This follows common assumption as these players are considered to have the least amount of grit and thus are less desirable. While many of the top players, such as Jalen Ramsey, had Grit

Scores near zero, we also saw a few elite cornerbacks that had negative Grit Scores. Top players with negative Grit Scores might be less consistent in their play, but it is far more desirable to improve in the face of adversity rather than the opposite. Through the process of this research, I have determined that players with a near-zero grit scores have the most grit. Negative grit scores show that a player might not be playing to their full potential, but they typically perform better after a poor result. A positive grit score means that a player cannot persevere through adversity as well, and their ensuing plays are hindered by it. With these in mind, I can say that those with the most talent, or highest PFF coverage score, and near-zero Grit Scores are the players that NFL management should target due to their top-tier talent and consistency.

Bibliography

[1] Angela Duckworth and Patrick Quinn. Development and validation of the short grit scale (grit-s). *Journal of personality assessment*, 91:166–74, 04 2009.

[2] Angela L. Duckworth, Christopher Peterson, Michael D. Matthews, and Dennis R. Kelly. Grit: Perseverance and passion for long-term goals. *Journal of Personality and Social Psychology*, 92(6):1087–1101, 2007.

[3] Lauren Eskreis-Winkler, Angela Duckworth, Elizabeth Shulman, and Scott Beal. The grit effect: predicting retention in the military, the workplace, school and marriage. *Frontiers in Psychology*, 5, 2014.

[4] Pro Football Focus. Pff's top 101 players from the 2018 nfl season. https://www.pff.com/news/pro-pffs-top-101-players-from-the-2018-nfl-season, February 7 2019. Accessed: May 7, 2023.

[5] Line From, Dorthe Kirkegaard Thomsen, and Martin Hammershøj Olesen. Elite athletes are higher on grit than a comparison sample of non-athletes. *Scandinavian Journal of Sport and Exercise Psychology*, 2:2–7, Feb. 2020.

[6] Linda S. Gottfredson. Mainstream science on intelligence: An editorial with 52 signatories, history, and bibliography. *Intelligence*, 24(1):13–23, 1997.

[7] The Painted Lines. Football 101: Cover 2 Zone. http://thepaintedlines.com/football-101-cover-2-zone/, 2023. Accessed: April 26, 2023.

[8] Brian Lyons, Brian Hoffman, and John Michel. Not much more than g? an examination of the impact of intelligence on nfl performance. *Human Performance*, 22(3):225–245, 2009.

[9] Emily A. Nothnagle and Chris Knoester. Sport participation and the development of grit. *Leisure Sciences*, 0(0):1–18, 2022.

[10] The Phinsider. Football 101: Defensive Cover Schemes, aka How a Quarterback Reads a Defense. https://www.thephinsider.com/2016/6/27/12039106/football-101-defensive-cover-schemes-aka-how-a-quarterback-reads-a-defense, 2016. Accessed: April 26, 2023.

[11] David Randolph. David randolph nfl game log. https://www.pro-football-reference.com/players/R/RandDa00/gamelog/, 2021. [Accessed: May 7, 2023].

[12] Statista. Average player career length in the national football league, n.d.

[13] Dennis Waldon Van Berkum. Selected traits for predicting success in football as determined by college football coaches. Electronic Theses and Dissertations, Paper 4858, 1975.