

Doinkalytics: A New Paradigm in National Football League (NFL) Placekicking Evaluation

Football Track
20251410

1. Introduction

Kickers account for over 30% of scoring in the NFL. All 20 of the NFL's all-time leading scorers have been kickers. Made and missed field goals commonly determine the outcomes of drives, games, and even entire seasons. Unfortunately, analyses of these vital performances remain inadequate, and conventional place-kicking metrics fail to accurately characterize some of the most noteworthy performances in professional football. [1]

This paper introduces two new methods designed to enhance the evaluation of kickers in the NFL.

First, we introduce a context-aware model that applies an expected-points-added framework to kicking attempts. After studying the successes and failures of over 7,000 NFL kicking attempts, we developed a model that estimates the probabilities and expected points totals for every kick attempt in the NFL by blending key inputs including yardage, precipitation, playing surface, elevation, wind, and temperature. Not all 50-yard attempts are equally difficult and any framework that fails to account for contextual effects in turn fails to accurately quantify the accomplishments of the top scorers in professional football.

Second, we introduce Command+, a new metric that quantifies kick accuracy by accounting for the precise location of the football as it crosses the back plane of the endzone. Some made kicks are more impressive than others, and by factoring in the spatial locations of attempts we assert we can further refine our understanding of kicking successes and failures.

Command+

2022 NFL Field Goal Attempts | Total Feet From Centerline




Figure 1: Command+: The "where" of field goal makes and misses

Finally, we conduct a case study that applies these methods to evaluate kicking performances in the NFL over the 2022 and 2023 seasons. Our results demonstrate that the new approaches can enrich our understanding of the best and worst kickers in the NFL.

2. Previous Work

While some previous research has examined how placekickers in pro football are affected by conditions like temperature and wind, these existing publications do not include context-aware modeling approaches that estimate the expected points of individual place-kicking attempts. [e.g., 2,3] Further, the lack of publicly available ball-tracking data for NFL games presents a significant obstacle for any research aiming to explore any matters of ball flight in the NFL. As a result, little previous football research has leveraged key flight variables like launch angles, launch velocities or cartesian ball locations to investigate passing, kicking, or punting performances.

One NFL staffer has revealed important proprietary data exist behind closed doors, “Using @NextGenStats ball tracking data, Justin Tucker's game winning (sic) field goal crossed the upright with a y-coordinate (width of field) of 26.52. The exact middle of the field is $y = 26.67$ That is, if the uprights were half a yard wide, the kick would have still been good”. [4]

We believe that integrating kicking context models and flight metrics will allow for a greater appreciation of field goal kicking in the NFL, in the same ways that concepts like shot quality and in-rim ball location data have advanced the appreciation of shooting performances in the NBA. [5]

Recent in-rim shot-tracking technology in basketball has allowed analysts to understand that shooting results are about a lot more than simply makes and misses. [6] By moving past a binary framework of performance, we may better appreciate those who are truly the greatest kickers in the NFL - beyond the limits of conventional stats.

3. Methods

3.1. Context Model

Field goal difficulty is about a lot more than just yardage. Adam Vinatieri's legendary 45-yard game-tying field goal in the 2001 NFL Playoffs against the Oakland Raiders was one of the most impressive kicks in league history. The kick wasn't unusually difficult because it was a 45-yarder in a high-leverage postseason moment - it was particularly challenging because it occurred in blizzard conditions amidst swirling winter winds. We argue that by accounting for variables including wind, elevation, playing surface, and other weather factors, we can better assess the difficulty of individual kicking attempts in the NFL, and then better assess the performances of individual kickers.

Conventional kicking statistics fail to account for key effects. As a result, conventional characterizations of vital kicking performances in the NFL are deficient. Here, we introduce Split-The-Uprights-Difficulty (STUD), a new metric derived from parametrized logistic regression model that estimates the make-probabilities and expected points of individual field goal attempts in the NFL.

We compiled field goal data from the nflFastR play-by-play data API back through the 2014 season. [7] Then, after integrating these events with precise geolocations of NFL stadiums, we used the open-meteo.com historical weather API to collect detailed weather conditions for each game including wind speed, temperature, and precipitation. [8] Each indoor game was set to a standard 0 mph wind speed, 72-degree temperature, and 0 inches of precipitation. The final dataset did not include overseas games due to limitations in the weather data API, however, it included 10,918 field goals from the 2014 to the 2024 season.

The modeling data was limited to observations from outdoor games. The final modeling data included 7,635 outdoor field goals. The data was split into a training set containing 75% of the data and a test set containing 25% of the data. The training data was used with cross validation for each model to prevent overfitting. Model hyperparameters were set using grid-search. Once hyperparameters were analyzed, the top performing version of each model was retrained with the entire training set. All modeling and data splitting was done using the caret library in R. Data was centered and scaled using caret's preprocessing functions. [9]

The kicking context variables used for prediction were field goal distance, wind speed in kilometers per hour, a binary marker for the presence of precipitation, temperature in Fahrenheit, field type (ex. turf), elevation in meters above sea level, and a binary marker for whether the temperature is freezing or not.

One concern during the modeling process was the quality of kickers inflating the probability of making difficult kicks. If only good kickers are attempting the hardest of kicks, this would inflate the true expected probability of a kick due to survivorship bias. To control for this, for each kick, the longest made field goal in the dataset from that kicker was used as a predictor variable. This was chosen over controlling for the kickers themselves because it allows for setting this variable to the median to assess the predicted probability of making a kick for a 50th percentile kicker. Summary statistics for each of these variables can be found in Tables 1 and 2 below.

TABLE 1: Summary Stats of Numeric Modeling Variables

Variable	Minimum	25%	Median	75%	Maximum	Mean
Kick Distance (yards)	18.0	30.0	38.0	47.0	68.0	38.1
Elevation (m)	-6	5	68	190	2200	169
Wind Speed (khp)	0.4	7.2	10.5	15.0	43.7	11.7
Temperature (F)	-7.9	42.1	55.2	68.2	94.0	54.4
Longest Made Kick by Player (yards)	30.0	56.0	58.0	61.0	66.0	58.3

TABLE 2: Summary Stats of Categorical Modeling Variables

Variable	Category 1 (count)	Category 2 (count)	Category 3 (count)
Field Surface	Grass (4231)	Hybrid (610)	Turf (2794)
Precipitation	TRUE (1368)	FALSE (6267)	N/A
Freezing	TRUE (703)	FALSE (6932)	N/A

The goal of STUD is to accurately evaluate kick difficulty for further application of kicker impact. For this to work successfully, it is important that when a kick is assessed as a 50% make probability, the kick is made 50% of the time. To assess this, models were evaluated on R^2 of a calibration curve where each cluster of test data had an equal number of samples. The test data (n=1908) was split into 100 clusters for this. The equal number of samples is particularly important due to skewness of the data - there are a lot more kicks taken that have a high make rate than a low make rate. As a result, bracketing by probability ranges strictly creates some small sample size issues with lower probability kicks. In evaluating kicks this way, we can feel confident that the model is correctly assessing the probability of a kick being made. The R^2 of the model calibrations can be seen in Table 3 below.

Table 3: Model Calibration Scores

Model	Test Calibration R^2
Logistic Regression	0.720
Gradient Boosted Tree	0.673
Neural Network	0.669
Regularized Logistic Regression	0.573
Random Forest	0.460

It may be surprising to find that logistic regression was the best performing model, however, a linear model allowed us to specify relationships between variables based on what we know about kicking directly rather than relying on a more complex model to infer these relationships. The general form of logistic regression is provided below in Equation 1, where X is the vector of predictor variables and β is the set of coefficients that define the relationship between the predictors and outcome. We tried a variety of variable and interaction combinations that resulted in equation 2 below as the top performing combination among available variables.

$$\Pr(\text{Made FG}) \sim \frac{1}{1 + e^{-\beta X}} \quad (1)$$

$$\Pr(\text{Made FG}) \sim \text{longestKickPlayer} + \text{kickDistance} * \text{temperature} * \text{windspeed} + \text{freezing} * \text{fieldSurface} * \text{precipitation} \quad (2)$$

Intuitively, this seems to make sense. In the interaction between distance and wind and temperature all variables are related to ball flight. Cold temperature can make a ball more difficult to kick as it becomes flatter due to the ideal gas law - limiting distance because the ball does not come off the leg as quickly. [10] Additionally, wind speed may also influence kicking that depends on distance - longer kicks have more flight time to be affected by the wind regardless of the direction. [11] The other set of interactions is between field surface, precipitation, and whether the temperature is freezing. These variables would all seem to affect the kicker's platform and footing - precipitation could make the surface slippery and freezing temperatures could change the hardness of the ground.

Another consideration in model selection, was whether the models tracked with some simple kicking heuristics: first, that longer kicks are more difficult than shorter kicks and second, that given clearly difficult kicking conditions, a kick of equal distance indoors would be evaluated as no easier than in difficult conditions outdoors. For our purposes, we assessed "difficult" conditions as the 95th percentile of wind, the 5th percentile of temperature, precipitation, and the 5th percentile of elevation observed among outdoor kicks in the dataset.

All predictions set the kicker's longest kick made to the median value of 58 yards, which is also very close to the mean value of 58.3. As indicated by the plot, the probability of making a kick under these separate conditions is very similar for the shortest kicks possible, but quickly diverge and increase in separation as kicks get longer. The probabilities again converge for the longest of field goals where we may be seeing some survivorship bias from the best kickers still, but also reach a point where the distance becomes the primary driver of difficulty - regardless of conditions.




Figure 2: The effect of distance on field goal make probability: indoor vs. worst case outdoor

Tree based models like random forests or gradient boosted tree models provided issues that during iterations of the modeling process presented predictions where the probability of a kick sometimes increased with a longer kick, holding all other predictors the same. While having some promising calibration scores, this ultimately became disqualifying when considering final model selection.

From this model we have created STUD (Equation 3) which allows us to assess kicker quality when controlling for conditions. A kicker's expected points added (EPA) over their set of kicks "P" are calculated by the difference between the outcome and probability multiplied by points available for a make, then summed (Equation 4). We also grade kickers on total performance over expected, adjusted for attempts and difficulty, with the result given as a Z-scored percentile of the raw score given below for a 0-100 scale. (Equation 5)

$$STUD = \text{Pr (Made Field Goal)} \quad (3)$$

$$\text{Expected Points Added} = 3 * \sum_{p=1}^N (\text{pointsScored}_p - STUD_p) \forall p \in P \quad (4)$$

$$STUDPlus (\text{raw}) = \frac{\sum_{p=1}^N (\text{kick_result} - STUD_p)}{\sum_{p=1}^N (STUD_p)} \forall p \in P \quad (5)$$

3.1.1. Limitations

There are two primary biases we expected to influence the model that could skew results. The first bias comes from the coaching side of the decision to kick a field goal and the second comes from the quality of the kicker themselves.

From a coach's perspective, attempting a field goal carries significant risks if they are not confident in making the kick. Beyond the absence of points, a missed field goal results in favorable field position for the opponent, as they take possession at the spot of the kick, typically 7 yards behind the line of scrimmage. Compared to punting or going for it, this often provides the opponent with the best starting field position among all possible outcomes, making the decision particularly consequential for long field goals. In general, this leads to a lack of in-game attempts of challenging kicks, particularly long ones.

As previously discussed, kickers vary in skill, and the top kickers in the league are often better at making difficult kicks. Consequently, the most challenging kicks are typically attempted only by the best kickers, which can artificially inflate the make probability for these difficult attempts.

Blocked kicks were not included in this analysis as blocks are infrequent and rarely factor into field goal kick difficulty for the kicker themselves.

Another limitation involves the weather data. Weather conditions for each game are based on a single data point, which may not accurately reflect the exact conditions at the time of each kick. For

instance, wind gusts or lulls occurring during the kick are not captured. Incorporating more granular, time-specific weather data could better reveal the relationship between kick context and field goal difficulty. This artificially decreases potential amount of variance in the data available for modeling, leading to decisions like making precipitation a binary variable when a measurement of amount rain was available. Without more granular weather data, additional data in general could help to solve some of these issues.

Lastly, while windspeed is clearly important for kicking, the predictive power of this variable is severely limited without the interaction of the direction of the wind and how that affects a specific stadium – which influences the kicker.[11] Future analysis that not only controls for the stadium, wind direction, but also which endzone the kicking team is kicking toward would certainly help to better understand these effects.

3.2. Command Score

Using the same play-by-play data as for the contextual analysis, we manually charted the location at which each kick passed through the uprights, outside of the uprights, or “doinked” off the uprights. This process involved recording the hashmark from which the kick was attempted and the cartesian coordinates where the ball crossed the uprights. To estimate these locations, we used a one-to-one representation of the uprights in DESMOS. [12]. Charting was done for all field goals in the 2022 and 2023 NFL seasons in their entirety.

Each observed field goal attempt used the DESMOS graphical representation to estimate the point where the ball crossed the plane, and then recorded the coordinates alongside the corresponding kick in a spreadsheet. The point where the crossbar meets its base serves as the origin, meaning any y-coordinate above zero indicates the ball was above the crossbar, and any x-coordinate between -9.25 and 9.25 falls between the uprights.

A normal distribution was created from the data to understand the dispersion. The true mean of the data was 0.025 feet to the left of the exact center of the goalposts. For the individual kick Location score, the mean was manually shifted to the exact center of the goalposts. Using the standard deviation of the fit distribution and the centered mean as the distribution for evaluation, the Location Score for each kick was determined by the formula given in Equation 6 below. Dead center was set to be a score of 100 and in each direction from the center the score decreased corresponding to the probability distribution function described such that the scores were symmetrical on either side of the center as shown in Figure 3.

$$Location\ Score = \begin{cases} 2 * normalizedLocationPercentile, & \text{if } percentile \leq 50 \\ 200 - 2 * normalizedLocationPercentile, & \text{if } percentile > 50 \end{cases} \quad (6)$$




Figure 3: Location Score vs. Location for all Field Goals from the 2022-2023 NFL Seasons

Two scores were created for each kicker from this assessment. First, a location-based score was created according to the mean kick location for each player across the sample (Equation 7). Second, a consistency score was created based on the standard deviation of kick locations across the sample for each player (Equation 8). These two scores were placed on a 0-100 scale based on a normal distribution percentile, and then averaged to create an overall Command+ score encompassing both the accuracy of and consistency of each kicker (Equation 9).

$$Location\ Score\ (raw) = \frac{1}{N} \sum_{i=1}^N kickLocation \quad (7)$$

$$Consistency\ Score(raw) = \sqrt{\frac{1}{N} \sum_{i=1}^N (kickLocation_i - Player\ Location\ Score)^2} \quad (8)$$

$$CommandPlus = \frac{Location\ Core + Consistency\ Score}{2} \quad (9)$$

3.3. Kicker Grading

As the final step to our analysis, we sought to bring together the context adjusted difficulty of a kick with where it goes through, around, or off the uprights. To do this, we first compared across all kicks the mean Location Score and Consistency Score across the kick difficulties. We confirmed that as kicks become more difficult, the average left/right location from which it crosses the uprights moves away from dead center. We also found that as kicks become more difficult, the standard deviation of the individual kick Command score increased. An indicator of the validity of our rating systems. In essence, it's harder to put a tough kick down the center than an easier kick.

Next, we combined the Command+ and STUD+ scores to rate all kickers. These scores are correlated at 0.689 meaning they share less than 50% of the same information. ($R^2 = 0.475$). By understanding them as separate constructs, we used each of these scores on their own axis and found the magnitude of the vector created by the Command+ and STUD+ scores as an overall kicker rating (Equation 10). The plots of each player's scores are in Figure 4 below. The final version of the kicker rating is a percentile 0-100 of a normal distribution amongst the raw Kicker Rating scores.

$$Kicker\ Score = \sqrt{CommandPlus^2 + STUDPlus^2} \quad (10)$$




Figure 4: STUD+ Score vs. Command+ Score for all Kickers across 2022 and 2023 seasons combined

4. Kick and Player Analysis

4.1. Dome Yards

The diversity of kicking conditions and related factors means that some NFL kickers should have lower expectations than others, and any framework that fails to account for the relative difficulty of attempts, will also fail to accurately characterize the relative success rates of kickers in pro football. To create a frame of reference for those analyzing or watching football, we present an additional new metric called *Dome Yards*. Dome Yards is the length of kick that is of an equivalent STUD to a given outside kick – a metric that normalizes kick difficulty on a familiar frame of reference for players, fans, and front office members trying to understand the quality of a kick.

For instance, in the final week of the 2023 NFL regular season, New York Jet's kicker Greg Zuerlein missed a 49-yard field goal about 3 feet left of the goalpost against the New England Patriots. At face value, a missed 49-yarder seems bad by today's NFL standards, however, the conditions say otherwise. During this game it was 24° F, with some precipitation, and 12+ MPH sustained winds. An equivalently difficult kick in a dome would have been 58 yards – a 9-yard difference! and the median longest kick made for any individual kickers going back to 2014.

4.2. Expected Points Added

Over the course of time, some NFL kickers attempt harder kicks than others. To demonstrate this vital effect, we first applied the STUD model to estimate the expected points and make probabilities of 2,150 field goal attempts in our two-season study period. According to our analyses, the average attempt in the set had an expected conversion rate of 84.1% and was valued at 2.54 expected points, but these values vary in meaningful ways for individual kickers around the league.

Among the 38 kickers that attempted at least 25 kicks in our study period, our results indicate that Brandon McManus who kicked for Denver and Jacksonville had the “hardest” kicking profile, while Robbie Gould, who kicked for the San Francisco 49ers in 2022 had the “easiest” set of attempts. We estimate that an average McManus attempt would result in a made field goal 79.4% of the time and yield just under 2.40 points. For Carlson these expectations are 89.0% 2.67 points respectively. Tables 4 and 5 include the results for most difficult and least difficult kicking profiles in our study period.

Table 4: Most difficult kicking profiles, 2022 and 2023 seasons (Min. 25 FGA, 38 qualifiers)

Kicker	Total FGA	Estimated Make%	Expected Points
Brandon McManus	70	79.35	2.38
Matt Prater	58	79.42	2.38
Graham Gano	59	79.60	2.39
Chris Boswell	59	80.47	2.41
Greg Zuerlein	74	80.84	2.43

Table 5: Least difficult kicking profiles, 2022 and 2023 seasons (Min. 25 FGA, 38 qualifiers)

Kicker	Total FGA	Estimated Make%	Expected Points
Robbie Gould	38	88.97	2.67
Riley Patterson	56	88.76	2.66
Anders Carlson	36	87.47	2.62
Michael Badgley	35	87.01	2.61
Will Lutz	62	86.96	2.61

Over the course of our study period, every kicker in the league attempted a unique set of field goal attempts that resulted in a similarly unique expected-points yield. In addition, every kicker in the study also accrued an actual tally of points scored on these attempts. By simply comparing the expected point yields and actual point yields of every kicker in the study we can identify which kickers performed above and below expectations. Tables 6 and 7 include the results for the kickers that produced the most and least points above expectations.

Table 6: Highest Expected Points Added Per Kick, 2022 and 2023 seasons (Min. 25 FGA, 38 qualifiers)

Kicker	FGA	Average EP	Expected Points	Points Per Kick	Actual Points	EPA Per Kick	Total EPA
B. Aubrey	38	2.54	96.25	2.92	111	0.388	14.75
C. McLaughlin	69	2.57	169.91	2.74	189	0.277	19.09
C. Santos	60	2.56	151.93	2.80	168	0.268	16.07
J. Elliot	50	2.51	149.64	2.80	165	0.260	15.36
K. Dicker	58	2.55	151.20	2.83	165	0.238	13.80
K. Fairbairn	62	2.52	159.91	2.81	174	0.227	14.09

Table 7: Lowest Expected Points Added Per Kick, 2022 and 2023 seasons (Min. 25 FGA, 38 qualifiers)

Kicker	FGA	Average EP	Expected Points	Points Per Kick	Actual Points	EPA Per Kick	Total EPA
C. Ryland	25	2.46	61.66	1.92	48	-0.547	-13.67
A. Carlson	36	2.62	94.40	2.42	87	-0.205	-7.4
B. Grupe	37	2.52	93.33	2.43	90	-0.090	-3.33
R. Bullock	62	2.60	161.7	2.56	159	-0.044	-2.74
W. Lutz	56	2.66	149.1	2.62	147	-0.038	-2.12

Expected Points Added Per Kick NFL Kickers | 2022 & 2023 Seasons




Figure 5: Individual Actual vs. Expected Points

No kicker in our study performed as well as Brandon Aubrey. The Cowboys' Pro Bowl kicker is more than just a reliable long-range threat. Relative to kickers around the NFL he outperforms league averages more than any player we evaluated. According to our models, Aubrey's set of 38 field goal attempts were expected to yield an average of 2.54 points, but Aubrey produced 2.92 points per attempt; that deviation of 0.382 was by far the biggest average surplus in the set of 38 qualifying kickers in the study period. Even more impressive, his total EPA of 14.75 is comparable to the other players at the top of EPA, but each of them kicked in both 2022 and 2023 while Aubrey was a rookie in the 2023 season and had at least a dozen fewer attempts than any other kicker included. Figure 5 visualizes the expected points per kick and the actual points per kick for all the qualifying kickers in the study.

Chad Ryland, the rookie kicker for New England in 2023, exhibited the least effective performance. According to our models, Ryland's set of 25 field goal attempts were expected to yield an average of 2.46 points, but Ryland produced 1.92 points per attempt - exactly 1 full point per attempt below Aubrey's 2.92; Ryland's expected-points-added deviation of -0.547 was clearly the most negative deviation in the set of 38 qualifying kickers in the study period.

4.3 Kicker Comparisons

Below in Figures 6 and 7 are our rankings of the top 5 and bottom 5 kickers in the league including their overall Kicker Score, individual Command+ and STUD+ scores, and their traditional number of kicks and make % during the two-year study period. Some insights appear when we look at the rest of the players. While Justin Tucker is tied with 5 players for 9th

in traditional make rate (91%), he's our 4th highest graded kicker, while those he is tied with ranked between 7th and 15th. Our 8th ranked kicker, Matt Gay, has a traditional make rate of 87% behind all kickers at 88-89% who we evaluate no higher than 13th and as low as 28th among 8 kickers.

These nuances can have practical implications. Harrison Butker signed a 4-year contract worth the most total money among all kickers in the league, 6.7% higher average salary than any other kicker in the league, and over \$3M more guaranteed than the next closest kicker. To simplify, he has a Kicker Score of 70.3, is ranked 13th of the kickers we assessed, and has average annual contract value of \$6.4M. We believe Ka'imi Fairbairn (74.9/12th/\$5.3M), Jake Elliot (91.7/3rd/\$6M), Cameron Dicker (90.1/5th/\$5.5M), Dustin Hopkins (82.1/7th/\$5.3M), Chase McLaughlin (77.5/9th/\$4.1M), and Nick Folk (77.3/10th/\$3.8M) are all better kickers who also signed contracts this last offseason. [13] While there are more things that go into an NFL player deciding to sign somewhere for a given amount of money (ex. taking a hometown discount to stay where they are familiar), we believe that there is a fundamental misunderstanding of how kickers are evaluated and that our evaluative methods have the potential to teach kickers and teams their true value.

2022 & 2023 NFL Seasons' Top Kickers					
Kicker	Kicker Grade	Command Grade	STUD+ Grade	FG%	Attempts
Brandon Aubrey					
	96.5	94.4	94.6	97	38
Cairo Santos					
	94.7	95.2	84.5	93	60
Jake Elliot					
	91.7	86.1	83.6	93	59
Justin Tucker					
	90.3	88.1	77.3	91	80
Cameron Dicker					
	90.1	85.3	79.9	95	58

Figure 6: Best Kickers

2022 & 2023 NFL Seasons' Worst Kickers					
Kicker	Kicker Grade	Command Grade	STUD+ Grade	FG%	Attempts
Greg Joseph					
	23.2	39.8	34.0	84	61
Jake Moody					
	19.7	32.4	36.2	84	32
Blake Grupe					
	11.8	30.1	23.0	81	37
Anders Carlson					
	10.6	34.5	10.2	81	36
Chad Ryland					
	1.2	3.9	0.1	64	25

Figure 7: Worst Kickers

5. Conclusion

For decades, placekickers have been judged by an overly simple make-miss framework that ignores both situational context and the actual flight quality of kicked footballs. Our results offer compelling evidence that the integration of contextual and spatial approaches promise to improve the status quo of place-kicking analytics in the NFL and beyond.

References

- [1] Pro Football Reference. (2024). *2024 NFL season*. Pro-Football-Reference.com. <https://www.pro-football-reference.com/years/2024/>
- [2] Lopez, M. J. (2016, January 8). It sucks to kick in the cold. *StatsbyLopez*. <https://statsbylopez.com/2016/01/08/it-sucks-to-kick-in-the-cold/>
- [3] Ledoux, J. (2016, January 14) Machine Learning and the NFL Field Goal. *Medium*. <https://medium.com/@jamesledoux/machine-learning-and-the-nfl-field-goal-dde242b6dd39>
- [4] Lopez, M. J. [@StatsbyLopez]. (2022, October 9). Using @NextGenStats ball tracking data, Justin Tucker's game winning field goal crossed the upright with a y-coordinate (width of field) [Tweet]. X. <https://x.com/StatsbyLopez/status/1579311461406769152>
- [5] Marty, R. (2017, March 3-4). *High-resolution shot capture reveals systematic biases and an improved method for shooter evaluation* [Research paper]. MIT Sloan Sports Analytics Conference, Boston, MA. https://cdn.prod.website-files.com/5f1af76ed86d6771ad48324b/5ff4ad56b18b323042079f8e_An%20improved%20method%20for%20shooter%20evaluation.pdf
- [6] *Science of Shooting*. Noah Basketball. (2024). <https://www.noahbasketball.com/methodology>
- [7] Carl S, Baldwin B (2024). *nflfastR: Functions to Efficiently Access NFL Play by Play Data*. R package version 4.6.1.9017, <https://github.com/nflverse/nflfastR>, <https://www.nflfastr.com/>.
- [8] Zippenfenig, P (2023). Open-Meteo.com Weather API [Computer software]. Zenodo. <https://doi.org/10.5281/ZENODO.7970649>.
- [9] Kuhn, Max (2008). "Building Predictive Models in R Using the caret Package." *Journal of Statistical Software*, **28**(5),126. [doi:10.18637/jss.v028.i05](https://doi.org/10.18637/jss.v028.i05), <https://www.jstatsoft.org/index.php/jss/article/view/v028i05>.
- [10] Drovetto, T. (2023, February 28). *Steven Hauschka explains how cold weather impacted the Seattle Seahawks' kicking game*. Seattle Seahawks. <https://www.seahawks.com/news/steven-hauschka-explains-how-cold-weather-impacted-the-seattle-seahawks-k-148292>
- [11] Branch, J. (2010, November 24). *Riddle of the wind at New Meadowlands Stadium*. The New York Times. <https://www.nytimes.com/2010/11/25/sports/football/25wind.html>
- [12] *Graphing calculator*. Desmos. (n.d.). <https://www.desmos.com/calculator>
- [13] Spotrac. (n.d.). *NFL Contracts*. spotrac.com. https://www.spotrac.com/nfl/contracts/_/position/k/sort/initial_guaranteed/dir/desc