# A Spatial Framework for Analyzing NFL Offensive Line Play

Paul Ibrahim University of Chicago pibrahim@uchicago.edu

# Abstract

In the NFL, the offensive line plays a crucial role in any offense. On passing plays inside the pocket (the scope of this paper), an offensive line's task is generally two-fold: (A) provide defender-free space for the quarterback and (B) sustain this operable space for as long as possible. However, current measures attempting to assess offensive line play fall short of providing a holistic reflection of a line's true effectiveness. Counting stats such as sacks and quarterback pressures have shown to be unreliable [1], and more advanced metrics such as Pass Block Win Rate (PBWR) [2] that are constructed around discrete "win-loss" outcomes do not provide a continuous measure that reflects the spatial fluidity inherent to offensive line play. In this paper, we seek to redefine the notion of pressure from a spatial perspective. Using publicly available tracking data from the 2020 NFL regular season [3], we design an analytical framework centered around spatial analysis within the pocket, using dynamic features of the quarterback's spatiotemporal setting to offer more meaningful information about the quality of an offensive line's play. The methodology outlined offers a more comprehensive understanding of what individual players on the offensive line, segments of the line, and the line as a whole contribute to the quarterback's protection<sup>1</sup>.

# **1 Past Approaches**

# 1.1 Sacks

In the past, there have been a variety of approaches to quantify offensive line performance. Principally, in an attempt to measure how well an offensive line protects a quarterback, the obvious counting measure to look towards is sacks allowed, how many times the defense was able to tackle the quarterback before a passing event. However, analyses have produced results discrediting the validity of sacks as a measure reflective of offensive line play, instead determining sacks to largely be a function of a quarterback dropping back to pass has been used to try and better capture offensive line failures outside of the quarterback's control, but similarly, pressure rate has been shown to be largely subject to a quarterback's decisions and play [4].

# 1.2 Pass Block Win Rate

<sup>&</sup>lt;sup>1</sup> The methods prescribed in this paper to analyze offensive line play are equally suited to analyze defensive line play

In an attempt to better assess offensive line play, in 2018 ESPN produced Pass Block Win Rate (PBWR), a metric generated from player tracking data [2]. PBWR is compiled on an individual player basis, assigning either a success (win) or failure (loss) outcome on each play; a win is defined as an offensive player "sustain[ing] their blocks for 2.5 seconds or longer", with a failure defined as a failure to do so. Thus, PBWR is the percentage of plays in which an offensive lineman records a win. This approach, while certainly providing a better assessment of offensive line performance, fails to represent the dynamism of the task of the offensive line. That is, the performance of an offensive lineman conceding a pressure 1 second after the snap is considered equal to that of a lineman who concedes pressure after 2.4 seconds, and likewise the metric offers no distinction for the degree of protection the quarterback enjoys. Of course, across a large sample of observations, such distinctions typically even out, but on the level of an individual game, play, or even frame, a more fluid approach proves useful.

# **1.3 Survival Analysis of Quarterback Pressures**

In 2019, seeking to analyze offensive line performance more comprehensively, Riske initiated the application of survival analysis, defining a failure event as the instance in which a pressure was recorded [5]. This is a good approach, and one we seek to build upon from a spatial perspective in section 5.

# 2 Data Processing

# 2.1 Normalizing the Data

In our analysis, we use a publicly available repository of tracking data from six games during the 2020 NFL regular season compiled by Yui [3]. The data contains x and y coordinates for all players on the field and the ball, with 10 observations recorded each second, and an x-coordinate range of (30,750) and a y-coordinate range of (30,350). For our purposes, we wish to normalize the data such that the initial position of the ball pre-snap is at x=0. The y-coordinate of the ball pre-snap can take on any of three different values, because depending on the ending location of the preceding play the ball can be snapped from the left hash, center, or right hash. We define the pocket as a 10x12 yd<sup>2</sup> rectangular space, being centered around the initial position of the ball. Additionally, we wish to normalize the coordinates to reflect yard-measurements. Thus, the line of scrimmage (initial position of the ball) is where x=0, and the furthermost edge of the pocket is x=-10, while the y-bounds of the pocket span 6 yards to either side of the ball but vary from play-to-play according to the hash mark at which the ball was snapped. Thus, the x-coordinate range becomes (-10,110) and the y-coordinate becomes (-26.67,26.67)<sup>2</sup>.

# 2.2 Sampling Data

Because our focus is on passing plays within the pocket, we sample only plays wherein the quarterback remains inside our determined pocket the entirety of the time between the snap of

<sup>&</sup>lt;sup>2</sup> In the raw data, (30,30) represents the bottom-left corner of the field, and (750,350) represents the top-right corner of the field. After our normalization, (-10,-26.67) represents the bottom-left corner of the field, and (110,26.67) represents the top-right corner of the field.

the ball and the time the ball is released<sup>3</sup>, and within these plays we sample the frames between the instance the snap of the ball is recorded and the instance a forward pass is recorded. Originally drawing from a dataset of six games, our specified sample consists of 310 plays and 7094 total frames. Additionally, we limit our analysis to the first 35 frames (3.5 seconds) of pocket play after the snap because the sample of plays where the guarterback takes more than 3.5 seconds after the snap to release the ball is scarce. Thus, compounding this relatively small sample of observations with the data's occasional noisiness<sup>4</sup>, we're unable to draw authoritative conclusions, but we can outline viable methods to analyze offensive line play using tracking data<sup>5</sup>.

### **3 Introductory Spatial Techniques**

The foundation of our spatial framework is constructing a Voronoi tessellation from players within the pocket. A Voronoi tessellation is a simple algorithm which computes, for any set of points in two-dimensional Euclidean space, all the space closer to a given point than any other point, denoting this space as that point's Voronoi cell.

Of course, spatial domination is a more complex idea than just what space is closest to who, a problem Javier Fernandez and Luke Bornn tackle [7], but for the purposes of constructing a straightforward framework, we use the simple definition of spatial occupation afforded by the Voronoi tessellation.

Thus, implementing the Voronoi tessellation in a pocket-passing setting, we produce the following diagram using data from the Ravens vs Chiefs game during week 3 of the 2020 NFL season. It is of note that on this play the Chiefs are the offensive team and a forward pass is recorded at frame 30, the final frame depicted.



#### Figure 1: Voronoi Diagram

<sup>3</sup> Unfortunately, this precludes many passing plays in which the quarterback scrambles outside of the pocket before making a throw. This is because we deemed it too difficult to control for the fact that different guarterbacks have different proclivities to exit the <sup>4</sup> For instance, ball location data only updates one in every ten frames unlike the player location data which updates every frame

<sup>5</sup> It is of note that because all quarterbacks in our sample were right-handed, the analysis exhibited is under the presumption of right-handedness. The same methods can of course be applied to analyze left-handed quarterbacks, but the data to achieve that in this paper is not available to us.

The Voronoi diagram offers an intuitive sense of the spatial dynamics at play. While at the instance of the snap the offense obviously dominates the vast majority of space within the pocket, we see in succession the defense increasing their occupation of the space, especially on the left side of the line wherein Matthew Judon's Voronoi cell functionally flanks Mitchell Schwartz, a threat sensed by quarterback Patrick Mahomes who in turn moves to the right. Likewise, we can observe pressure up the middle by way of Derek Wolfe, who wedges himself behind center Austin Reiter and pushes guard Mike Remmers backwards, compressing both Remmers' Voronoi space and that of quarterback Mahomes.

Observing the change in area in Patrick Mahomes' Voronoi space over the course of this play and that of the offense as a whole, shown in the table below, we see it expectedly decrease as the play progresses, with the biggest change in both offensive and quarterback Voronoi area occurring between frames 13 and 19, reflecting Judon's flank on the edge and Derek Wolfe's ingress into the center of the pocket.

	frame1	frame7	frame13	frame19	frame25	frame30
Offensive Voronoi Area	114.00634 / 112.3581	108.02242 / 105.7387	95.63331 / 92.34656	73.99175 / 77.94803	75.92791 / 71.11568	63.26031 / 70.35549
QB Voronoi Area	60.27334 / 49.32199	54.42512 / 45.79621	45.77989 / 38.6502	29.45247 / 28.93712	19.22202 / 20.4432	11.02796 / 18.01495

Format: Observed Value in Play / Sample Mean

#### Table 1: Voronoi Areas Throughout Play

Likewise, from this table we can see that by the measure of both Voronoi areas, the offensive line was providing more space than expected, until near the end of the play, when the values of both the offensive Voronoi area and the quarterback Voronoi area fall below expectation.

The Voronoi tessellation is a broad, but useful approach. It allows us to capture a general picture of the occurrences of the play. As exhibited, we can observe the offensive line as a unit, seeing the rapidity of the decay in total area, and comparing it to other plays and other offensive lines, we begin to develop a mode of understanding an offensive line's performance.

However, the broad Voronoi approach lacks the scalability we desire in offensive line analysis. That is, while the Voronoi space of the entire line gives meaningful information, the Voronoi space of an individual linemen is susceptible to variations, either positive or negative, not indicative of their actual performance. To illustrate this point, between frames 25 and 30, center Austin Reiter's Voronoi space grows from 12.631 yd<sup>2</sup> to 18.063 yd<sup>2</sup>. However, as any basic understanding of football makes apparent, his contribution to the quarterback's protection between these frames hasn't improved whatsoever. Likewise, Eric Fisher's assignment on this play evidently seems to be blocking Pernell McPhee. From frames 7 to 13, Eric Fisher's Voronoi space decreases from 12.152 yd<sup>2</sup> to 11.563 yd<sup>2</sup>, not so much because of the increased intrusion of Pernell McPhee, but rather because Brandon Williams got past Kelechi Osemele who was neighboring Eric Fisher on the line, and thus Eric Fisher's Voronoi space decreased through no fault of his own.

# 4 Bin Method

To rectify the lack of scalability and differentiation in spatial value in the summarized Voronoi approach, we introduce the bin method. Partitioning the Voronoi tessellation of the  $12x10 \text{ yd}^2$  pocket into 30 2x2 yd<sup>2</sup> bins, yielding the arrangement visualized below in *Figure 2*, with the number inside each bin denoting that bin's assigned id.

0.0 -						
	1	6	11	16	21	26
-2.5 -	2	7	12	17	22	27
× -5.0-	3	8	13	18	23	28
-7.5 -	4	9	14	19	24	29
10.0-	5	10	15	20	25	30
-10.0 -	6	-3	Ċ		3	6

Figure 2: Bin Arrangement Guide

We proceed to chart the mean Voronoi occupation in each bin within our sample, the results of which are displayed in *Figure 3* below:



Figure 3: Mean Bin Voronoi Occupation

These results match the general perception of offensive line play. As expected, the offense usually occupies the vast majority of space in the pocket in the beginning of the play, and their space tends to gradually lessen, becoming increasingly concentrated around bins 14 and 19 near the center-bottom of the pocket where the quarterback usually is. Within 22 frames (2.2 seconds) of the snap, the periphery of the pocket is largely dominated by the defense, reflective of the compression of the offensive line, causing the quarterback's blindside [7] to become contested space.

Applying the bin framework, we can deduce areas where an offensive line, on a particular play or as a whole, perform above or below average. Recalling the play profiled in *Figure 1*, we can compare the offensive line's effectiveness in various bins against that of the league average, as exhibited in *Figure 4* below which computes a simple bin difference for frame 30 (instance a forward pass was recorded in the play).



Figure 4: Bin Profile Comparison at Time of Throw

From this figure, we can see very clearly the areas of the pocket wherein the offensive line overperformed, and likewise the areas they underperformed. Affirming the observations made in section 3, at the time of throw the offensive line is overperforming at the top-center of the pocket, the region occupied by Austin Reiter and Kelechi Osmele, and likewise the line is overperforming at the bottom of the pocket. However, most critically, we see that in a diagonal region extending from the top-left corner of the pocket to the center, the offensive line is notably underperforming, especially in bins 13 and 18, indicating pressure to the quarterback's center. Similarly, the offensive line is shown to be underperforming on the right side of the pocket, indicating edge pressure being induced on the guarterback. Thus, from this plot, we can glean an understanding of the different values carried by various space in the pocket across time. That is, the offensive line may appear to be performing better than expected at the top of the pocket, but we see this space doesn't hold much value when the defense has infiltrated the space underneath them, closer to where the guarterback is. As a proxy for protection around the guarterback, we know that Mahomes is in bin 24, and thus we can sum the area of all neighboring bins to gauge the degree of protection he enjoys, finding this figure to be 19.096 vd<sup>2</sup>, 3.866 vd<sup>2</sup> below expectation.

Though, we again refrain from using an individual player's contribution to bin space, for the rationale outlined in section 3, we now have a framework wherein with a general understanding of a given player's responsibility, we can develop metrics to assess their performance. For instance, it's evident that a left tackle's responsibility, in the vast majority of circumstances, is to protect the quarterback from pressure originating on the lefthand side of the pocket. Thus, analyzing the relevant bins, we have a strong proxy for the left tackle's performance on any given play. Compiling these metrics with others, we produce the following table, allowing us to compare performance in certain groupings of bins relative to expectation:

	total_area	dp_neignboring_bin	len_tackie_bin	len_guard_bin	center_bin	right_guard_bin	right_tackie_bin
Actual	63.26031	19.09637	16.78023	21.11644	26.50767	29.689420	19.97241
Expected	70.35549	22.96242	21.82427	25.62589	27.52479	25.789150	21.00643
Difference	-7.09518	-3.86605	-5.04404	-4.50945	-4.50945	3.900272	-1.03402

### Table 2: Grouped Bin Profiles

Of course, the comparative approach of bin profiles isn't only limited to an individual frame against the league average. You can compare against the average bin profile of that offensive line throughout the season, against the defensive line's average, against past seasons, et cetera. However, the size of our sample precludes this, and thus such comparisons are outside the scope of this paper.

In summary, the bin method allows us both to observe the performance of the offensive line as a whole throughout the entire pocket, but it also allows us to employ a more microcosmic approach, capturing nuances of offensive line play and shifts in spatial dynamics that are overlooked by the broader Voronoi approach.

# **5 Survival Analysis**

As referenced in section 1.3, Building upon Riske's work [5], we seek to provide a framework for analyzing the spatial dynamics of offensive line play using survival analysis [8].

### **5.1 Survival Framework**

Similar to Riske's work, we define the failure event as an offensive line conceding a pressure, and our analysis concerns the relationship between the time to this event and the likelihood of survival. Critically, our approach differs from that of Riske by our method of defining pressure. While his analysis uses data produced by observers of the game charting when they perceive a pressure occurs, our approach defines pressure based upon relevant spatial features within the pocket. Furthermore, we analyze distinct types of pressures faced, as pressures from different directions and pressure at various locations vary significantly in their survival curves.

We define a pressure at the instance in which the area of the space in question falls below a certain threshold. This threshold is relatively arbitrary and is largely subject to the aim of the motivating analysis. For our purposes, we observe the pocket at all instances of a forward pass within our sample, and using the distribution of the area of the space in question, we select the value at the 30<sup>th</sup> percentile of this distribution as our threshold.

Because much of the data we have is right-censored (i.e. often a quarterback throws the ball before a pressure is observed), and because the time to event/censor data is right-skewed, we use the nonparametric Kaplan-Meier estimator. It is important to note that because our censor time was set at 35 frames, in most cases our survival probability doesn't descend all the way to 0.

### 5.2 Survival Analysis of General Voronoi Space

Applying the Kaplan-Meier estimator to both the offense's Voronoi space as a whole and just the quarterback's Voronoi space, we retrieve the survival curves and hazard functions displayed below in *Figure 5*<sup>6</sup>.



\*Note difference in scales of y-axis Figure 5: Offense and Quarterback Voronoi Space Survival and Hazard Curves

Evidently, both the survival curves and hazard functions approximately emulate each other in shape, and likewise these curves reflect what we would expect. That is, throughout the first twelve frames, the likelihood of survival decreases quite slowly, indicating that a quarterback can reliably expect to have at least one second in the pocket without facing pressure. However, by frame 15, both curves begin to fall precipitously. Of note, as reflected by

<sup>&</sup>lt;sup>6</sup> Because of the overlap between the survival and hazard curves, we decided to forego the traditional graphical representation of the Kaplan-Meier estimator using an indexed step function and rather have elected to use a smoothed graph for the sake of clarity and visual readability

the hazard functions, after 30 frames, the probability of a pressure in the subsequent frame becomes exceedingly likely.

### 5.3 Survival Analysis of Pocket Features

To obtain a more precise understanding of spatial pressure, we seek to perform a similar analysis on two relevant features of pressure within the pocket; that is, pressure on the quarterback's blindside and pressure up the middle. Rather than denote the blindside and middle using the bins defined in section 4, we chose to use an angular approach to allow us to consider both offensive and quarterback directional Voronoi space with greater precision (bin survival analysis in section 5.4). Thus, taking an angular profile of the quarterback at every frame, we define Voronoi space between 60 and 120 degrees of the quarterback's position to be the space in the center, and Voronoi space between 150 and 300 degrees to be their blindside<sup>7,8</sup>. The resulting survival and hazard curves are illustrated in *Figure 6* below:



Figure 6: Pocket Features Survival and Hazard Curves

These curves have a number of distinct features worth noting. First comparing the groups (blindside vs center) qualitatively, immediately evident is how low the initial hazard is for pressure up the middle, especially when using the offensive Voronoi space threshold. Presumably, this is due in large part to the fact that quarterbacks often line up directly behind the center to receive a snap before dropping back deeper into the pocket to throw the ball, and thus even slight initial pushes from the defensive line up the middle in the beginning of the play register by this definition as a pressure. However, even omitting the first 7 frames after the snap from the survival analysis, more than enough time for the quarterback to drop back into the pocket, and re-calculating time to failure, we still observe only a 68.734% survival rate of the offensive Voronoi space in the center at frame 8, indicating that this space is substantially more vulnerable to pressure soon after the snap than the other spaces observed. Throughout the majority of the play, we observe that the hazard functions are all analogous, only observing evident divergence around frame 24.

<sup>&</sup>lt;sup>7</sup> A measure of 0 degrees denotes space exactly to the right of the quarterback, 90 degrees denotes space directly in front of the quarterback, et cetera. These angular definitions of center and blindside space are based upon a standard human field of view, with 60 degrees being the range of near-peripheral view, and 210 degrees being the range of far-peripheral view, hence the blindside ranges 150 degrees.

<sup>&</sup>lt;sup>8</sup> Information about player direction was not provided in the data source. However, with this information, one could define pressure specific to the quarterback's pose, producing many interesting variations of a spatially-dynamic survival analysis.

To analyze these differences statistically and assert a difference between the survival curves for blindside and central space, we perform a log-rank test. The log-rank test is a hypothesis test with the null hypothesis that the hazard functions are equivalent, and an alternative hypothesis that the hazard functions differ. Using a log-rank test to assess the difference between the hazard functions of the offensive blindside Voronoi space and the offensive center Voronoi space, we compute a chi-squared statistic of 99.2, and accordingly a p-value less than 2e-16, allowing us to confidently affirm a difference in the two hazard functions. Likewise, performing a log-rank test on the corresponding quarterback Voronoi spaces, we calculate a chi-squared statistic of 20.4 with a p-value of 6e-6.

### 5.4 Survival Analysis of Pocket Bins

Conducting the same analysis by bin, rather than determine a threshold for pressure unique to each bin, we elected to define a universal threshold for all bins, recording a pressure in that bin at the instance that the bin's area fell below  $\sim$ 1.51 yd<sup>2</sup>. For the purposes of our survival analysis, we're not seeking to judge each bin relative to the expected performance in that bin, but rather we seek to identify the moment that bin's area became contested to a point of inducing pressure in that location, and thus a universal threshold is appropriate. Using a Kaplan-Meier estimate, we produce the survival curves displayed below in *Figure 7*.



Figure 7: Bin Survival Curves

As expected, the deeper in the pocket a bin lies, the longer the expected survival time. However, notable from this graph is the consistent order of survival curves observed. That is, where the central bins of the pocket consistently have the highest survival probability across time, the survival curves of the edges of the pocket (zones typically occupied by the offensive tackles) have the lowest survival probability across time, with the bins typically occupied by the guards situated in the middle of these curves. These observations are affirmed by performing a log-rank test, a summary of which is presented below in *Table 3*.

Bin.Depth	Chi.Sq.Statistic	p.value
0-2 Yards	425	<2e-16
2-4 Yards	424	<2e-16
4-6 Yards	413	<2e-16
6-8 Yards	341	<2e-16
8-10 Yards	234	<2e-16

Table 3: Grouped Bin Log-Rank Results

# **6** Conclusion

In this paper, we have outlined a methodological framework through which one can dynamically evaluate offensive line play. Within our sample, we have used the applied methods to observe spatial tendencies in offensive line play within the pocket, and have offered examples as to how an offensive line's performance on a given play can be assessed. These methods can be used to evaluate the performance of the offensive line in its totality, a specific region within the pocket, or an individual player when provided with schematic understanding. Each of these entities can be assessed on varying time-scales, ranging from an individual frame to an entire season, affording a nuanced and comprehensive spatial understanding of offensive line performance.

### References

- [1] Lisk, Jason. "Sacks Are a Quarterback Stat." *The Big Lead*, The Big Lead, 19 Oct. 2018, www.thebiglead.com/posts/sacks-are-a-quarterback-stat-01dxqapkgvw9.
- [2] Burke, Brian. "We Created Better Pass-Rusher and Pass-Blocker Stats: How They Work." *ESPN*, ESPN Internet Ventures, 5 Oct. 2018, <u>www.espn.com/nfl/story/ /id/24892208/creating-better-nfl-pass-blocking-pass-rushing-stats-analytics-explainer-faq-how-work</u>.
- [3] Yui, Lao Sze. "NFL-Live-Tracking-Data/output\_file at Main · 903124/NFL-Live-Tracking-Data." *GitHub*, github.com/903124/NFL-live-tracking-data/tree/main/output\_file.
- [4] Eager, Eric. "Quarterbacks in Control: A Pff Data Study of Who Controls Pressure Rates: NFL News, Rankings and Statistics." *PFF*, PFF, 14 Aug. 2019, <u>www.pff.com/news/pro-z-quarterbacks-in-control-a-pff-data-study-ofwho-controls-pressure-rates</u>.
- [5] Riske, Timo. "PFF Quantitative Insights: Investigating Pass Protection with Survival Analysis: NFL News, Rankings and Statistics." *PFF*, PFF, 17 Dec. 2019, <u>www.pff.com/news/nfl-pff-quantitative-insights-investigating-pass-protection-with-survival-analysis</u>.
- [6] Fernández, Javier & Bornn, Luke. (2018). Wide Open Spaces: A statistical technique for measuring space creation in professional soccer.
- [7] "Blind Side Definition Sporting Charts." *SportingCharts.com*, <u>www.sportingcharts.com/dictionary/nfl/blind-side.aspx</u>.
- [8] Moore, Dirk F. Applied Survival Analysis Using R. Springer Science+Business Media, 2016.