PRACTICAL APPLICATIONS OF SPACE CREATION FOR THE MODERN NFL FRANCHISE

A PREPRINT

Alexander C. Stern
Undergraduate Departments of Statistics and Computer Science
University of Virginia
Charlottesville, VA 22904-4135
acs4wq@virginia.edu

December 20, 2019

1 Introduction

Establishing the run: a fundamental tenant of every NFL offense. Even in today's pass-heavy league, many offensive coordinators would prefer to run the ball when they can to increase their time of possession and not risk giving up an interception. This is indicated by the fact that the vast majority of teams have chosen to run $> 40\%$ of the time on early downs in 2019 [1]. Therefore, running the ball well is a key part of a successful offense. That said, this should not be confused with the fallacy that "running the ball more times makes a successful offense." The relationship between the number of carries in a game and a greater percentage of games won is a prime example of correlation falsely indicating causation. Teams that are already winning run the ball with greater frequency to waste the clock and ensure ball security.

When the ball carrier runs out of the backfield, the primary job of the other offensive players is to create space for the runner to move through. This is done in one of two ways: pulling defenders with you (away from the runner) or blocking. Both of these actions are performed with the intent of creating space ahead of the runner. This is especially true for members of the offensive line, whose jobs on the vast majority of designed run plays are to create and own the initial space through which the runner will move. 2018-2019 Big Data Bowl participant Kyle Burris [2] succinctly describes space ownership as:

> "... a space is owned by the player who can beat every other player to that space."

The main function of this paper is to evaluate the creation of space by the offense and the limitation of that endeavor by the defense. While the conceptual idea of space creation has existed for as long as the game itself, concrete formulas for identifying space creation are few and far between. Recently, some research papers have touched on the topic, most deriving their ideas from Bornn and Fernandez [3]. However, this paper is deeply rooted in concepts from professional soccer. Control of space over a much larger area of the field, usually referenced in "thirds," is required while being close to the player in possession of the ball is not. The paper describes a method for evaluating space ownership based on the player's velocity and distance to the ball as factors of a bivariate normal density function. In football, a player's distance to the ball has significantly less to do with their concrete ability to own space at any given point on the field.

In this paper, details are provided for the implementation of a football-specific space creation model using an algorithmic approach rooted in logarithmic transformations of statistical distributions and a multivariate normal density function. Further, it includes an array of useful and practical applications. Space ownership of each point on the field by each player will be defined as a function of that player's speed, direction, and distance to that point. Essentially, it represents the ability of that player to turn towards any point p and get there while running at their current speed, relative to the other 21 players on the field. The modeling of space ownership in football could materially affect the design of offenses, situational play calling, player grading across a wide range of positional groups, and front office decisions for years to come.

2 The State of NFL Statistics

Before getting too far into this paper, the way mainstream NFL player statistics are collected and analyzed needs to be examined. Currently most statistical categories are the result of simple counts: how many yards a running back accumulates in a game, a quarterback's completion percentage (the number of completed passes divided by the number of passes attempted), the number of sacks a defense produces, etc. Thus far, these stats have provided a helpful indication of which players are better than others but they only go so far. 14 games into the 2019 NFL season, the Jacksonville Jaguars have allowed 5.1 rushing yards per attempt and the New York Jets have allowed 3.3 [4]. In Week 2 this year, Texans running back Carlos Hyde rushed for 90 yards on 20 carries (4.5 yards per carry) against the Jaguars [5]. A stat line like this would be more than respectable on any given Sunday, but in this case, Hyde ran for 0.6 yards less per carry than the average allowed by the Jaguars. Coupled with the fact that the Jaguars allowed Colts running back Jonathan Williams later in the season to run for 116 yards on just 13 carries (8.9 yards per carry), Hyde's performance becomes far less impressive. In Week 8 this year, Jaguars running back Leonard Fournette ran for a modest 76 yards on 19 carries (4.0 yards per carry). At 0.7 more yards per carry than the average allowed by the Jets, Fournette's performance was extremely impressive. Conversely, in a hypothetical situation, if Fournette had run for 4.0 yards per carry against the Jaguars defense, the performance would have been abysmal. Judging players not just based on counts in certain categories but instead on where those counts fall in the "distribution of counts" allowed by a team across a season paints a much better picture of what went on in a game and who is performing above or below expectation given a certain opponent. When grading players and making acquisition decisions as a front office, this sort of ranking system would allow teams to easily identify players who truly excel and out-perform their peers at their position.

3 Space Ownership

3.1 The Data

The data used for this research is from the NFL's 2^{nd} Annual Big Data Bowl competition. It consists of player tracking data at the point of hand-off for thousands of running plays over the course of the 2017-2018 and 2018-2019 NFL seasons. This includes a variety of things from player speed, acceleration, direction, position, etc. The data set required extensive cleaning and feature engineering to derive predictors that would be helpful in calculating player ownership of space on the field at the point of hand-off. This includes, but is not limited to: manipulating player position and velocity to find player trajectory and projected future position, classifying play types, standardizing player direction, classifying defensive alignments, clustering runs by ball carrier trajectory, and filtering out plays that occur in garbage time or too close to the goal line to qualify for comparison to standard running plays. Since player movement as the play develops is extremely difficult to speculate about due to the highly dynamic nature of the game, this paper will focus on player intentions at the time of hand-off since this is the only granular data that was made available. The moment the hand-off occurs, the defense is fully aware that the offense is running the ball. Some defenders have been dragged downfield by receivers and offensive linemen are moving in their predetermined motions to block in front of the intended path of the rusher. The amount of space they create for the runner is one of the main actions this paper will look to evaluate. The runner, at this time, is likely still moving in the path predetermined by the original play call and not yet reacting to the defense. This is not necessarily true for counter plays, but at the point of hand-off, the runner's velocity vector is usually already headed in the countered direction. Soon after the hand-off, the runner will have to react to the defense's convergence on him but for now these are all safe assumptions about player intentions and movement. What about those special few known as "patient runners" such as Le'Veon Bell? While these players may wait for more space to be created as the play continues, they still must initially move through some amount of the space created for them at the time of hand-off to ensure they aren't tackled behind the line of scrimmage.

3.2 Calculating Space Ownership

By adapting Bornn and Fernandez's [3] calculation of space ownership to the game of football and the player tracking data provided by the competition, a player's "ownership" of each point on the field can be calculated using the probability density function of the multivariate normal (Gaussian) distribution. Do's research on the relationships between Gaussian distributions [6] defines the probability density function of any multivariate Gaussian distribution as the product of the probability density functions of independent Gaussian distributions[1].

$$p(x; \mu, \Sigma) = \frac{1}{(\sqrt{2\pi})\sigma} e^{-\frac{1}{2\sigma_1^2}(x_1-\mu_1)^2} * \frac{1}{(\sqrt{2\pi})\sigma} e^{-\frac{1}{2\sigma_2^2}(x_2-\mu_2)^2} \tag{1}$$

[1]All assumptions were verified and skewed distributions were appropriately transformed to normal distributions.

As noted in Section 1, the distribution assesses how quickly a player could get to each point on the field compared to the other 21 players on the field at the time of hand-off using their current position and velocity–$speed \times direction$. For any player, these values are maximized as he travels faster; becomes closer to the point, p, on the field; and as the angle his position, projected direction, and p create gets closer to zero.

Another idea adapted from Bornn and Fernandez [3] was to define team ownership of p by applying the sigmoid activation function

$$\Phi(z) = \frac{1}{1 + e^{-z}} \tag{2}$$

to the net difference between the summed ownership of p by each defensive player subtracted from the summed ownership of p by each offensive player on the field. A value closer to zero indicates complete defensive control of that space on the field, and vice versa a value closer to one indicates complete offensive control of that space on the field. Ideally, runners should travel through the space with the highest level of offensive ownership as they move into the defensive backfield in order to travel as far downfield as possible before they are tackled by a defender.

3.3 Player Grades

After each player's ownership of every point on the field in each play has been calculated, player grades were calculated. Rushing zones were defined as a triangle formed by the location of the rusher at the point of hand-off and lines projecting from his direction angle +/- 45° up to four yards past the play's line of scrimmage. Angles that were too steep were cut off at the sideline four yards past the line of scrimmage. Each player's ownership of space within this rushing zone was calculated as the sum of his ownership at each point within the triangle. After consulting with Division I coaches and reviewing available work in the field, Offensive Line duties further on into the play lie in a gray area. However, the first four yards past the line of scrimmage provides a good baseline for determining the extent of the impact of offensive blocking for the runner. The amount of initial space created has a clear, positive impact on the chance the runner has of finishing his run at a point further down the field (see section 3.5). This implies that, to a certain extent, even without taking into account who the runner is, if he makes it at least four yards past the line of scrimmage, the blocking was above-par and if he doesn't then it was below-par. With a more comprehensive data set, this approach could be taken further to evaluate runner's abilities to move within the space created for them, but for the purposes of grading offensive linemen: runner ability, decision-making skills, and speed does not need to be taken into account.

Figure 1: James White rushes for 5 yards

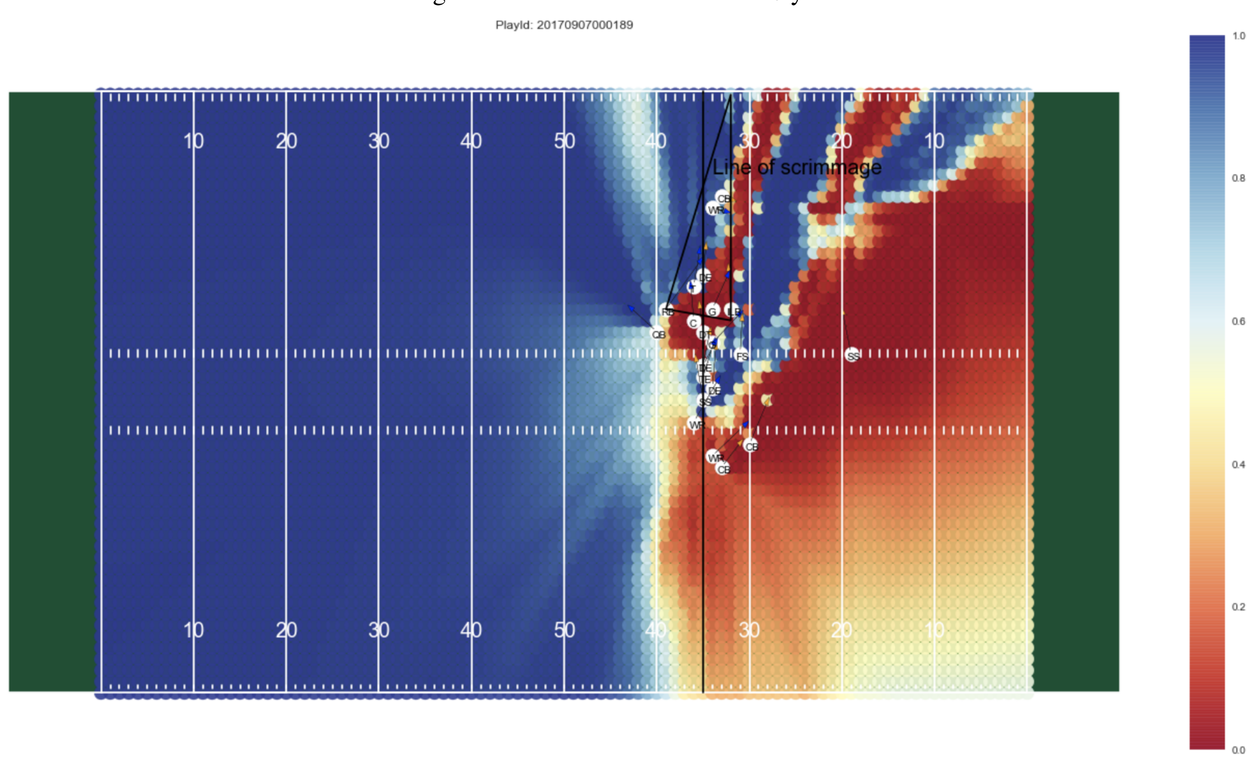

In the plot above, each white dot represents a player on the field and is labeled with their position abbreviation. Arrows extending from players represent their velocity vectors which are the result of their current speeds and directions. In black is the line of scrimmage and the rushing zone described above for this play. Points with a value closer to 1.0 are blue and those closer to 0.0 are red. Blue indicates total offensive ownership of that point and red indicates total defensive ownership of that point. The ball carrier, White, moves toward the left side of the field through a relatively blue area. This indicates that the New England offense owns the majority of the space White initially plans to move through and therefore one would expect him to be successful. Evidently he is and White is initially able to run free. The play then breaks down and he is tackled for a five yard gain. The space created by the offense combined with the play design allowed White to gain the initial four yards past the line of scrimmage. This perfectly encapsulates one of the main purposes of this paper: to prove that the quantified amount of offensive space owned in front of the runner has an incredible and direct impact on the ability of a runner to move at least four yards past the line of scrimmage. Unequivocally, this is the goal of almost every designed run play in the history of football.

3.4 Standardizing Player Grades

As discussed in Section 2, one of the main problems with current player grading systems is that they fail to account for the situation in which each event occurs. Football is an entirely situational game and each time something happens, the question needs to be asked: "How much better or worse was that action compared to other players facing a similar situation?" If team A plays in a division with three teams that are all do a poor job stopping the run, they will play each of them twice throughout the season and their running back's stats will look phenomenal compared to others. However, this doesn't necessarily mean he's an impressive running back. Therefore, plays were clustered by defensive formation, run direction, and the specific team on defense. The mean and standard deviation of rushing zone space ownership was calculated for each player position at each combination of these situational factors. Player ownership was then standardized by these values to help give each player an unbiased grade of their abilities relative to their peers in similar situations. For example, one can assume that if the runner is going to the outer left side, the left tackle would inherently create more space in the rushing zone than the right tackle. Therefore, if you are a right tackle creating space for a ball carrier moving to the left end, you can still be properly rewarded in this grading system for doing a superior job creating space than your peers who may slack off given the less important nature of their job on that play. In addition, players facing defenses that fare better against the run are expected to own less space in the rushing zone than those facing more porous rushing defenses. Similarly, players facing five defensive linemen on a play are expected to own less space than players facing three defensive linemen on a play. Standardization by all combinations of these groups allows the algorithm to account for all of these situational factors and compare player performances in very similar play situations. After player scores on each play are standardized, individual player scores are averaged across all snaps played. Only players with greater than ten snaps played were included in this analysis to ensure the variances across players remained realistic. The resulting score can be interpreted as the player's average number of standard deviations above or below the mean in the distribution of player scores from the same position in similar play situations.

3.5 Relationship to Yards Gained

Whether or not a runner gets at least four yards past the line of scrimmage depends on two things: how much space is created for the runner to move through and how efficiently he runs through that space[2]. The space model will attempt to quantify the former. A highly predictive random forest ensemble model was fitted using extreme gradient boosting (XGBoost) to predict the yards gained by a rusher based on a number of predictors describing player position and velocity, offensive scheme, defensive setup, game situation, down, distance, offensive team, defensive team, and standardized offensive ownership of space within the rushing zone defined above. Random forest models allow for predictor importance to be evaluated and compared across all decision trees, making it ideal for this type of analysis. Offensive ownership of space was one of the top predictors in order of importance, with an average information gain of 9.329 across all decision trees within the random forest model. In addition, it outperformed a completely random variable by 259%. This implies that space created by the offense is an extremely important and consistent factor in running the ball well, even when the identity of the runner himself is not taken into account. However, as a stand-alone predictor, it is not able to explain 100% of the variance in yards gained per play. This means that a portion of that variance relies on the running back's personal ability to create space for himself, and more importantly, move successfully through the space created for him by the offensive line and other members of the offense.

[2] The runner can also create additional space for himself by changing his direction or speed, juking, spinning, etc.

4 Applications

4.1 Offensive Linemen

Standardized average player grades allow players to be concretely ranked by position based on how well they create space on running plays compared to their peers. This will be referred to as a player's *SpaceGrade*. Currently, many generic player grades are based off counts of binary events happening such as pressures allowed, sacks, whether or not the pocket collapsed, etc. Play-by-play comparisons of how one player does vs. their peers in similar play situations is a much more holistic and substantial method for analyzing player performance.

Figure 2: Tackle Andre Smith blocks for Running Back Joe Mixon. Mixon ran through the space Smith created and gained 19 yards on the play.

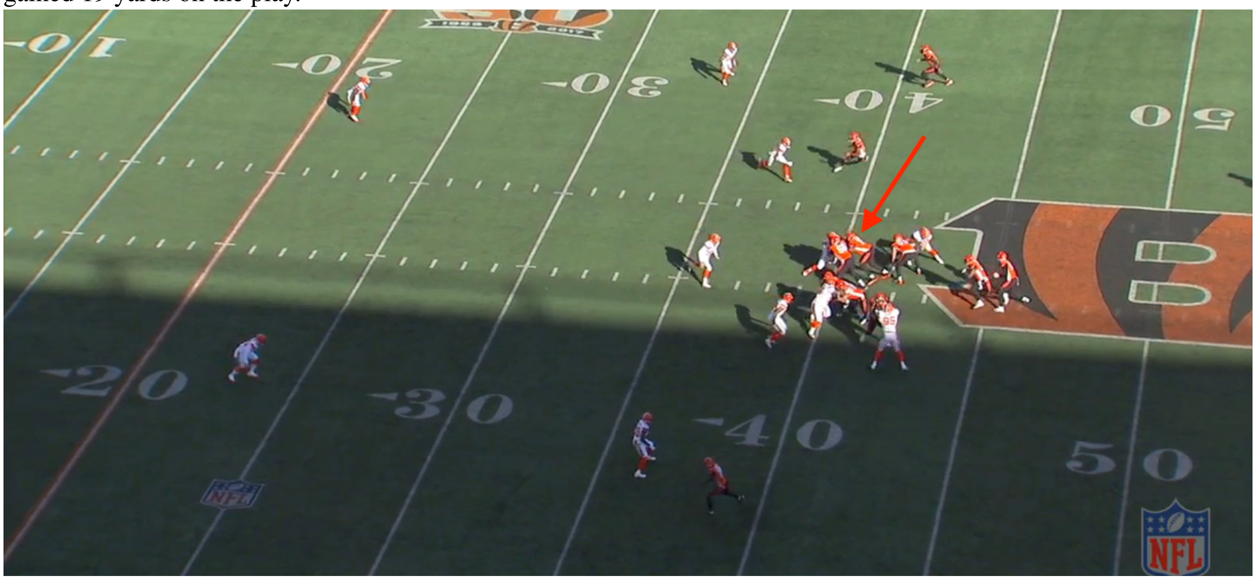

Smith, one of the algorithm's highest graded tackles, does a tremendous job blocking on this play. This was one of the largest instances of space creation by a lineman the algorithm found among all plays on which it was trained.

After these grades were calculated, players and their respective *SpaceGrade* were matched up with their Pro Football Focus *RunBlock* grades. As expected, the two grades had mild correlations for centers, guards, and tackles. A correlation is the numerical measure of the statistical relationship between two variables. This indicates that space created is tangentially involved with calculating run blocking grades for offensive linemen currently, but likely isn't being measured quantitatively. In addition, many offensive linemen grades at the moment fail to take into account specific individual performances and actions on plays, using play-level response variables such as yards gained as a baseline metric. This is very similar to the plus-minus idea in basketball. Joseph Sill [7] and a plethora of others have gone into detail as to why team-wide statistics such as this can be very misleading. They fail to account for the situation the offense is in and who the opponent is, among other confounding factors. One of the ideas behind the space creation algorithm is to recognize that even when the team doesn't operate effectively, individual linemen who do their jobs effectively and more successfully than their peers will be properly rewarded by the grading system. For example, Smith is ranked 3^{rd} in *SpaceGrade* for tackles despite ranking 54^{th} at the position in run blocking, the largest discrepancy among all tackles.

Figure 3: Center Max Unger blocks for Running Back Alvin Kamara. Kamara gained only 2 yards on the play.

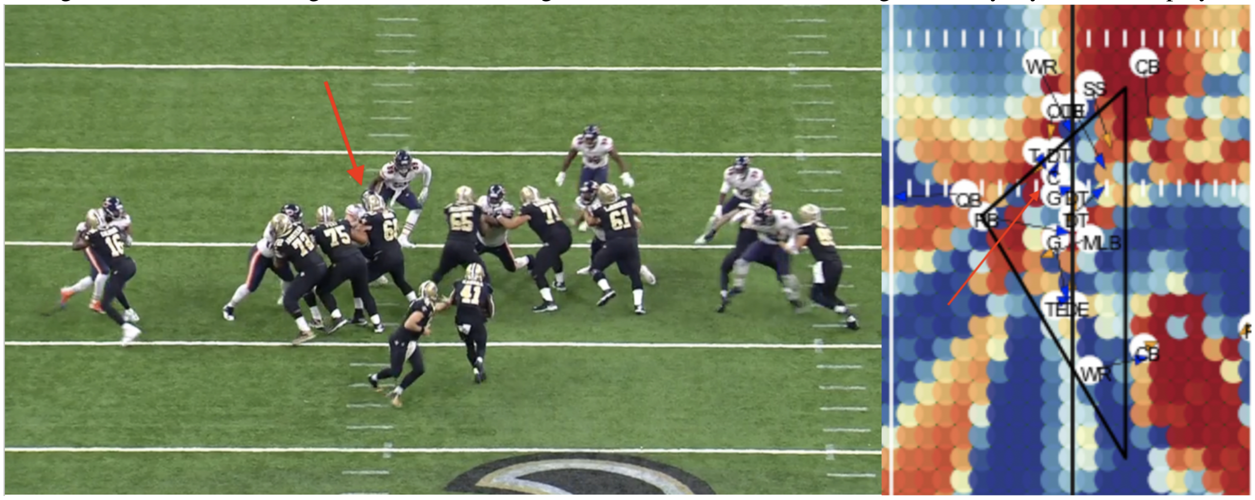

Unger claims a top five *SpaceGrade* ranking for centers despite being the 31^{st} highest graded run blocker by Pro Football Focus, the largest discrepancy among all centers. This play only gained 2 yards due to poor performance from the right half of the offensive line, as well as the tight end. However, Unger performs his duties quite well. Plays like this are a perfect example of the current discrepancy in the evaluation of offensive lineman and their actual performances.

4.2 Tight Ends

Similar to the offensive line, the tight end is often required to create space for the runner as well. In line with the analysis above, tight end scores were calculated and compared to the player's respective run blocking rank. The correlation between *SpaceGrade* and run blocking rank for tight ends was similarly mild.

Figure 4: TE Dallas Goedert helps block for Running Back Josh Adams. Adams gained 2 yards on the play.

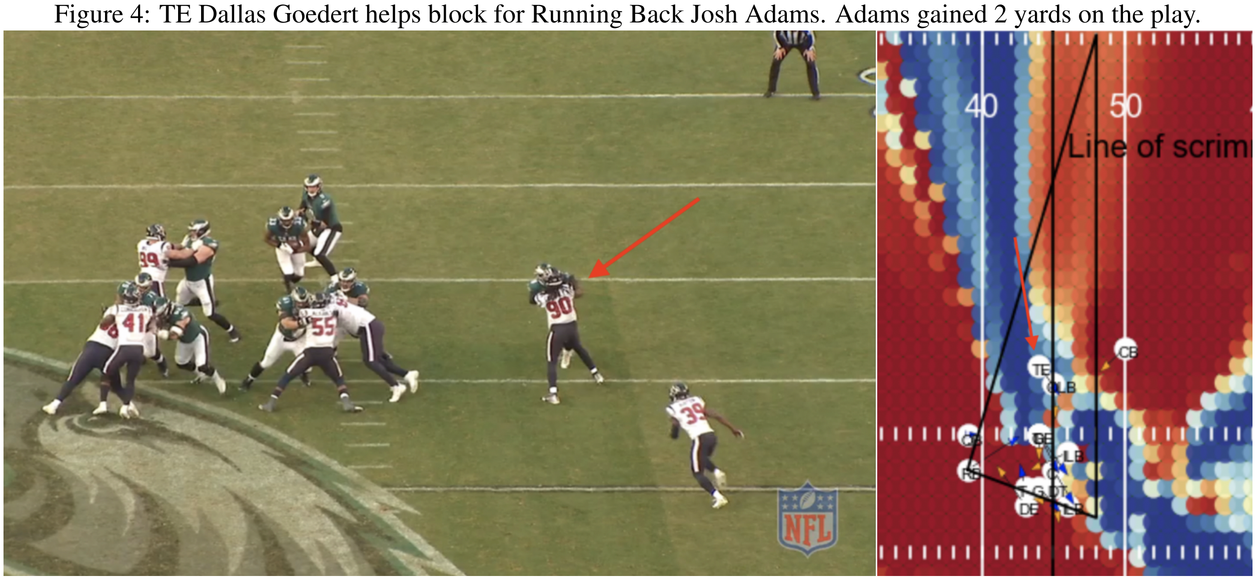

After watching the full play, Goedert blocks 2014 #1 overall pick Jadeveon Clowney very impressively throughout this play. Almost the entirety of the blue running space in the diagram is created by Goedert. Plays like this earned Goedert a top 5 *SpaceGrade* for the tight end position. Despite Adams only gaining 2 yards on the play by running up the middle into heavy traffic, Goedert still does his job very effectively. All blockers for the Eagles would be discounted for this poor play by many conventional grading measures that simply saw Adams get tackled early into the play for only 2 yards.

4.3 Defensive Linemen and Linebackers

Similar to each of the positions above, defensive linemen and linebackers can also be judged on space occupation. In their case, the goal shifts to taking space away from the offense, eliminating running lanes the runner can possibly take, and therefore minimizing the yards he can gain on the play. The correlations between ability to own space inside the rushing zone and Pro Football Focus run defense rank by position for nose tackles, defensive tackles, defensive ends, outside linebackers, inside linebackers, and middle linebackers were collectively higher than that of the offensive linemen. This is likely due to the fact that when a defender is rewarded for a tackle statistically, he has converged on the ball carrier and therefore limited the space created for him to run through. Current grading systems for defensive players on running plays are likely a better heuristic for the space creation algorithm than those for offensive linemen and tight ends. The implementation of the space creation algorithm also ensures that the results of double-teams are better accounted for on the individual level.

4.4 The Front Office

The greatest application of *SpaceGrade* analysis lies in the front office. Player salary data from 2018 was scraped from Over the Cap [8] and merged with the *SpaceGrade* ranking system. Players were then ranked based on their salary vs. their respective position. The correlations between a player's *SpaceGrade* and salary rank by position were 0.48, 0.18, and 0.24 for centers, guards, tackles, and tight ends respectively. The center position's higher correlation is reflected in Unger's 13^{th} ranking among salaries of centers. However, Smith ranks at 75^{th} among all tackles. The average absolute difference between player's *SpaceGrade* and salary rank by position is 31 for guards, 29 for tackles, and 15 for centers. Multiplying these by the median difference between each salary rank and the one directly below it translates to the average center, guard, and tackle being either over- or under-paid by $\$2,122,500$, $\$1,445,737$, and $\$1,984,000$ per year respectively based on their true space creation abilities relative to their peers. Based on these results, a league-best offensive line could be built for far less than most franchises currently pay their offensive lines. Space creation and ownership accounts for far less of what tight ends and defensive players are asked to do across all plays. Therefore, it is much harder to quantify pay discrepancies at those positions due to the larger variety of tasks they are asked to perform.

5 Discussion

Current methods for grading offensive linemen lack the ability to properly account for individual performances and how much better or worse those performances were compared to players at the same positions in similar play situations. Players graded highly in *SpaceGrade* can easily overshadow players on the same offensive line who are graded poorly in *SpaceGrade*. Those two player's offensive line as a whole could easily be deemed middle-of-the-line or even above average, even if only one of them actually did their job well. Space creation makes for a far more effective way of grading player performances on running plays, especially offensive linemen. On running plays, this is their primary job, and fewer outside factors need to be taken into account. For other positions, space creation should be taken into account proportionally based on how important creating space for others on offense or dominating player space as a defender is to the player's position and role in the given play. With full frame-by-frame player tracking data throughout the play, this would be very feasible to implement.

Further research into spatiotemporal modeling should combine this model with an expected points or win probability model. This is a better method for evaluating specific runners than just yards gained on a play. In addition, space created for and by receivers on passing plays could prove to be a useful metric. Given a certain amount of space or separation on a play, how much better or worse is receiver X's catch rate than receiver Y's given a similar play situation (number of receivers on the play, type of coverage faced, route run, skill level of opposing CB, etc.)? Does receiver X do a better job at creating space and separation for himself to catch the ball than receiver Y, given a similar play situation? Space creation is also a great way to identify how much space secondary running backs or fullbacks create while blocking for the primary rusher due to their high level of mobility in the backfield and ability to travel further down the field with the primary rusher.

These ideas are just the beginning for spatiotemporal modeling in the game of football. Player scouting and grading will certainly be revolutionized in the next few years with more accurate player tracking data, innovative methods for analyzing different aspects of player skill sets, and more advanced predictive machine learning techniques. Unlike the NBA or MLB, a hard annual salary cap makes predicting player abilities in various types of offenses absolutely crucial. Therefore, building an ideal roster under this cap at a time when an increasingly expensive market for skill players (QB, RB, WR) exists will require increasing levels of analytical optimization.

6 Acknowledgements

A huge thank you to UVA football's analyst Matt Edwards for providing me with the incredible opportunity to perform undergraduate statistical research with the UVA football team and for his immeasurable help, support, and professional insight into the sport throughout this project. In addition, a big thank you to Professor Jeffery Woo, Professor Abigail Flower, and Professor Jeffery Holt at UVA for providing me with invaluable guidance and acumen.

References

[1] Baldwin, Ben (@benbbaldwin). "Tweeting random stuff I need to link for a piece. Here's pass rates in neutral game script situations." Dec 16, 2019, 10:34 AM. https://twitter.com/benbbaldwin/status/1206613614586662912.

[2] Kyle Burris. A Trajectory Planning Algorithm for Quantifying Space Ownership in Professional Football. Big Data Bowl Competition, 2019.

[3] Luke Bornn and Javier Fernandez. Wide Open Spaces: A statistical technique for measuring space creation in professional soccer. MIT Sloan Sports Analytics Conference, 2018.

[4] "NFL Stats" TeamRankings.com: Sports Predictions, Rankings, and Stats. [Online]. Available: https://www.teamrankings.com/. [Accessed: 20-Dec-2019].

[5] "NFL Schedule - 2019" ESPN. [Online]. Available: https://www.espn.com/. [Accessed: 20-Dec-2019].

[6] Chuong Do. The Multivariate Gaussian Distribution. Stanford University, 2008.

[7] Joseph Sill. Improved NBA Adjusted +/- Using Regularization and Out-of-Sample Testing. MIT Sloan Sport Analytics Conference, 2010.

[8] Overthecap.com, '2018 Salaries', 2018. [Online]. Available: https://overthecap.com/position/. [Accessed: 11- 11-2019].

Python Notebooks can be found at `https://github.com/alexcstern`