

RESEARCH

Open Access



Modeling safe passing regions in American football: integrating ball trajectories and player coverage

Matthias Schilling^{1*}, Rudy Milani¹, Maximilian Moll¹ and Stefan Pickl¹

*Correspondence:

Matthias Schilling
matthias.schilling@unibw.de
¹Department of Computer Science,
University of the Bundeswehr
Munich, Werner-Heisenberg-Weg
39, Neubiberg 85579, Germany

Abstract

Spatiotemporal analysis of professional sports like American football is a research area that has grown in recent years with the emergence of reliable tracking technologies for both players as well as the ball. Although established research focuses on analyzing which team controls which parts of the playing field, usually the ball is not taken into account in this type of analysis. We propose a nuanced quantification of passing target regions by which parts of the field can be reached by each player before the ball arrives when a pass is thrown. A safe target area in American football is characterized as a part of the field that can be reached by a receiving offensive player but no defender. A physics-inspired method and a data-driven method for estimating the area each player can cover are presented. Each is implemented using two approaches for discretizations, one based on squares and one based on hexagons. This analysis makes the spatial dynamics in passing plays explicit and thereby provides coaches, academics, and fans a more precise understanding and a novel perspective. The utility of the resulting areas is evaluated using a publicly available dataset of passing plays from the 2018 NFL season. Passes to a safe target area as identified by the physics-inspired method using hexagons with 0.1 yard long sides are completed 89% of the time compared to 59% if the pass was not considered safe, and only 2 out of 410 interceptions occurred when the pass was considered safe. Furthermore, passes that are considered safe according to the data-driven method, using 0.1 yard by 0.1 yard squares, result in 2.02 additional yards per attempt and 0.25 more expected points added per attempt. The identified safe passes are less likely to result in bad outcomes and on average lead to more productive results.

Keywords Spatio-temporal data, Player tracking data, Football, American football, Pitch control, Passing, Data-driven, Motion modeling, Grid-based

Introduction

The competitive environment of professional sports is at the cutting-edge of the latest technological solutions to gain a better understanding of their own players as well as the opponents [1]. The largest sports venues are nowadays equipped with systems that

© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

track every movement of all players on the field at all times. This technology has also opened new avenues for research in the field of sports analytics, as this tracking data have become publicly available [2].

While spatial analysis using player positions in other sports, e.g. soccer, has a long history of established research [3], the largest US sport, American football, still lags behind. The National Football League (NFL) also recognizes this problem and created the annual NFL Big Data Bowl (see [4]) to facilitate more tracking data analysis in American football.

In many team sports, including soccer and American football, space is an important concept, as there are multiple players but only one ball - most players, most of the time do not control the ball but focus on positioning and on controlling space. Tracking data covering player as well as ball movements allows for advanced analyses of player movements and positioning. Consequently, tracking data is used to investigate different aspects of space, including which team controls how much of the limited available area, improve tactical understanding, and enable advanced passing analysis ([2]).

In American football there are two main ways to advance the ball and score: passing and running. In the 2018 season, there were 121, 737 yards¹ gained by 17, 671 passing plays and only 58, 601 by 13, 272 rushing plays. Passing plays are not only more frequently used, but also more productive. Each play resulted on average in 7.4yd compared to 4.4 for the rushing offenses in 2018 according to [5]. However, passing can also result in the worst case in an interception by the defense. An interception does not only mean no progress towards scoring points, but the offense even loses possession all together without scoring any points and furthermore potentially leaves the opponent in a good position to score themselves. It is therefore paramount avoid these interceptions as much as possible.

An unique characteristic of passing in American football requires that a pass is caught before hitting the ground or it is ruled incomplete leading to the play being over without any getting closer to scoring. While a successful pass is generally a pass that is completed to a teammate, both in the American football literature [6, 7] and other sports like soccer [3, 8, 9]. For a successful pass, the possible target areas are therefore limited by how fast and how far the ball can be thrown on one hand and by how far the players can travel during that time on the other. This sets American football apart from other team sports.

However, passing analysis in American football focuses on player positioning, including relative position between receiver and defender [6, 10, 11], receiver and passer [11] and spatial control [11]. Furthermore, these approaches are concerned with the outcome of the pass, i.e. the likelihood of a completed pass. Furthermore, in soccer, [12] analyze passing safety of target regions on the field based on the difference between the time of arrival of teammates and the time of arrival of opponents without taking the ball into account, since the ball in soccer does not need to be caught directly. There is some work analyzing passes using the balls trajectory in soccer [9, 13, 14], demonstrating the utility and feasibility of passing analysis using ball trajectory data.

In this work, we identify safe passing target regions by combining the area that each player can cover with the range of the pass. This approach allows automatic identification of areas on the pitch that are safe to target with a pass and those that are not. Given

¹ In American football distance is given in yards. 1 yard $\approx$ 0.915 meters.

a time interval for the pass to reach a receiver, the area each player can reach at the end of that interval can be estimated. Calculating the union of reachable areas for offense and defense results in the area where a receiver could catch the ball and the area covered by the defense respectively. The difference between offense and defense describes the portion of the reachable area where offensive players are expected to catch the ball unbothered by any defender. Finally, the overlap between that area and the throwing range of the football is the portion that can be safely targeted with a pass.

American football is a complex sport, even if focusing solely on passing. There are different ways of evaluating the quality of a passing attempt and there are various aspects influencing its outcome. For example, the further the pass is away from the endzone, the larger the area the defense has to cover and thus the more difficult to defend all receivers well. If it is raining heavily and the ground is wet, quick direction changes or steep acceleration are more difficult and in a two-minute-drill, the offense often prefers passes close to the sideline. All of these different factors build on the passes that are thrown to be actually caught, and the competition between the receivers to get open against the defenders trying to cover them. Understanding when a receiver is open or which areas covered better or worse follows the same idea independent on whether it is a short or a long pass or if the drive is just starting with 90 yards to go or already in the redzone.

Partitioning the playing area according to which team can reach the different spots on the field provides a novel perspective on spatial dynamics that not only allows coaches, analysts, scientists, and fans to understand positioning and spacing better when analyzing a single play. At the same time, it enables objective analysis of defensive coverage and offensive success of creating open areas for passing. The same concept can be used to help coaches easily and transparently identify patterns and tendencies of opponents as well as their own teams performance.

The presented discrete set based approach is designed to be adaptable to practitioners needs, as it allows to balance efficiency and granularity. Since the grids are consistent across plays, this applies both to how quickly and fine-grained a single play can be evaluated. When analyzing multiple plays, the grids can be easily combined for many plays with a single aggregation operation.

The area each player can cover is estimated building on previous work modeling player movement, either using physics-inspired models [3, 12, 15, 16], or data-driven approaches [17, 18]. Either way, the resulting regions need to be aggregated. Since the union of several circles or polygons is not necessarily convex and calculating the difference between non-convex planar sets is complex to solve, we provide an efficient alternative by using a discrete grid-based approach that transforms this complex problem into simple set operations.

This work is structured as follows: an overview of the literature is given in Section "[Related work](#)". A description of the proposed methodology is presented in Section "[Methods](#)", consisting of two parts; first, we describe two alternative approaches to estimate the area that players can reach in continuous space. This approach provides a novel perspective on large-scale data-driven analysis of passing in American football that respects the practical restrictions of the game. The second part details how this approach is implemented using either a square grid or a hexagonal grid. Subsequently, Section "[Computational experiments](#)" describes the evaluation and implementation followed by an analysis of the results for the 2018 NFL season in Section "[Results](#)". Finally,

Section "[Conclusion and further work](#)" provides concluding remarks and an overview of limitations and possible directions for future work.

Related work

There are several components that form the foundation of this work, starting with the concept of spatial control in team sports. Table 1 lists the relevant literature and shows that spatial control is a common avenue of research. Voronoi diagrams are a common starting point for spatial analysis of team sports, both as a main tool for analysis [19–21] and as a building block for a more complex model [3, 15, 22]. The evaluation of space in team sports started with [3], who propose a quantitative analysis of teamwork in soccer based on enhancing Voronoi diagrams with a physics-inspired motion model. [15] build on the concept of dominant regions proposed by [3] with a modified model of motion for field hockey. [23] investigate the accuracy of the model proposed by [15] for soccer players by comparing the model to the observed trajectories and find that [15] accurately predicts the trajectories to end within circular regions.

[16] compare four different physics-inspired player movement models for estimating the boundary of the reachable area of soccer players within a set time interval. These models, based on physical laws of motion, include constant speed, uniformly accelerated motion with bounded and unbounded velocity, as well as a non-uniform model with the first segment based on the initial velocity and the based on uniform acceleration. They conclude that while circles around the initial positions result in an upper bound of reachable areas, the estimate can be improved if the maximum speed of the players and initial velocity are considered.

The research of player tracking accelerates with growing technical capabilities, both in frequency and variety. As an alternative to physics-inspired modeling, [24] and [17] employ a data-driven probabilistic formulation of spatial control and player movement

Table 1 Overview of approaches in sorted by year of publication

Reference	Year	Pitch Control	Pass	Player Traj.	Ball Traj.	Sport	Dataset
[3]	2002	✓				S	Not mentioned
[15]	2005	✓				FH	Multiple scenes
[19]	2016	✓				BB	Not mentioned
[26]	2017		✓			S	726 matches
[13]	2017		✓	✓	✓	S	38 games
[18]	2018	✓		✓		S	20 matches
[28]	2018		✓			S	82 matches
[17]	2019	✓		✓		S	1 match
[27]	2019		✓			S	18 matches
[25]	2019	✓		✓		AF	91 games
[29]	2020		✓			AF	154,908 frames
[6]	2020		✓			AF	91 games
[24]	2021	✓				S	4 matches
[12]	2021	✓	✓	✓		S	45 games
[11]	2021	✓	✓			AF	91 games
[16]	2022			✓		S	3 games
[9]	2022		✓	✓	✓	S	918 games
[20]	2023	✓				S	2,243 games
[14]	2024		✓		✓	S	34,938 shots
This work		✓	✓	✓	✓	AF	253 matches

S: Soccer, FH: Field Hockey, BB: Basketball, AF: American Football

that grants control to the player with the highest likelihood of arriving at a given point on the field first. [20] combine insights from a Voronoi tessellation and manually created zones to analyze patterns in pitch control. [25] propose a Deep Learning approach for estimating pitch control based on predicted player trajectories in a competition entry to the NFL Big Data Bowl 2019.

While other approaches aim to create a binary partition of the pitch, assigning points either to one team or the other, [18] use a multivariate Gaussian to quantify occupation as a value between 0 and 1 instead of a binary assignment to one team or the other. Furthermore, [12] extend the binary view by distinguishing between safe and risky controlled areas. They characterize those areas as risky, in which the minimum arrival time of the opposing side is close to the minimum arrival time of the controlling team. This concept of risky implies the potential of interference by an opponent when receiving a pass.

The investigation of passing is another topic of interest in a variety of team sports. As Table 1 shows, analysis in passing is done both with and without relying on pitch control. [26, 27] and [28] design spatio-temporal features to quantify passing in soccer. [9] builds on event data including a model of player movement to analyze passes and [12] and [24] connect pitch control and pass evaluation directly by evaluating the passes depending on which team controlled the target area of the pass.

While soccer passing analysis can fall back on decades of spatial analysis, the NFL has a monopoly on the data and only recently made it available to researchers as part of the Big Data Bowl. Consequently, passes in American football are less thoroughly investigated. In general, the player who passes the ball is always the Quarterback (QB). Investigating QB performance, and therefore passing, can be done in different ways. On one hand, [29] evaluate QB decision making including catch probability. On the other hand, [11] evaluate QBs based on receiver openness, which they define as a combination of pitch control (as proposed by [18]), distance to the sideline and to the closest defender. [6] introduce a Bayesian approach to hypothetical completion percentage of passes that were never thrown. On the defensive side [30] use tracking data to evaluate defensive players, including against passing plays.

It needs to be mentioned, that there are several companies that are providing data analysis as a service to teams, either using tracking data or video data. Companies like Catapult, Chyronhgo, Skill Corner or Kinexon have a variety of products including Soccer and American football. However, it not publicly available, what their products capabilities are and how they compare to the academic state of the art represented in the scientific literature.

In summary, there are a variety of approaches analyzing passing and space, both individually and in combination. Table 1 provides an overview and shows that most of the literature does not focus on American football. This work aims to bridge the gap between spatial and passing analysis for other team sports and the specific requirements of American football passing.

While many existing approaches model passing safety based on spatial analysis of players, we model the interaction of the range of the football with the area the players can cover while the ball is in the air. Identification of safe passes provides an avenue for advanced analysis of passing in American football that goes beyond available metrics.

The proposed approach compares two intuitive and transparent approaches to estimate reachability of players and ball. First, a model based on ideal kinematics, which is anchored in prior research but due to its simplified nature can overestimate the reachable area. Second, a data-driven approach based on historic observations in order to capture a more realistic estimate of relative movement.

Methods

The core idea of this work is based on the area each player can cover within a given time interval in combination with the distance the ball can be thrown. The reachable area of a player can be estimated as a function of the initial position of the player and the length of the time interval.

Following the literature of the field, we compare two approaches of estimating the reachable regions; the first approach is inspired by the laws of motion while the second is a data-driven approach based on the actual trajectories observed from players. In both cases the areas are mapped to discrete masks of the playing field with a square grid and a hexagon grid, where each cell can either be reachable or not reachable by receiver, defender and ball.

Physics inspired region estimation

The area each player $p = 1, \dots, P \in \mathbb{N}$ can cover can be estimated based on the widely known formula of uniformly accelerated motion $x_t = x_0 + v_0 t + 0.5 a t^2$, where x_0 is the initial position, v_0 the initial velocity and a the acceleration. The goal is an estimation of reachability and taking initial velocity and direction of motion into account leads to a smaller reachable area. We assume that the initial velocity of the players is $v_0^p = 0, \forall p = 1, \dots, P$ and they can accelerate in every direction equally, which overestimates the reachable area, but for the purpose of this approach, overestimation is preferable to underestimation. The area that player p covers is therefore estimated as a circle centered at their starting position (x^p, y^p) with a radius r^p depending on their acceleration a^p and top speed v_{max}^p as follows:

$$r^p = \frac{1}{2} a^p \bar{t}^2 + v_{max}^p (\max\{0, t - \bar{t}\}) \quad (1)$$

where $\bar{t}$ denotes the time at which the maximum velocity is reached. We use speed to mean the magnitude of the velocity.

Furthermore, the acceleration a^p and maximum velocity v_{max}^p are assumed to be constant and identical for all players, i.e., $a^p = a, v_{max}^p = v_{max}, \forall p = 1, \dots, P$. The values for a, v , and v_{max} are set based on the data as described in Sect. "Experimental setup".

Using constant and identical values provides stable values for all players, even if there is not a lot of data available at the cost of slightly overestimating some players athletic prowess or their capabilities to accelerate even if already moving at a high speed.

Applying these parameters to an example interval of $t = 2.5$ seconds, chosen as a midpoint of considered interval durations, results in the estimated boundary shown in Fig. 1. The figure shows the relative positions from the tracking data of players after 2.5 seconds of movement as blue points. When not taking initial speed and movement direction into account, as overall more than 99.9% of players end within this estimated boundary. This is also true for other time intervals. From the second time step on, more




Fig. 1 Random sample of relative player positions after 2.5 seconds starting from point (0, 0) and disregarding initial movement speed and orientation. The circle describes the calculated reachability boundary. The diagram shows 1,000,000 randomly drawn observations of player movement from passing plays of 253 NFL games with a duration of 2.5 seconds contained in the dataset without further restrictions on players or game state

than 99.5% of player movement falls within the estimated boundary. After a single time step, 1.5% of movements exceed this boundary due to the initial velocity being high. The impact of the initial velocity decreases with increased interval length, which is reflected in the portion of movement ending outside of the estimated boundary declining strictly monotonically (Fig. 2).

This provides an estimated boundary of the area each player can reach within a given time interval. The union of all players of one team is the region they can reach as a whole within the given time interval and provides a clear understanding of the area that team can cover. Figure 3 shows the resulting areas for both teams in the top row.

The same process is applied to the football, with acceleration a^f and maximum velocity v_{max}^f , resulting in the image on the bottom left of Fig. 3. Combining the different reachable regions provides a spatial division of the pitch based on which team can reach each area within the flight time of the ball. In Fig. 3, red means the offensive team can reach that area, green means the defensive area can reach that area, blue means the football can be thrown that far. An overlap between green and blue areas (but not red) is shown as teal, which is the mix of blue and green, and means that both ball and defense can reach that area but not the offense. Similarly, purple is a combination of blue, indicating the ball can reach that area and red, indicating a receiver can reach that area, but not green, indicating no defender can reach that area. So purple areas are safe to target.




Fig. 2 Random sample of relative positions after 2.5s, observed of all players, color coded by initial velocity, given in yd/s, and corrected for initial direction of movement to align vertically downwards. For reference, the circle of the pixel based method is shown in pink. Clearly, the outliers have to be eliminated to get an accurate region

This physics-inspired method is based on the assumption that the initial velocity is zero and the player can accelerate in all directions equally. This approach provides an efficient and transparent way to calculate the area each player can reach. However, as [16] already observed, it overestimates the area players actually end up in. In order to estimate the minimal reachable regions, a data-driven alternative approach is employed.

Data-driven region estimation

The estimation of areas on the pitch reachable by some players can be based on historic observations of past player trajectories, consisting of player and ball position, speed and direction of motion. When directly computing the reachable region as the minimal convex area containing all past player trajectories, this results in a tighter boundary and therefore in a more precise analysis.

When disregarding initial speed and direction of movement, the physics-inspired approach provides a good boundary. In contrast, if those variables are taken into account, the resulting areas are no longer circular. Figure 2 shows the starting point and previously estimated boundary, but the points are colored according to the initial velocity and the relative positions are corrected for the initial movement direction. This figure suggests that the larger the initial velocity the less appropriate it is to model the reachable region of that player as a circle, as larger initial velocities coincide with increasingly




Fig. 3 Reachable Regions for each team based on circular boundaries of the first passing play of the season. The colored areas represent the area that can be reached within the time interval. The first three plots are aggregated together to get the bottom right image. The triangle corresponds to the position of the football at the time the pass is thrown. The offense is playing right to left

smaller regions towards the edge of the reachable region, i.e. the previous approach overestimates the area this player can cover.

Instead of estimating each region on a physics-inspired model of motion, an alternative approach is to solely rely on the information provided in the data. The goal is to estimate the boundary of the area a player can reach within a time interval. This can be estimated by computing the smallest convex polygon containing the set of all vectors of relative position after a time interval of the same length. This set can be extracted directly from the data and the smallest polygon containing all these vectors is the convex hull.

Before computing the convex hulls, the relative positions for each player at each play sequence and for all time intervals need to be calculated, where the intervals of interest are bounded by the maximum duration in time steps of a pass. Each trajectory is grouped by game, player and play, resulting in a single group of data points $g_{gameId,playId,playerId}$, i.e. one group of points corresponds to one trajectory per player and play, can then be divided into $|g| - t - 1$ sub-intervals of length t , with $|g|$ the length of the group, or trajectory. The relative position (x_r^p, y_r^p) of a player p starting at the initial position (x_0^p, y_0^p) and being at (x_t^p, y_t^p) after t time steps is given by

$$\begin{pmatrix} x_r^p \\ y_r^p \end{pmatrix} = \begin{pmatrix} x_t^p \\ y_t^p \end{pmatrix} - \begin{pmatrix} x_0^p \\ y_0^p \end{pmatrix} \quad (2)$$

resulting in a set of vectors $(x_r^p, y_r^p), \forall p = 1, \dots, P$ that contains one vector for each unique interval of length t and each player playing at that time.

This set can be further improved by aligning the vectors according to the initial direction of movement, given by the angle $\theta \in [0^\circ, 360^\circ]$ to the y -axis. The relative position after correcting for initial direction of movement (x_n^p, y_n^p) is obtained by rotating (x_r^p, y_r^p) by $-\theta$

$$\begin{pmatrix} x_n^p \\ y_n^p \end{pmatrix} = \begin{pmatrix} \cos(-\theta) & -\sin(-\theta) \\ \sin(-\theta) & \cos(-\theta) \end{pmatrix} \cdot \begin{pmatrix} x_r^p \\ y_r^p \end{pmatrix} \quad (3)$$

Subsequently, the resulting points are then grouped according to the initial velocity. This is done by discretizing the velocity into equal-width buckets and assigning each point to the bucket according to its initial velocity value. As an optional step, it is recommended to eliminate outliers using depth-based outlier detection as described in [31, section 8.2.3]. The convex hull of each set of points is then computed. The resulting convex hulls are used to estimate the limit of the relative positions a player can end up after a specific time interval, i.e. their reachable region. The estimated hulls computed from the points shown in Fig. 2 are visualized in Fig. 4. An overview of this algorithm for data-driven calculation of the regions is provided in Algorithm 1, given a dataset D of player tracking data including sequences of player positions (x, y) , direction of movement θ and velocity v and the maximum duration of a pass $max_pass_duration$. each point is




Fig. 4 Convex hulls of all positions observed after 2.5s color coded by initial velocity, given in yd/s , and corrected for initial direction of movement to align vertically downwards. For reference, the circle of the pixel based method is shown in pink

associated to a game with a *gameId*, a play with a *playId*, and player with a *playerId*. That data is divided into groups using an arbitrary *groupby* function and an algorithm ConvexHull is used to compute the convex hull of a set of points, this work uses [32]. The Algorithm is presented without elimination of outliers.

Require: $\mathbf{D}$, the player tracking data

Require: *max_pass_duration*, maximum number of time steps

```

1: for  $t \in \{1, 2, \dots, \text{max\_pass\_duration}\}$  do
2:    $\text{points\_collection}[t] \leftarrow$  an empty collection (e.g., list or dictionary)
3:    $\text{hulls\_collection}[t] \leftarrow$  an empty collection (e.g., list or dictionary)
4: end for
5:  $\mathbf{G} \leftarrow \text{groupby}(\mathbf{D}, \{\text{gameId}, \text{playId}, \text{playerId}\})$ 
6: for  $t \in \{1, 2, \dots, \text{max\_pass\_duration}\}$  do
7:   for  $g \in \mathbf{G}$  do
8:     for  $i \in \{0, 1, \dots, |g| - t - 1\}$  do  $\triangleright$  Loop over indices where a  $t$ -step is possible
9:        $x_r \leftarrow x_{i+t}^g - x_i^g$ 
10:       $y_r \leftarrow y_{i+t}^g - y_i^g$ 
11:       $x_n \leftarrow \cos(\theta_i^g) \cdot x_r - \sin(\theta_i^g) \cdot y_r$ 
12:       $y_n \leftarrow \sin(\theta_i^g) \cdot x_r + \cos(\theta_i^g) \cdot y_r$ 
13:      Append  $(x_n, y_n)$  to  $\text{points\_collection}[t]$   $\left[ \lfloor 0.5v_i^g \rfloor \right]$ 
14:    end for
15:  end for
16:  for  $\text{points} \in \text{points\_collection}[t]$  do
17:     $h \leftarrow \text{ConvexHull}(\text{points})$ 
18:    Append  $h$  to  $\text{hulls\_collection}[t]$ 
19:  end for
20: end for
```

Algorithm 1 Data-driven region estimation

These convex hulls are then used as estimates of the reachable region of a player. As the hulls are computed with the points that were corrected for the initial movement direction, the resulting hulls need to be rotated to align with the actual direction of movement of the player in question.

Similar to the circle regions, combining the individual hulls of one team leads to the reachable region of the whole team. Figure 5 shows the resulting regions for offense, defense, and ball in the top right, top left, and bottom right, respectively. In comparison to Fig. 1, it shows a much more jagged outline of the regions, clearly influenced by direction of movement of the players. Furthermore there is a red area around the 50 yard line at the top, where the receiver is open behind the defense. However, in this example play the ball cannot reach the receiver in time.

Grid-based mapping of reachable regions

Both the physics-inspired and the data-driven approach require calculating a set of reachable regions, either in the form of circles or polygons, for players and ball. These regions are then combined by calculating the union for each team and subsequently calculating the overlap between ball and offense and subtracting the defenses area to identify the safe target regions. Calculating the union, overlap and difference is a complex problem in itself. A more efficient approach is to solve the problem in a discretized setting numerically.

Previously we already presented images of the estimated regions, each image can also be interpreted as a binary mask. Each pixel indicating whether a pixel-sized portion of the field is within the reach of a team or not. An American football field is 120 yards (yd)




Fig. 5 Reachable Regions for each team based on observation based convex hull boundaries for the first passing play of the season. The colored areas represent the area that can be reached within the time interval. The first three plots are aggregated together to get the bottom right image. The offense is playing right to left

long and 53.3yd wide, so in an image with a size of 1200×533 pixels, one pixel corresponds to a square patch of size $0.1\text{yd} \times 0.1\text{yd}$. We define $o_{ij}, d_{ij}, b_{ij} \in \{0, 1\}$ as the binary variables that represent whether the pixel (i, j) is reachable by an offensive player, by a defensive player, and by the ball, respectively. As the safest passes are those, that can be caught by a receiver but not by a defender, then we define the *safeness* p_{ij} of a pixel (i, j) as follows:

$$p_{ij} = o_{ij}(1 - d_{ij})b_{ij}. \quad (4)$$

Therefore the safe target areas are those, where the mask is $o_{ij} = 1$ for offensive players and $d_{ij} = 0$ for defensive players. In addition, the ball needs to be able to reach the target as well, i.e., $b_{ij} = 1$. In Fig. 3 the aggregated mask shows the safe regions where $p_{ij} = 1$ in purple, since those areas with both red($o_{ij} = 1$) and blue($b_{ij} = 1$) but not green ($d_{ij} = 0$) result in purple pixels.

The masks for the different regions are computed utilizing drawing functions, which is equivalent to analytically calculating the overlapping regions followed by a numerical evaluation at the center of each pixel. The image size can be chosen such that one pixel is equivalent to a single grid cell of the chosen size. The circle shaped regions can be mapped to the grid in this case utilizing a disk-drawing algorithm, like the Bresenham midpoint circle algorithm [33]. It follows the idea, that all pixels on a circle are roughly the same distance from the center and directly adjacent to two other pixels. Similarly,

the convex hulls can be mapped to the grid by drawing a polygon, e.g. using Bresenham's line drawing algorithm [34].

Mapping the regions on a mask using disk or polygon drawing is implicitly utilizing a square grid, where each pixel corresponds to one grid cell. However, as [35] observed, in the context of geoinformation systems hexagonal grids provide superior geometric properties when trajectories and shapes with non-straight edges are considered, as they are closer to circles and have equidistant neighbors. Therefore, a similar approach based on a hexagonal lattice is tested as well.

Compared to the square grid used previously, also referred to as *pixel* based, this approach requires a few modifications. First of all, the coordinates (x, y) have to be translated to hexagonal grid coordinates (h_i, h_j) , with $h_i, h_j \in \mathbb{N}$ representing the hexagon with the center point at $(h_i \cdot 1.5 \cdot \ell, h_i \cdot 0.5\sqrt{3} \cdot \ell + h_j \cdot \sqrt{3} \cdot \ell)$, where ℓ denotes its side length. When directly computing this distance on the hexagon grid, i.e. touching hexagons have a distance of 1, hexagons with the same distance from a center point form concentric rings of hexagons around that center. Therefore, since we cannot use drawing directly, we instead compute all k concentric rings with distance to the center less or equal to k to obtain the reachable grid cells. k depends on the radius r and side length ℓ , and is computed as

$$k = \left\lceil \frac{r}{\frac{2\ell + \sqrt{3}\ell}{2}} \right\rceil \quad (5)$$

where the ceiling function $\lceil \cdot \rceil$ is used to overestimate the resulting region and the mean of the width and height of the hexagon is used as an approximation of the diameter.

Mapping a circle to the grid this way results in a region that is not circular, but a hexagon inscribed in the circle. This approach is easy to compute, as the indices of all hexagons that are part of the region can be directly determined from the index of the center hexagon with index (h_i^c, h_j^c) and the number of rings k . Following the idea that all hexagons indices (h_i, h_j) must satisfy

$$h_i^c - k \leq h_i \leq h_i^c + k \quad (6)$$

$$\max(h_j^c - k, h_j^c - h_i + h_i^c - k) \leq h_j \leq \min(h_j^c + k, h_j^c - h_i + h_i^c + k) \quad (7)$$

where the constraints represent the bounding edges of the hexagon. The intuition is that the hexagon is a parallelogram with the two corners furthest from the center trimmed away, so that all indices on the boundary are equally far from the center. Figure 6 shows a mask calculated following this approach.

In order to map the convex hulls to the hexagon grid, each hexagon center is checked whether it is inside the hull or not. Hexagons having center points within the hull are considered part of the reachable region while the rest is considered outside. This approach allows for the creation of a mask, from which point on the aggregation of the different masks and subsequent spatial classification is performed analogous. Figure 9 shows the results for the four proposed methods for comparison.




Fig. 6 Combined hexagonal regions with polygon side length $\ell = 1.0\text{yd}$. Offense is playing right to left, the safe target area is shown in purple

Table 2 Summary statistics of the *raw* tracking data of the 2018 NFL season

Variable	Mean	Median	Min	Max
x [yd]	59.8	59.8	-3.88	124
y [yd]	26.7	26.7	-9.30	62.7
velocity [$\frac{\text{yd}}{\text{s}}$]	3.04	2.77	0.00	42.7
acceleration [$\frac{\text{yd}}{\text{s}^2}$]	1.94	1.73	0.00	70.2
θ [°]	180	180	0	360
pass duration [s]	13.7	12.0	0.1	4.6

Computational experiments

The described methods are validated on passing data of the 2018 NFL season (see [36]). This section starts by describing dataset and evaluation metrics we use. Afterwards we describe the results and compare the approaches.

Dataset

This work uses data published by the NFL for the NFL Big Data Bowl 2021 [36], as it is publicly available and remains relevant for the purpose of this analysis, since the analysis focuses on player movement. It contains 2D-position, speed and movement direction every 100 ms for all players (except linemen) and for the football. An overview of the 18,309,388 rows of raw tracking data is provided in Table 2. Data cleaning procedures, e.g. removing extreme and unlikely high values for velocity and acceleration, are described in Section "Experimental setup".

Additional annotations for game events like the ball snap, a forward pass being thrown and the play outcomes are recorded, including the time step at which they occur. The tracking data is complemented by further information for games and players as well as individual plays that provide further context. Specifically, the data for each play's result is relevant here, as it provides the basis for the analysis.

The dataset contains passing plays of 253 of 256 games of the 2018 NFL regular season, with on average 76 plays each. In total there are 19,239 plays in the dataset. After

eliminating plays without a forward pass or when no QB is on the field, there remain 17,643 plays, consisting of 11,208 complete, 6,025 incomplete and 410 intercepted passes. The dataset does not contain other types of plays such as running plays.

Evaluation metrics

There is no direct metric that we can use to quantify how well the resulting regions are modeled. In this section, we provide a set of passing metrics that allow us to measure the quality of the proposed methods. For a pass to be safe, we expect several properties to be true. Defending players should be unlikely to interfere with the receiver catching the ball. This suggests that the completion percentage (Eq. (8)) should be higher for safe passes than for non-safe passes. Furthermore, this means that the ball is less likely to be intercepted. Since no defender is close enough to challenge the catch, they can also not immediately tackle the catching player, leading to an increased output of yards after catching the ball. Safe passes, due to the lack of bad outcomes and the potential for yards after the catch, should result in a higher average Expected Points Added (EPA) than unsafe passes.

To evaluate the results of the experiments, we define the following metrics based on whether a pass was classified as *safe* or *not safe* on the one hand and the result of the pass on the other. Possible outcome are complete, incomplete or intercepted. Thus, we define C, I, I_N as the sets of complete, incomplete, and intercepted passes, respectively. The most obvious metric for evaluating passes is based on the success of the pass.

Completion percentage

The completion percentage is the percentage of passes that are completed of all attempted passes. The fraction of completed passes is computed as

$$Comp = \frac{|C|}{|C \cup I \cup I_N|} \quad (8)$$

where $|\cdot|$ represents the cardinality of the specific set. From this value (always in the interval $[0, 1]$), we define the completion percentage as $C\% = 100 \cdot Comp$. When only those passes classified as *safe* are considered, this is equivalent to the widely used concept of precision

$$\begin{aligned} Prec_C &= \frac{|S \cap C|}{|(S \cap C) \cup (S \cap \bar{C})|} \\ &= \frac{|S \cap C|}{|S|} \end{aligned} \quad (9)$$

where S is the set of passes that are classified as safe, $S \cap C$ is the set of completed passes that are classified as safe, $S \cap \bar{C}$ is the set of passes that are classified as safe but were not completed, i.e., incomplete or intercepted $\bar{C} = I \cup I_N$.

Specificity

The other important aspect of a safe target is the possibility of the pass being intercepted by a defender. A safe pass is expected to have a much lower chance of being intercepted. The metric capturing this, the *specificity* with respect to the interceptions, is computed as

$$\begin{aligned}
 Spec_{I_N} &= \frac{|\bar{S} \cap I_N|}{|(\bar{S} \cap I_N) \cup (S \cap I_N)|} \\
 &= \frac{|\bar{S} \cap I_N|}{|I_N|}
 \end{aligned} \tag{10}$$

where $\bar{S}$ is the set of passes classified as not safe, $S \cap I_N$ is the set of safe passes that are intercepted and $\bar{S} \cap I_N$ is the set of not safe passes that are intercepted. The specificity describes the portion out of all interceptions, that happen as the result of an unsafe pass. A specificity close to 1 means interceptions occur mostly on passes that are not considered safe. Since the pass can be tipped or the receiver can make a mistake while catching, there remains a small chance of an interception occurring even on a safe pass.

Recall

The third metric is the portion of all completed passes that are considered safe. This is captured by the recall (or sensitivity) of safe passes.

$$\begin{aligned}
 Rec_C &= \frac{|S \cap C|}{|(S \cap C) \cup (\bar{S} \cap C)|} \\
 &= \frac{|S \cap C|}{|C|}
 \end{aligned} \tag{11}$$

Although most safe passes are expected to be completed, not in every play will there be a targeted safe receiver. As the proposed approach is designed to identify passes with a low chance of failure, the precision is more important than the recall. Therefore the results are expected to reflect this with a precision value close to 1.0 and a recall that is lower.

EPA

A football-centric approach to comparing the different methods can be achieved using an established advanced metric called *Expected Points Added (EPA)*, which has been previously used by [10] and [37]. EPA is an approach that estimates the expected points of an offensive possession based on distance to the end zone, down, yards to the next first down and home-field advantage. For each play the difference in expected points can then be calculated, which is the EPA of that play. With the expected points for touchdown set at seven, i.e. six points for the touchdown plus one point after touchdown, a drive is between -7 points for a defensive touchdown and 7 points for an offensive touchdown. Therefore the Expected Points can have a value between -14 and $+10$, where $+10$ corresponds to a scoring drive from the own 1 yard line starting at -3 expected points. In this data set, the observed values of EPA are between -11.94 and 8.63 . EPA lends itself as an evaluation metric for the proposed approaches, because it quantifies the outcome of a play in terms of scoring points and is readily available as it is part of the dataset.

Yards per attempt

A more direct analysis is based on the yards per attempt of each play to provide an application focused interpretation of the distinction between safe and not safe passes. On one hand providing a tangible quantification of passes and allowing a comparison of the

different approaches on the other. Three different metrics are used in this section: Yards per attempt, attempted air yards and completed air yards per attempt.

Yards per attempt (Y/A) is a widely used passing metric calculated as the sum of yards divided by the number of passing attempts. This number provides the average gained yards of the play including both the pass and the running with the ball after the catch. The other two metrics are only taking the passing distance into account, not the gains after the catch. *Air Yards* are calculated as the difference between the yard line where the ball was caught and the yard line where the ball was snapped, not the actual euclidean distance between the throwing position of the QB and the catching position of the receiver. Attempted air yards (aAY/A) describe the attempted pass distance regardless of result, the completed air yards (cAY/A) incorporate whether or not a pass was completed. Comparing the completed air yards with the total yards per attempt as expressed by Y/A provides insight into the potential for yards after the catch, i.e. the amount of yards the catching player can gain after catching the ball by running and carrying the ball. The yards after the catch reflect the idea that a pass thrown to a receiver that is unbothered by a defender can gain more additional yards than a receiver that is closely guarded.

Experimental setup

We compare the proposed methods to a pitch control baseline approach, a Voronoi diagram, which assigns each spot on the pitch according to which player is closest. We choose it as a baseline because it is both the starting point for many pitch control approaches and commonly used in the literature for comparison ([17, 20, 24]). The decision of whether a pass is safe or not is based on the Voronoi diagram calculated at the time of the pass. The pass is safe exactly if it ends in a Voronoi region of an offensive player.

The first proposed method estimates the reachable area as a circle, where the radius depends on the velocities and the acceleration of the player or ball of interest. Since the actual highest velocity and acceleration occurring in the data are unrealistic values, $42.74 \frac{\text{yd}}{\text{s}}$ and $70.15 \frac{\text{yd}}{\text{s}^2}$, respectively, the maximum speed for the players is set to the 99.9th percentile of the players accelerating over $1 \frac{\text{yd}}{\text{s}^2}$ and moving over $1 \frac{\text{yd}}{\text{s}}$ respectively, resulting in $a^p = 8.31 \frac{\text{yd}}{\text{s}^2}$ and $v_{max}^p = 9.86 \frac{\text{yd}}{\text{s}}$. The speed and acceleration of the football is set analogously to $a^f = 22.17 \frac{\text{yd}}{\text{s}^2}$ and $v_{max}^f = 25.91 \frac{\text{yd}}{\text{s}}$. It has to be noted that the dataset uses yards per second as a unit of velocity and yards per second squared in the case of acceleration, which has the advantage of being easily mapped to the grid since the American football field, and thus the grid, is set in yards as well. With the acceleration and maximum speed given, the radius for the interval of interest can be computed and the circular area for each player is mapped to the grid, with the center of the circle given by the current player position.

The second method is based on the observed player trajectories in the data which allows for incorporating the initial velocity and direction of movement. In this approach, the initial velocities are grouped in intervals of 2 and the orientation of the movement is used as an integer value. The mask is once again mapped directly to the square grid using polygon drawing as visualized in Fig. 7, while for the hexagons are considered part of the region if their center points are inside the hull. The resulting mask for the first pass of the season is shown in Fig. 8.




Fig. 7 Mask using convex hulls with $\ell = 1.0$ yd square grid with the offense playing right to left. The safe target area is shown in purple




Fig. 8 Mask using convex hulls with $\ell = 10$ hexagon grid. The safe target area is shown in purple

Both methods are implemented on two discrete grids, namely, a square grid and a hexagonal grid. Each grid is parametrized by ℓ . The highest resolution square grid with $\ell = 0.1$ consists of $A = 0.1\text{yd} \times 0.1\text{yd} = 0.01\text{yd}^2$ squares. Similarly the smallest hexagon grid with $\ell = 0.1$ consists of $A = 1.5 \times \sqrt{3}\ell^2 = 0.026\text{yd}^2$ hexagons. All computations are run for $\ell \in \{0.1, 0.2, 0.3, 0.5, 1.0\}$ yards.

Results

The results for the passing plays of the 2018 NFL regular season are described in Table 3. The average for all passing plays that have been considered is provided for comparison in the first row.

The passes classified as *not safe* based on the Voronoi diagram have an average completion percentage of 39%, which is 25% less than the overall average whereas those




Fig. 9 Comparison of the different region estimation methods with $\ell = 1.0\text{yd}$. The safe target area is shown in purple

classified as *safe* are completed at a higher rate 78% of the time. Additionally, the Y/A are almost a yard less than the average.

Figure 10 shows the recall and specificity for the different models with $\ell = 0.1\text{yd}$ in comparison to the Voronoi based classification. While the recall using Voronoi is much higher, the specificity is much lower. With only 0.65, meaning 65% of interceptions occurred on not safe passes, i.e. 35% of interceptions happened on supposedly safe passes, the specificity of this method is mostly filtering out bad attempts rather than identifying good passes. The method presented in this work are conversely showing a high specificity, meaning the passes identified as safe are more unlikely to lead to bad outcomes.

The difference between this work and Voronoi based analysis is further supported when comparing the values for recall. While 77% of completions are classified as safe according to the Voronoi model, our models classify 12%–19% as safe. The reachable region estimation expressly identifies areas that are expected to be reached before the ball by a receiver but the defenders only arrive after the ball. The Voronoi tessellation on the other hand classifies all passes as safe, that are thrown to a location closer to a receiver than any defender, even if a defender is very close and can easily reach the target location before the ball arrives. This is consistent with 11,086 passes classified as safe and 6,557 unsafe according to Voronoi compared to only 1,535 to 2,449 based on reachable region estimation.

Looking at the physics-inspired model on the square grid, as shown in Table 3 in the first block, the first major difference is that the unsafe passes are close to the overall

Table 3 Resulting metrics, ℓ : side length in yard, Y/A: Yards per Attempt, aAY/A: attempted Air Yards per Attempt, cAY/A completed Air Yards per Attempt, C%: completion percentage, $RecC$: Recall with respect to completions, $Spec_{IN}$: Specificity with respect to interceptions, Count: number of observations passes

Method	ℓ	Safe	A/Y	aAY/A	cAY/A	EPA	C%	$Spec_{IN}$	$RecC$	Count
All			7.26	10.64	5.02	0.14	64			17643
circle square	0.1	False	7.20	11.39	5.24	0.13	61	1.00	0.14	15846
		True	7.77	4.08	3.08	0.19	85	1.00	0.14	1797
	0.2	False	7.21	11.40	5.24	0.13	61	1.00	0.14	15847
		True	7.68	4.01	3.04	0.19	86	1.00	0.14	1796
	0.3	False	7.20	11.38	5.23	0.13	61	1.00	0.14	15847
		True	7.78	4.13	3.14	0.19	86	1.00	0.14	1796
	0.5	False	7.21	11.38	5.23	0.13	61	1.00	0.14	15855
		True	7.67	4.13	3.12	0.18	86	1.00	0.14	1788
	1.0	False	7.21	11.38	5.23	0.13	61	1.00	0.14	15851
		True	7.67	4.10	3.14	0.17	86	1.00	0.14	1792
	0.1	False	7.12	11.68	5.26	0.12	59	0.995	0.19	15194
		True	8.14	4.22	3.54	0.25	89	0.995	0.19	2449
	0.2	False	7.11	11.65	5.25	0.12	60	0.998	0.19	15281
		True	8.22	4.14	3.51	0.25	89	0.998	0.19	2362
circle hex	0.3	False	7.12	11.59	5.24	0.12	60	0.998	0.18	15370
		True	8.23	4.25	3.51	0.25	89	0.998	0.18	2273
	0.5	False	7.13	11.49	5.22	0.12	60	0.998	0.16	15567
		True	8.22	4.28	3.47	0.24	88	0.998	0.16	2076
	1.0	False	7.16	11.23	5.15	0.13	61	0.995	0.12	16108
		True	8.32	4.51	3.59	0.21	87	0.995	0.12	1535
hull square	0.1	False	7.03	11.16	5.07	0.11	61	0.998	0.15	15625
		True	9.05	6.62	4.63	0.36	82	0.998	0.15	2018
	0.2	False	7.03	11.17	5.07	0.11	61	0.998	0.15	15617
		True	9.04	6.62	4.64	0.36	82	0.998	0.15	2026
	0.3	False	7.03	11.17	5.07	0.11	61	0.998	0.15	15620
		True	9.04	6.58	4.61	0.36	82	0.998	0.15	2023
	0.5	False	7.03	11.17	5.07	0.11	61	0.998	0.15	15605
		True	9.01	6.64	4.59	0.36	82	0.998	0.15	2038
	1.0	False	7.03	11.19	5.08	0.11	61	0.998	0.15	15565
		True	9.02	6.58	4.57	0.36	82	0.998	0.15	2078
hull hex	0.1	False	7.03	11.17	5.07	0.11	61	0.998	0.15	15623
		True	9.05	6.62	4.61	0.36	81	0.998	0.15	2020
	0.2	False	7.02	11.16	5.06	0.11	61	0.998	0.15	15617
		True	9.12	6.65	4.68	0.36	82	0.998	0.15	2026
	0.3	False	7.02	11.16	5.06	0.11	61	0.998	0.15	15614
		True	9.13	6.67	4.70	0.37	82	0.998	0.15	2029
	0.5	False	7.02	11.16	5.06	0.11	61	0.998	0.15	15638
		True	9.09	6.62	4.64	0.36	82	0.998	0.15	2005
	1.0	False	6.99	11.16	5.06	0.10	61	0.998	0.15	15564
		True	9.29	6.77	4.71	0.38	82	0.998	0.15	2079
hull		False	7.03	11.16	5.07	0.11	61	0.998	0.15	15623
		True	9.03	6.69	4.62	0.35	81	0.998	0.15	2020
Voronoi		False	6.32	16.23	5.05	−0.00	39	0.65	0.77	6557
		True	7.81	7.34	5.00	0.22	78	0.65	0.77	11086

average, i.e. Y/A is within 0.1yd, EPA differs by 0.01 and completion percentage only by 3%, suggesting that unsafe passes are the rule. This is consistent with the idea that a majority of targets are competently defended and therefore unsafe. Furthermore, safe passes being the exception is also reflected in the split of 1, 792 to 1, 797 safe passes




Fig. 10 Recall and specificity for different region estimation methods with $\ell = 0.1\text{yd}$ and Voronoi

out of 17, 643 total passes. Compared to overall average, the Y/A, EPA and completion percentage of safe passes are higher, by 0.5yd, 0.05 points and 21% respectively, marking them the better option when available. This is further supported by the Specificity, which is 1.00 so all interceptions occur on unsafe passes.

Another observation can be made regarding air yards. Figure 11 contrasts safe and not safe passes for different methods. The attempted safe passes are thrown for an average gain of 4.01 to 4.13 yards, that is much shorter than both the overall average pass, 10.64yd, and the unsafe passes of even longer 11.40yd. When taking completion into account, the average distance gained by a pass is only 0.96–1.0 yards less than the attempted distance, while it is 6.15 to 6.16 yards less for unsafe passes. Furthermore, the average of gained yards per attempt is 0.56 to 0.58 more for safe passes than for unsafe passes, suggesting a bigger opportunity for additional yards after the catch if the pass is considered safe. This is consistent with the concept of no defender being near the point of the catch, and therefore also not being able to quickly tackle the receiver to end the play. In total the safe passes are more productive on average while at the same time being rarer.

There are several key differences comparing square and hexagonal grids for the physics inspired approach. The difference between safe and unsafe passes is larger when using the hexagonal grid. In the case of circles mapped to the square grid, safe passes result on average 0.56yd – 0.58yd more per attempt while for the k-rings on the hexagonal grid the safe passes result in 1.02 – 1.16 higher Y/A.

The latter also identifies more passes as safe, up to 2, 449 or 652, which at the same time are also completed at a higher percentage of 652 more, and the 0.89%, and 0.85% should be 89% and 85% respectively. This is also reflected in the recall, with the square grid consistently showing a recall of 0.14, while the hexagon grid starts with higher recall of 0.19. Especially at finer resolutions, the EPA difference as well as the completion




Fig. 11 Comparison of pass outcomes for the different region estimation methods with $\ell = 0.1\text{yd}$ and Voronoi

percentage and Y/A show a higher average gain by passes classified as safe with the based on the hexagon grid rather than the square grid. This is true for both safe and unsafe passes, suggesting that the implementation of hexagon rings rather than mapping the circle directly to hexagons influences the results. As the latter are smaller at the edges of the region, this indicates that the full circle is more likely to overestimate the area each player can cover.

On the other hand, the circles mapped to the square grid provides consistent performance for different grid resolutions while using k-rings on the hexagonal grid results in fewer passes being identified as safe as the cells become larger. This is reflected in the recall dropping from 0.19 at $\ell = 0.1$ to 0.12 at $\ell = 1.0$ and the number of passes classified as safe decreasing from 2,449 to 1,535. This suggests that when mapping a continuous shape to a discrete grid, the discretization error has little effect on the results. However, when calculating the region directly using the hexagonal grid distances, the difference in cell size is important. This is corroborated when comparing the square and hexagonal grids to discretize the hulls in the data-driven approach. When comparing the square and hexagonal grid version of the data-driven approach, they are almost identical for $\ell < 0.5\text{yd}$ and differ only slightly for larger grid cells. Safe passes using hexagon cells are slightly more productive, resulting on average in 0.27 additional yards and 0.02 more points with $\ell = 1.0\text{yd}$.

The differences between safe and unsafe passes are more pronounced, when hull-based regions are used for classification compared to the physics inspired models. Safe passes provide an average gain of 2 yards more than non-safe passes and EPA of around 0.25 more compared to up to 1.1 yards and 0.13 EPA for the hexagon k-rings. Both the model with the highest EPA and Y/A are based on hulls. Both the attempted and completed air yards are also higher, $6.62 - 6.77$ aAY/A and $4.61 - 4.71$ cAY/A compared to $4.01 - 4.51$ aAY/A and $3.04 - 3.59$ cAY/A. At the same time, the completion percentage is 3%–8% lower while the specificity remains at near perfect 0.998. This suggests that considering the initial movement direction and velocity allows to better identify longer passes as safe.

Compared to Voronoi and circle-based region estimation, using the hulls as a basis for estimating the boundaries of the reachable regions for players and ball provides better results in terms of difference between the safe and unsafe passes overall, which is expected based on the fact that the hull is calculated including initial velocity and direction of movement. An exception is the completion percentage, where the circle computed on the hexagon grid provides the highest completion percentage and the lowest portion of open cells in the unsafe case.

Overall, the methods presented here identify passes as safe that are highly unlikely to lead to interceptions. Additionally, while the throwing distance tends to be below average, the passes result in above average yardage gains because safety in this context means that there is no defender close by and therefore the potential to gain more yards after the catch is higher.

Conclusion and further work

The study at hand shows that the combination of reachable regions of players in combination with passing range provides an informative and novel perspective leading to a better understanding of the interaction between defensive backs, receivers, and the football in passing plays. It allows to clearly identify the areas on the field that can be safely targeted by a pass. Those passes thrown to a safe target area are more productive and at the same time very unlikely to result in an interception. The presented method allows teams, fans, academics, and analysts to improve their understanding of spatial dynamics between receivers, defensive coverage and the ball by moving from individual player paths to interacting areas. While player behavior and tactical approaches change depending on the game context, the safe target areas as identified in this work remain the same across situations: An area is safe to target with a pass if there is a receiver able to arrive in time to catch the football, but defenders are out of range.

We compare a physics-inspired model based on simplified kinematics with a data-driven approach leveraging the amount of data collected over a whole NFL season. The former is to provide good results, especially in terms of recognizing passes that are more likely to be completed and very unlikely to be intercepted. The latter, data-driven, approach is superior at identifying a wider variety of passes as safe, which is expected due to the estimated reachable regions closer aligned with the observed player trajectories.

While the proposed methodology is well suited to understand which parts of the pitch are safe to target with a pass, these safe passes are not necessarily the best options available to the quarterback. Often times, a safe option is a quick checkdown to a Tight End

or Running Back. However, while that pass might be the safe option, it can be a better alternative to target a receiver further down the line that is open enough to justify the extra risk. Especially given game context, e.g. if time is tight or the play is a third down, a pass might not always be the *best* available pass just because it is *safe*. In further work, incorporating additional context such as down and distance or other game context can provide a path to understand and analyze QB decision-making better. Additionally, a more detailed kinematic model [16, 23, 38] can narrow down the reachable regions even more and give a more precise estimation and understanding.

There is a variety of work using sports tracking data to understand reachable areas based on player movement observations [3, 15–17, 23, 38], the concept of pitch control [3, 12, 15, 18, 19, 21, 22, 25] and passing analysis in American football [6, 7, 29, 39] and other sports, e.g. [8, 9, 13, 26–28]. However, using the travel time of the football to understand how far players can travel in the same time interval provides a novel perspective on spatial dynamics and passing. Taking into account the throwing range of the ball is especially relevant in American football, since the ball cannot touch the ground, or the pass will be ruled incomplete. This allows us to clearly identify the areas where a pass is possible and, more importantly, highly unlikely to be intercepted, adding to the existing passing analysis focused on completion [6], targeted player [40], or pass interference fouls [7].

The approach of reachable regions results in a binary view of whether a defender can interfere with a receiver catching the ball or whether a receiver can even reach the ball to catch it. This does not allow any reasoning about the quality of the pass or whether it was the best available option. Especially at the edges of the reachable regions, a probabilistic model of player movement as proposed by [18] or [17] might allow for a more nuanced analysis. Similar to the data-driven approach, for each player a two-dimensional discrete distribution can be estimated, which then allows for how likely a player will end up at a given grid cell. Those can then be aggregated following the same logic as presented in this work to create a continuous rather than a binary concept of safety.

Alternatively, the approach of [12], which uses the difference between the time of arrival of teammates and opponents to analyze passing in soccer, can also provide more information. They conclude that the pass is more likely to succeed if the difference in arrival times between teammate and opponent increases. They do, however, not consider time of arrival of the ball in their analysis. This can extend our analysis by quantifying *how* safe a target area is rather than deciding whether it is safe or not, e.g. by computing a value between zero and one instead of a binary classification of safety. A continuous quantification of safety can open new directions of research such as risk-reward-tradeoffs.

In order to focus on the concept of safe regions, we only consider the interval between the passing of the ball and when the outcome of the pass is decided. This approach constrains the analysis to one interval per play and does not allow recognition of missed passing opportunities or passes that are significantly longer or shorter in terms of time needed to reach the target. Therefore, another interesting avenue for research is an analysis of a whole play continuously. This requires a more complex approach to efficiently identify and update reachable regions but could provide valuable insights into defensive coverage or offensive structure.

While the presented methods are efficient enough to be executed in real time, continuous analysis of complete plays can lend itself well for a real-time decision-support tool in game contexts, adding a layer of understanding the current game, to coaches and analysts, can provide all the edge needed in close contests and can provide insights not only into actually thrown passes but also missed opportunities. In order to facilitate a continuous and real-time application, the next step would be investigating how much intermediate results can be saved and reused and how to compute the difference between two subsequent time steps instead of individually computing one time step after the other.

Given a continuous analysis, the results can be provided to the coaching staff using a dashboard around a 2D-Player visualizing player movement and game context, such as down and distance, score, lineups in combination with the results of the analysis for individual plays. Another interesting avenue is aggregating different plays within a game, e.g. in the form of a heatmap, to better understand patterns in the coverage or receiver movements by aligning the different plays using the line of scrimmage or the first down line. However, as different plays can vary a lot, depending on personnel and game situation, this probably needs to be combined with a filtering mechanism such as formations, personnel on the field, play type or outcome, and field position. A dashboard can also be used to extend the analysis with other metrics, such as individual player metrics like Quarterback Rating, or other team or play focused metrics like win probability added (WPA) or expected completion percentage to quantify the performance of teams, lineups, or players, and achieve more well rounded and comprehensive understanding of the game or situation at hand.

The competitions hosted by the NFL that provide the data for this work are designed with specific questions in mind. Although public availability is helpful for scientific progress, the published data set is limited. In the data for this work, only games of a single season are taken into account, the tracking data at 10 Hz is only available for a subset of players, and no external factors such as weather are taken into account. The high speeds lead to imprecise measurement for players and the ball, especially if they are moving fast, and for some aspects, such as when a pass has arrived, the annotations can be imprecise and may not be available in a real time decision support setting.

Extending the dataset can help to make the results more robust, especially with respect to rare situations or conditions. Adding missing players provides more context for understanding the passer. More importantly, adding data of more seasons allows provides more robust insights, and allows for a better estimate for the data-driven approaches. Lastly, a higher resolution of tracking, e.g. 30 Hz instead of 10 Hz could alleviate imprecision in measurements caused by high speeds as well as interactions between players and other players or the ball. Overall, a more complete and larger dataset could provide a more robust and more well rounded analysis, which is likely to be beneficial in game as complex as American football.

We propose a methodology for computing safe passing target regions and evaluating them by contrasting passes between safe and unsafe regions. The same regions can also be used for a variety of other analysis. On the one hand, the same analysis can be used to evaluate the passing behavior of teams or individual quarterbacks based on their tendency to play safe or unsafe passes. However, spatial analysis can also be used to analyze the capabilities of offensive lineups to generate safe target areas following the idea that the amount of open space is indicative of a good offensive scheme. Following the

same logic, our methodology can be used to quantify defensive performance, as a good defense prevents the development of safe target areas.

Acknowledgements

We acknowledge financial support by Universität der Bundeswehr München. Rudy Milani is funded by dtec.bw—Digitalization and Technology Research Center of the Bundeswehr project RISK.twin. dtec.bw is funded by the European Union-NextGenerationEU.

Author contributions

Conceptualization M.S., R.M., M.M.; Methodology, Investigation, Formal analysis, Software, Visualization M.S.; Writing - original draft M.S.; Writing - Review and Editing. M.S., R.M., M.M., S.P.; Funding acquisition. M.M. and S.P.; Supervision M.M. and S.P.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Data availability

All player tracking data used for this study is publicly available at <https://www.kaggle.com/competitions/nfl-big-data-bowl-2021/data>

Declarations

Competing interests

The authors declare no competing interests.

Received: 17 April 2025 / Accepted: 15 March 2026

Published online: 09 April 2026

References

1. Torres-Ronda L, Beanland E, Whitehead S, Sweeting A, Clubb J. Tracking systems in team sports: a narrative review of applications of the data and sport specific analysis. *Sports Med Open*. 2022;8(1):15. <https://doi.org/10.1186/s40798-022-00408-z>.
2. Kovalchik SA. Player tracking data in sports. *Annu Rev Stat Appl*. 2023;10(1):677–97. <https://doi.org/10.1146/annurev-statistics-033021-110117>.
3. Taki T, Hasegawa J. Visualization of dominant region in team games and its application to teamwork analysis. In: *Proceedings Computer Graphics International 2000/2000*, pp. 227–235. <https://doi.org/10.1109/CGI.2000.852338>.
4. Howard A, Reid JE, Lopez M, Cukierski W. Big Data Bowl. <https://operations.nfl.com/gameday/analytics/big-data-bowl/>. Accessed: 2025-4-4; 2018.
5. Pro-Football-Reference.com: 2018 NFL Standings & Team Stats. <https://www.pro-football-reference.com/years/2018/index.htm>. Accessed: 2025-11-20; 2025.
6. Deshpande SK, Evans K. Expected hypothetical completion probability. *J Quant Anal Sports*. 2020;16(2):85–94. <https://doi.org/10.1515/jqas-2019-0050>.
7. Skoki A, Lerga J, Stajduhar I. ML-based approach for nfl defensive pass interference prediction using gps tracking data. In: *2021 44th International Convention on Information, Communication and Electronic Technology (MIPRO)*; 2021. pp. 1038–1043. <https://doi.org/10.23919/MIPRO52101.2021.9596877>.
8. Li H, Zhang Z. Predicting the receivers of football passes. In: *Brefeld U, Davis J, Van Haaren J, Zimmermann A, editors. Machine Learning and Data Mining for Sports Analytics*. Cham: Springer; 2019. p. 167–77.
9. Anzer G, Bauer P. Expected passes: determining the difficulty of a pass in football (soccer) using spatio-temporal data. *Data Min Knowl Discov*. 2022;36(1):295–317. <https://doi.org/10.1007/s10618-021-00810-3>.
10. Yurko R, Ventura S, Horowitz M. nflWAR: a reproducible method for offensive player evaluation in football. *J Quant Anal Sports*. 2019;15(3):163–83. <https://doi.org/10.1515/jqas-2018-0010>.
11. Reyers M, Swartz TB. Quarterback evaluation in the national football league using tracking data. *ASTA Advances in Statistical Analysis*. 2021;107(1):327–42. <https://doi.org/10.1007/s10182-021-00406-8>.
12. Narizuka T, Yamazaki Y, Takizawa K. Space evaluation in football games via field weighting based on tracking data. *Sci Rep*. 2021;11(1):1–8. <https://doi.org/10.1038/s41598-021-84939-7>.
13. Spearman W, Basye A, Dick G, Hotovy R, Pop P. Physics-based modeling of pass probabilities in soccer. In: *MIT Sloan Sports Analytics Conference*; 2017. 2017.
14. Cao S. Passing path predicts shooting outcome in football. *Sci Rep*. 2024;14(1):9572. <https://doi.org/10.1038/s41598-024-60183-7>.
15. Fujimura A, Sugihara K. Geometric analysis and quantitative evaluation of sport teamwork. *Syst Comput Jpn*. 2005;36(6):49–58. <https://doi.org/10.1002/scj.20254>.
16. Renkin M, Bischofberger J, Schikuta E, Baca A. Validation and optimisation of player motion models in football. In: *Computational Science – ICCS 2022. Lecture notes in computer science*; 2022. pp. 26–32. Springer, Cham. https://doi.org/10.1007/978-3-031-08754-7_4.
17. Brefeld U, Lasek J, Mair S. Probabilistic movement models and zones of control. *Mach Learn*. 2019;108(1):127–47. <https://doi.org/10.1007/s10994-018-5725-1>.
18. Fernández J, Bornn L, Fraser S. Wide open spaces: a statistical technique for measuring space creation in professional soccer. In: *MIT Sloan Sports Analytics Conference*; 2018. 2018.
19. Bornn L, Goldsberry K. NBA court reality. In: *MIT Sloan Sports Analytics Conference*; 2016. 2016.

20. Higgins L, Galla T, Prestidge B, Wyatt T. Measuring the pitch control of professional football players using spatiotemporal tracking data. *J Phys Complex*. 2023;4(2):025008. <https://doi.org/10.1088/2632-072x/acb67d>.
21. Raabe D, Biermann H, Bassek M, Memmert D, Rein R. The dual problem of space: relative player positioning determines attacking success in elite men's football. *J Sports Sci*. 2024;42(19):1821–30. <https://doi.org/10.1080/02640414.2024.2414363>.
22. Franks A, Miller A, Bornn L, Goldsberry K. Characterizing the spatial structure of defensive skill in professional basketball. *The Annals of Applied Statistics*. 2015;9(1):94–121. <https://doi.org/10.1214/14-AOAS799>.
23. Narizuka T, Takizawa K, Yamazaki Y. Validation of a motion model for soccer players' sprint by means of tracking data. *Sci Rep*. 2023;13(1):865. <https://doi.org/10.1038/s41598-023-27999-1>.
24. Caetano FG, Barbon Junior S, Torres RDS, Cunha SA, Ruffino PRC, Martins LEB, et al. Football player dominant region determined by a novel model based on instantaneous kinematics variables. *Sci Rep*. 2021;11(1):18209. <https://doi.org/10.1038/s41598-021-97537-4>.
25. Burris K. A Trajectory Planning Algorithm for Quantifying Space Ownership in Professional Football. *operations.nfl.com*; 2019.
26. Power P, Ruiz H, Wei X, Lucey P. Not all passes are created equal: objectively measuring the risk and reward of passes in soccer from tracking data. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, New York, NY, USA; 2017. <https://doi.org/10.1145/3097983.3098051>.
27. Goes FR, Kempe M, Meerhoff LA, Lemmink KAPM. Not every pass can be an assist: a data-driven model to measure pass effectiveness in professional soccer matches. *Big Data*. 2019;7(1):57–70. <https://doi.org/10.1089/big.2018.0067>.
28. Kempe M, Goes FR, Lemmink KAPM. Smart data scouting in professional soccer: evaluating passing performance based on position tracking data. In: *2018 IEEE 14th International Conference on e-Science (e-Science)*; 2018. pp. 409–410. <https://doi.org/10.1109/eScience.2018.00126>.
29. Yurko R, Matano F, Richardson LF, Graner N, Pospisil T, Pelechris K, et al. Going deep: models for continuous-time within-play valuation of game outcomes in American football with tracking data. *J Quant Anal Sports*. 2020;16(2):163–82. <https://doi.org/10.1515/jqas-2019-0056>.
30. Eager E, Seth T. Investigating trade-offs made by American football linebackers using tracking data. *J Quant Anal Sports*. 2023;19(3):171–85. <https://doi.org/10.1515/jqas-2022-0091>.
31. Aggarwal CC. *Data Mining*. 2015 edn. Cham, Switzerland: Springer; 2015. <https://doi.org/10.1007/978-3-319-14142-8>.
32. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat Methods*. 2020;17:261–72. <https://doi.org/10.1038/s41592-019-0686-2>.
33. Bresenham J. A linear algorithm for incremental digital display of circular arcs. *Commun ACM*. 1977;20(2):100–6. <https://doi.org/10.1145/359423.359432>.
34. Bresenham JE. Bresenham: algorithms computer control digital plotter. *IBM System Journal: Algorithms Computer Control Digital Plotter*; 1965.
35. Birch CPD, Oom SP, Beecham JA. Rectangular and hexagonal grids used for observation, experiment and simulation in ecology. *Ecol Modell*. 2007;206(3–4):347–59. <https://doi.org/10.1016/j.ecolmodel.2007.03.041>.
36. Howard A, Reid JE, Lopez M, Bliss T, Cukierski W. *NFL Big Data Bowl 2021*. <https://kaggle.com/nfl-big-data-bowl-2021>. Accessed: 2025-4-4; 2020.
37. Pelechris K, Winston W, Sagarin J, Cabot V. Evaluating NFL plays: expected points adjusted for schedule. In: *Machine Learning and Data Mining for Sports Analytics. Lecture notes in computer science*; 2019. pp. 106–117. Springer, Cham. https://doi.org/10.1007/978-3-030-17274-9_9.
38. Grehaigne J-F, Bouthier D, David B. Dynamic-system analysis of opponent relationships in collective actions in soccer. *J Sports Sci*. 1997;15(2):137–49. <https://doi.org/10.1080/026404197367416>.
39. Dutta R, Yurko R, Ventura SL. Unsupervised methods for identifying pass coverage among defensive backs with NFL player tracking data. *J Quant Anal Sports*. 2020;16(2):143–61. <https://doi.org/10.1515/jqas-2020-0017>.
40. Dick U, Link D, Brefeld U. Who can receive the pass? A computational model for quantifying availability in soccer. *Data Min Knowl Discov*. 2022;36(3):987–1014. <https://doi.org/10.1007/s10618-022-00827-2>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.