Check for updates

Expected passes

Determining the difficulty of a pass in football (soccer) using spatio-temporal data

Gabriel Anzer[1,2] · **Pascal Bauer**[2,3]

Abstract

Passes are by far football's (soccer) most frequent event, yet surprisingly little meaningful research has been devoted to quantify them. With the increase in availability of so-called positional data, describing the positioning of players and ball at every moment of the game, our work aims to determine the difficulty of every pass by calculating its success probability based on its surrounding circumstances. As most experts will agree, not all passes are of equal difficulty, however, most traditional metrics count them as such. With our work we can quantify how well players can execute passes, assess their risk profile, and even compute completion probabilities for hypothetical passes by combining physical and machine learning models. Our model uses the first 0.4 seconds of a ball trajectory and the movement vectors of all players to predict the intended target of a pass with an accuracy of 93.0% for successful and 72.0% for unsuccessful passes much higher than any previously published work. Our extreme gradient boosting model can then quantify the likelihood of a successful pass completion towards the identified target with an area under the curve (AUC) of 93.4%. Finally, we discuss several potential applications, like player scouting or evaluating pass decisions.

✉ Gabriel Anzer
gabriel.anzer@sportec-solutions.de

Pascal Bauer
pascal.bauer@dfb.de

1 Sportec Solutions AG, subsidiary of the Deutsche Fußball Liga (DFL), Munich, Germany

2 Institute of Sports Science, Department of Sport Psychology and Research Methods, University of Tübingen, Tubingen, Germany

3 DFB-Akademie, Deutscher Fußball-Bund e.V. (DFB), Frankfurt, Germany

1 Introduction

Passes are a crucial part of modern football (soccer) matches. However, traditionally player's passing performance is quantified using a binary pass completion metric. This means that — regardless of the quality or difficulty of a pass — completed passes are rewarded a "1/+", and incomplete passes rewarded a "0/−". A player's pass completion rate is thus calculated as the ratio of completed passes to total passes played. This ratio neglects to take the complexity or the reward of a pass into consideration. Nevertheless, pass completion rates are regularly used as performance indicators on team and player levels in literature (Bradley et al. 2013; Król et al. 2017) and in the daily business of professional football teams. Whenever a player is in possession of the ball, they may choose to pass to any of their teammates — and each option comes with a unique set of risks and rewards. This decision can only be evaluated by considering both the risk and the reward of each option.

The relevance of passes in football was investigated with annotational analysis of passing patterns (Reep and Benjamin 1968) and through experimental studies analyzing influencing factors for passes (Williams 2000). The increasing availability of granular football data unlocked new avenues for the analysis of passes. Event data, following the idea of Reep and Benjamin (1968), describes a log of all on-the-ball-actions (e.g. shots, passes, tackles) and are systematically acquired in most professional football leagues. Several studies used this event data to analyze passes on a much larger scale than previous experimental studies would allow. For example, Szczepański et al. (2016) used 253, 090 open-play passes and McHale and Relton (2018) analyzed 960, 000 events including passes. While manually collected event data provides relevant information about one or two players involved in the current ball action, recent improvements in computer vision allow to accurately track the positions of all 22 players and the ball at any time of the match. This type of data is typically referred to as tracking, positional or movement data (Stein et al. 2017; Andrienko et al. 2019; Bauer and Anzer 2021; Anzer et al. 2021).

While some studies quantified the reward of a pass using event data only (Brooks et al. 2016; Power et al. 2017; Bransen et al. 2019), combining the manually tagged event data with the automatically acquired positional data allows for a more granular analysis of the reward of a pass. Several studies addressed this reward-quantification of passes in different ways (Rein et al. 2017; Chawla et al. 2017; Goes et al. 2019; Gómez-Jordana et al. 2019; Steiner et al. 2019; Anzer and Bauer 2021), but they typically measure how much a pass would increase the chance of scoring if successful.

As highlighted in Power et al. (2017) and Goes et al. (2021) the quantification of pass decisions has two dimensions: The reward of a pass, as discussed above, and the difficulty of the pass, usually measured in the completion probability. This risk of a pass is often referred to as expected pass (xPass) values in literature (Spearman et al. 2017; Power et al. 2017; Fernández et al. 2020; Arbués-Sangüesa et al. 2020; Alguacil et al. 2020; Stöckl et al. 2021). An xPass model tries to estimate the probability of

a given pass being successfully completed to a teammate, based on various factors describing the pass — usually derived from positional and/or event data. Furthermore, Li et al. (2019) and Vercruyssen et al. (2016) have explored the target identification of passes as a standalone problem. To quantify the risk, Power et al. (2017) built a logistic regressor for 571, 278 passes, whereas Spearman et al. (2017) modelled 10, 875 passes as Bernoulli trials. In order to retrieve the missing information regarding the intended receiver of a pass (at the moment the pass was played), they modelled both ball and player trajectory based on physical simulations first. This allows them to calculate an xPass value, as the predicted probability of a pass being completed, at the moment when the pass is played. The physics-based models were slightly improved by Alguacil et al. (2020) through taking friction for ground-passes into consideration.

Stöckl et al. (2021) later slightly improved the accuracy of the xPass model by using Graph Neural Networks (Battaglia et al. 2018) to overcome both the feature extraction and the ordering-problem of using spatio-temporal tracking data in a dynamic sport like football. Arbués-Sangüesa et al. (2020) showed that a player's body orientation (typically not included in off-the-shelf tracking data) has a significant influence on pass completion probabilities as well. Several further extensions, built on top of xPass models, exist in the literature: Fernandez et al. (2018) and Spearman et al. (2017) include xPass models as central ingredients for computing their expected possession values, and Hubáček et al. (2018) use it to try to predict which pass will be played next in any given situation.

But overall the literature is lacking a thoroughly described method of synchronizing pass events with tracking data, a highly accurate intended receiver estimation and a properly (manually) evaluated xPass model. Our work fills this gap, while keeping the individual modules completely separated and introduces novel concepts, like blocking probabilities.

Our goal is to train a machine learning model on the binary classification, of whether a pass will be successful or not using all the information available at the time of the pass. While the data set (described in Sect. 2) is extremely detailed, it is missing one piece of essential information, namely the targeted recipient of unsuccessful passes. Our work consists of the following four steps:

(1) **Synchronization of pass events:** We synchronize both the location and the exact timing of pass events from manually annotated event data with automatically acquired tracking data (similar to the method introduced for shot events in Anzer and Bauer (2021)). Details of the approach can be found in the Appendix A.
(2) **Estimate the intended receiver:** First, we use a state-of-the-art movement model to derive the potential positions of all players within a certain time window according Brefeld et al. (2019) (see Sect. 3.2), and second, combine this with a physics-based ballistic ball trajectory model as described in Spearman et al. (2017) (see Sect. 3.1). Given the ball positions within the first 0.4 seconds of a pass, this model uses the results from aerodynamic investigations (Asai et al. 2007; Oggiano and Satran 2010) to predict the trajectory of the ball. The combination of both steps provides us an accurate prediction of the intended receiver for unsuccessful passes (see Sect. 3.3).

(3) **Pass probability:** In Sect. 4, we train a machine learning model to estimate the probability of a pass (that was not blocked immediately) based on the information derived from (2) and from expert-based features describing the pass.

(4) **Blocking model:** In order to get unbiased estimates for the probabilities of all passes, we further calculate the likelihood that a pass is blocked (see Sect. 5). This is also approached using a supervised machine learning model with hand-crafted features.

Finally, we can compute the probability of any potential pass being completed. By combining and slightly improving previor work, we exceed the accuracy of all previously presented results for the prediction of the pass receiver as well as the classification of played passes being successful or not.

2 Data and definitions

In the official match-data catalog of the German Bundesliga,[1] a pass is defined the attempt to switch ball control from one player to a teammate. For each pass detected, trained operators annotate a variety of sub-attributes describing the pass in detail. Among others, they annotate who played and (in case of a successful pass) received the ball, whether it was a *high* or a *low played* pass, as well as, whether the pass was played over a *short*, *medium*, or *long* distance. Of course, all of the sub-attributes underlay detailed definitions, defining high passes as passes played above knee height and setting thresholds to differentiate short passes (< 10 m), passes of medium length ($10 - 30$ m) and long passes (> 30 m). All attributes are collected for both successful and intercepted passes, meaning that the intended height and the intended length is estimated by the human operator in case of intercepted passes. While this manually acquired event data underlays strict quality checks, especially for incomplete passes it can be quite subjective.

More objective and more granular information can be found in the positional data, capturing the positions of all 22 players and the ball at 25 Hz. In each Bundesliga-stadium, up to 20 installed HD-cameras record any action on the pitch and serve as input for computer vision algorithms estimating the 2D-positions of all players as well as the 3D positions of the ball. In the Bundesliga, data from Chyronhego's optical Tracab system is collected.[2] Several studies evaluated the accuracy of this data (Linke et al. 2020, 2018).

We excluded fair-play passes, in which a player voluntarily relinquishes his team's ball-control, passes that accidentally end up with a teammate who was not the intended target, as well as throw-ins from our analysis. This information is captured within the event data and can thus be simply filtered out for our investigation. We end up with positional and event data of 840, 386 passes from 918 Bundesliga games from the 2017/2018, 2018/2019 and 2019/2020 seasons, with an average completion rate of 85.2%.

---

[1] https://s.bundesliga.com/assets/doc/10000/2189_original.pdf (accessed March 27, 2021).

[2] https://tracab.com/products/tracab-technologies/tracab-optical/ (accessed April 14, 2021).

The necessity to synchronize the two independently acquired data-sources, is detailed in the literature (Anzer and Bauer 2021; Spearman et al. 2017). We apply a slightly modified methodology to synchronize pass events as described for shots in Anzer and Bauer (2021). The outcome of the synchronization is manually evaluated in Sect. 6 as a part of the xPass evaluation, finding that 99.1% of all pass events are identified correctly. Further details on the synchronization methodology as well as a more thorough validation study are provided in the Appendix A.

For reproduction, Pettersen et al. (2014) present a publicly available set of positional data, and open source event data can be found in Pappalardo et al. (2019).[3]

3 Estimating the target

While the receivers of successful passes are included in the event data, the intended target of unsuccessful passes is missing. This is a crucial point of information necessary for determining the difficulty of a pass, since otherwise only the surrounding circumstances of the passer could be taken into consideration. Therefore, we need to determine who the intended receiver was, to later extract features for both successful and unsuccessful passes. For that purpose we first use a physics-based approach to estimate the ball trajectory based on the first couple of frames after the pass is played (Sect. 3.1). Second, we compute a movement model, to estimate the area on the pitch, players could potentially reach in the next n frames, based on their movement direction and velocity (Sect. 3.2). Third, by combining both the estimated ball trajectory and the reachable area, we identify the teammate most likely to reach the ball first as the intended recipient of the pass (Sect. 3.3). Furthermore, we discuss (Sect. 3.4) how this can be used to derive physics-based features describing the difficulty of a pass.

3.1 Modelling the ball trajectory

Knowing that a football adheres to physical laws, we can use these laws to determine the path a ball will travel on (until it is touched again) based on its initial direction and velocity. As suggested in Spearman et al. (2017), we use the first 10 frames (equivalent to 0.4 seconds) after a pass was played, to receive a stable estimate of its initial direction (x, y, z) and exit velocity. Therefore, we exclude all passes blocked within the first 0.4 seconds, since we are unable to determine the necessary starting values for them reliably. Using a physical trajectory model, including gravity, air drag and rolling drag (with the simplification that as soon as a ball lands, it is grounded), we can estimate for every following frame, where the ball will be. As presented in Spearman et al. (2017), the trajectory of the ball is consequently modelled as:

$$\ddot{r} = -g\hat{z} - \frac{1}{2m}\rho C_D A \dot{r}\dot{r}$$

[3] Other (non-scientific) open-source data sets can be accessed from Skillcorner (https://github.com/SkillCorner/opendata), Metrica sports (https://github.com/metrica-sports/sample-data) or Statsbomb (https://github.com/statsbomb/open-data).
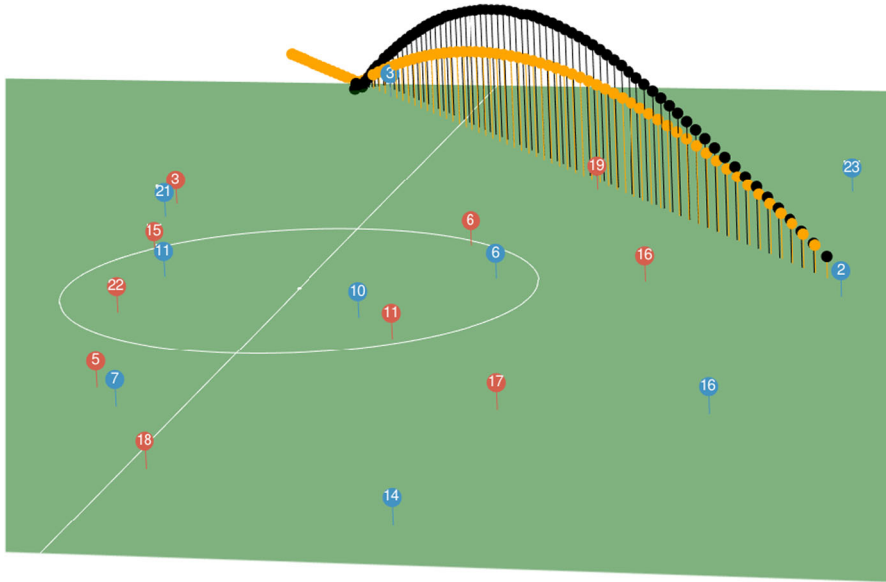
Fig. 1 Estimated ball trajectory (yellow dots) compared to the measured data-points from the tracking data. The video footage of the pass can be found here: https://dfb-my.sharepoint.com/:v:/g/personal/pascal_bauer_dfb_de/EUJra9f8i6BCl2-mzpuHtacBvgHNx7cnCH9P8y5taozDnQ?e=nIhgkC (Color figure online)

All physical values are set to the respective standard.[4] The Bundesliga-ball has a weight of 0.4 kg (m) and a cross-sectional area of 0.038 m^2 Further background information regarding the aerodynamics of balls in football can be found in Asai et al. (2007); Oggiano and Satran (2010).

Figure 1 shows both the observed ball path from the tracking data (black dots) and the estimated ball trajectory from our physics-based model (yellow dots) for a played pass. The physics-based model yields a smooth and realistic ball path, while the observed ball path shows some jumps (e.g. around the highest point of the trajectory), frequently present when tracking small fast moving objects from large distances. Furthermore, it can be used to model where a ball might have ended up, had it not been deflected or intercepted.

3.2 Movement model

The movement model predicts what area of the pitch a player can reach within a defined time-window. Again, Spearman et al. (2017) presented a physics-based approach for this. Additionally, they gave a first outlook towards data-driven movement models which were later built upon by several studies. The positions a player can reach within a certain time frame depend on his current speed and direction (Brefeld et al. 2019; Fernandez et al. 2018). With these assumptions we use movement data from three

---

[4] Gravitational force $g = 9.8 \frac{m}{s^2}$; air density $\rho = 1.22 \frac{kg}{m^3}$; drag coefficient $C_D = 0.25$.

seasons of Bundesliga data. First, we transform the data, so that all players are traveling in the same direction. Next we compute the convex hull of all observed locations players traveling in a certain speed interval were able to reach after n-frames. Due to our large data set of tracking data, we are able to use much smaller speed intervals (of 0.5 km/h) compared to Brefeld et al. (2019). With this information we fit our movement model to estimate the center of the circle and its diameter, based on speed and time.

Now we can calculate for any player on the pitch what area they could theoretically cover in the next seconds based on their movement vector. This is displayed for some players (#18/#22 red team; #7 blue team) in Fig. 3. Each circle represents the area the respective player can reach within 0.5, 1, 1.5 and 2 seconds.

3.3 Target estimation

To estimate the intended target of a given pass, we combine the physics-based ball trajectory model with the data-driven player movement model. To incorporate the ball height, we additionally assume that a ball is only reachable below a height of 1.5 m. This threshold was obtained by optimizing for the accuracy of the intended receiver prediction for successful passes on the training data set introduced in Sect. 4. Thus, we can calculate which team member of the passer could theoretically be the first to reach the pass, and declare them as the intended receiver.

We are able to predict the correct player for successfully completed passes with an accuracy of 93.1%. For unsuccessful passes, we conducted an evaluation study, described in Sect. 6, showing that we are able to predict the estimated target with an accuracy of 72.0%.

3.4 Physics-based passing features

We can quantify what direction and how fast a pass would need to be played to arrive at the target receiver. For that purpose we compute hypothetical passes by varying the initial starting parameters of a pass, i.e. initial velocity and initial direction of a pass. Combined with the movement model, we can determine if the target player is still the most likely player to receive each hypothetical pass. This step is done by performing a grid based search varying the velocity and the direction of the pass noting for every (reasonable) combination[5] if the intended target is likely the first player to potentially reach the pass. From this we can compute the direction window, defined as the width of the reachable angles. The speed window is defined as the difference between the maximal relative increase and decrease of a baseline exit velocity, with which hypothetical passes would still be reachable by the intended receiver. An example of a direction window is indicated in Fig. 2. Hypothetical passes, with slightly modified x-/y-directions (assuming an ideal speed and launch-angle) that could be received by Emre Can (#3 of the blue team) according to our model are displayed as grey lines. The total width of theses potential pass angles amounts to 21 degrees in this example.

---

[5] The velocity range is between $[-100\%, 100\%]$ of the average pass speed, and the direction window range is between $[-25°, 25°]$ of the direct connection line between the passer and the intended target.
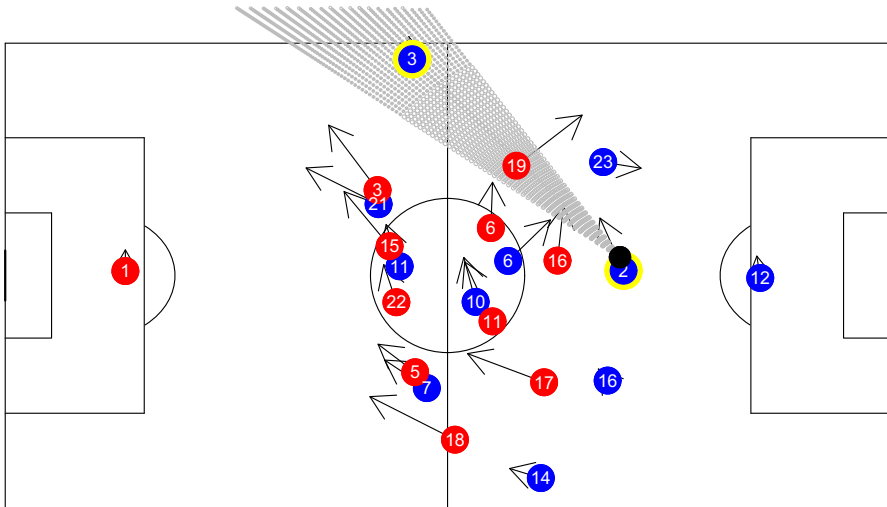
Fig. 2 Visualization of potential pass angles reaching the intended target. Players' movement vectors are displayed as arrows. The passing player (Robin Koch, #2 of the blue team) as well as the receiving player (Emre Can, #3 of the blue team) are highlighted in yellow. The same pass is displayed as in Fig. 1 and in this video: https://dfb-my.sharepoint.com/:v:/g/personal/pascal_bauer_dfb_de/EUJra9f8i6BCl2-mzpuHtacBvgHNx7cnCH9P8y5taozDnQ?e=nIhgkC (Color figure online)

Table 1 Speed scalar window width (as percentage) and direction window (in degrees) of passes

Passe Outcome	Speed window	Direction window	Number of Passes
Successful	0.99 (±0.48)	32.7 (±16.3)	291, 700
Unsuccessful	0.26 (±0.36)	11.6 (±12.6)	56, 570

The observed standard deviation is denoted in parentheses

Table 1 shows that unsuccessful passes have both a much narrower window of potential directions (in degrees) as well as in speed values (in percentage difference compared to a baseline speed value). This aligns with expert opinions that the less accurate a pass needs to be played, the easier it is and the higher chance that it will be completed.

The interplay of the target prediction using the physics-based ball trajectory model and the data-driven movement model is displayed in Fig. 3. In this situation Robin Koch (#2 of the blue team) plays a diagonal ball to his teammate Julian Draxler (#7 of the blue team). The curvature of the ball trajectory (yellow dots) shows the trajectory of the played diagonal pass from the mid-point (player #2) to the left attacker of the blue team (player #7).[6] Due to its height, the ball can only be reached towards the end of the projected trajectory. Matching possible intersections of the trajectory after n frames with each teammate's reachable area after the same time period, reveals, that the first player to possibly reach the ball is the attacker on the left wing after 2.48

---

[6] The video of the pass can be found here: https://dfb-my.sharepoint.com/:v:/g/personal/pascal_bauer_dfb_de/EWIWkaF8Gp5CjPCdQRs5KXsB6Rt0LKHoKomXUFogNsR2Wg?e=u6EziX..
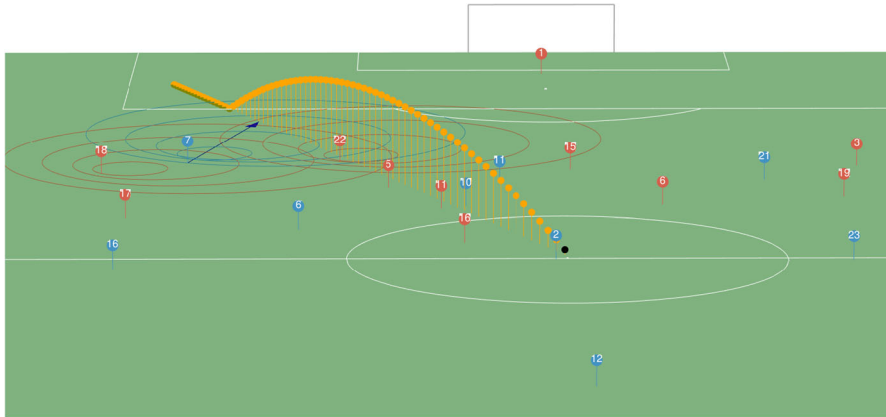
Fig. 3 Estimated target of a pass with ball-trajectory and movement models. The combination of the estimated ball trajectory (yellow dots) and the player movement model (blue and red circles) predict Julian Draxler (#7 of the blue team) reaching the ball first, with the arrow indicating the first point where he could potentially intercept the pass. The respective video sequence can be viewed here: https://dfb-my.sharepoint.com/:v:/g/personal/pascal_bauer_dfb_de/EWIWkaF8Gp5CjPCdQRs5KXsB6Rt0LKHoKomXUFogNsR2Wg?e=u6EziX (Color figure online)

seconds. The point where he could first reach the ball is indicated by the black arrow, but this does not mean that it is always the optimal strategy to do so.

4 Pass probability estimation

For successful passes we know the recipient of a pass from the event data. With the approach described in Sect. 3, we can identify the intended target of unsuccessful passes (as long as it is not blocked). This allows us to compute tailored features influencing the pass difficulty and train supervised machine learning models estimating pass completion probabilities. We build on the features describing passes presented recently (Power et al. 2017; Spearman et al. 2017; Mchale and Lukasz 2014; Hubáček et al. 2018). Table 2 shows an overview of all features we compute for every pass. As in Spearman et al. (2017), we are interested to train a predictive model, i.e. all features must be available at the time of a pass. This will allow us later to compute hypothetical pass probabilities, or evaluate if a player is over-/under-performing. The first eight features describe the pass origin and the situation around the passer, e.g. where on the pitch the passer is located, how far the next opponent is away from them and how much pressure they are receiving according the pressure-model introduced in Andrienko et al. (2017). The next set of features (rows $9 - 14$, Table 2) describe the receiver and their surrounding environment. The third block of features (rows $15 - 21$) describe further context information around the pass itself taking full advantage of the positional data. The manual collected features *height* and *distance* (rows 22 and 23) are described in Sect. 2. The last two features in Table 2 (rows 24 and 25) are calculated based on the logic described in Sect. 3.

Table 2 Hand-crafted features used to train the xPass-model

	Feature	Description	B
1	Location pass	x- and y- coordinate of the pass.	x
2	Distance sideline	Distance between the location of the pass and the closest sideline.	x
3	Distance goal	Distance between the location of the pass and the opposing goal.	x
4	Distance opponent	Distance between the location of the pass and the nearest opposing player at the time of the pass.	x
5	Speed passer	Speed of the passing player at the time of the pass.	x
6	Ball height	Ball height at the time of the pass.	
7	Opponents closer to goal	Number of opponents closer to their own goal than the passer at the time of the pass.	x
8	Pressure on passer	Pressure exerted on the passer according Andrienko et al. (2017).	
9	Location receiver	x- and y-coordinate of pass receiver at the time of the pass.	
10	Distance receiver sideline	Distance between the receiving player and the closest sideline at the time of the pass.	
11	Distance receiver goal	Distance between the receiving player and the opposing goal at the time of the pass.	
12	Distance receiver opponent	Distance between the receiving player and the nearest opposing player at the time of the pass.	
13	Opponents closer to goal receiver	Number of opponents closer to their own goal than the receiver at the time of the pass.	
14	Speed receiver	Speed of the receiving player at the time of the pass.	
15	Possession phase	The time the passer was in ball possession before attempting the pass.	x
16	Bypassed opponents	Opposing players that would be bypassed with the pass (Steiner et al. 2019).	
17	Angle	Directional angle of the pass compared to the playing direction (i.e. 0 is directly towards the opposing goal line, and 180 would be backwards towards the own goal line).	
18	Opponents in path	Number of opposing players in the passing path. The path being defined as a corridor between the pass location and the receiver location with a width of 10 meters.	
19	Nearest defender pass line	Nearest defender to a straight line connecting the pass location and the receiver location.	
20	Distance pass	Distance between passer and receiver when the pass is played.	
21	Dead ball	Binary information whether the pass originated from a set-piece (e.g. freekick, goalkick, ...) or not.	
22	Height	Manually annotated binary feature describing the intended ball height included in the event data.	x

Table 2 continued

	Feature	Description	B
23	Distance	Manually annotated intended pass length (short, medium, long) included in the event data.	x
24	Speed window	Mean speed window as defined in Sect. 3.4.	
25	Direction window	Mean direction window as defined in Sect. 3.4.	

Column "B" notes all features that are also used for the blocking model described in Sect. 5

For our model training we use 840, 386 passes from 918 Bundesliga games. We split the data into training (504, 232 passes), validation (168, 077 passes) and test set (168, 077 passes) and use different subsets of features from Table 2 to train various supervised machine learning models (logistic regression, extreme gradient boosting, random forest). For each feature set the best performing models on the test set were extreme gradient boosting (XGBoost) models (Chen et al. 2016). For all XGBoost models we applied Bayesian hyperparameter optimization on the validation set (Nazareth 2004). The accuracy metrics of the XGBoost models for the different feature sets are displayed in Table 3. Since precision, recall and F_1-score are not ideally suited to evaluate probabilities in an imbalanced data set, we focus on the metric area under the receiving operator curve (AUC), mean square error (MSE) and the Brier skill score (BSS, Brier (1950)). All relevant metrics indicate that the model using the full set of features (line 1, Table 3) provides the best results.

We implemented three simple baseline models for comparison: First, we trained a model using only the features that can be derived from event data (row 4, Table 3). In order to evaluate the necessity of identifying the intended receiver, we trained a receiver-agnostic model (row 5, Table 3). Lastly, row 6 in Table 3 presents a trivial baseline model, that assumes a constant pass success probability of 85.2% (the completion rate in the training data set). Table 3 shows, that both using positional data and the estimation of the targeted receiver (see Sect. 3) significantly improve the prediction accuracy.

The final hyper-parameter configuration for the model using all features is provided in the Appendix B (Table 8). In the complete model, according to the overall SHAP-values (Lundberg et al. 2017), the possible speed window in which the pass could be completed has by far the highest influence on the success prediction followed by the distance of the closest opponent to the receiver. Purely absolute position related features like the x-/y-coordinates of the pass origin and the receiver position, as well as the distance to the sideline/goal exhibited the lowest influence on the prediction. More details regarding the feature importance and SHAP-values can be found in Appendix B.

5 Blocking model

In order to get a reasonably reliable target identification in the previous sections we focused on passes that were not blocked immediately. However, this inflates the likeli-

Table 3 Outcome of the XGBoost models on the test set using different feature sets

	Model	Precision	Recall	F_1-score	AUC	MSE	BSS
1	All features included	0.944	0.958	0.951	0.955	0.059	0.528
2	Speed and direction window only	0.887	0.963	0.923	0.873	0.091	0.270
3	Speed and direction window excluded	0.923	0.953	0.937	0.920	0.075	0.397
4	Event-data only	0.894	0.960	0.926	0.832	0.095	0.241
5	Receiver Agnostic	0.898	0.957	0.927	0.864	0.091	0.027
6	Average pass probability	–	1	–	0.5	0.126	0

Table 4 Accuracy metrics for the blocking model on the test set

	Model	Precision	Recall	F_1-score	AUC	MSE	BSS
1	Gradient Boosting	0.643	0.031	0.060	0.753	0.059	0.084
2	Naive Blocking Model	–	0.000	–	0.500	0.065	0.000

hood of a pass being completed, since it ignores that about 3.12% of the passes in our data set are blocked. Therefore, we need to adjust our conditional passing probabilities by the likelihood that it is not blocked: For all successful passes A and all blocked passes B, the probability of any non-blocked pass being completed, $P(A \cap \bar{B})$, can be computed as follows:

$$P(A \cap \bar{B}) = P(A|\bar{B}) * P(\bar{B})$$

The probability of a pass success — provided that it is not blocked — $P(A \cap \bar{B})$, is calculated in Sect. 4. However, at the time of the pass we do not know whether it will be blocked or not. Consequently, to get an unbiased pass completion probability, we need to calculate the probability that it is blocked, $P(B)$. Rather then simply discounting all passes with the average rate in which passes passes are blocked (3.12%), we determine the likelihood of each pass being blocked individually, based on some of the features described earlier. We define a blocked pass, as a pass where an opposing player touches it within the first 0.4 seconds. A problem is that the exact initial direction of blocked passes cannot be accurately derived from tracking data. Therefore, we simply assume that if a pass was blocked the intended direction was towards the point where the opponent touched it.

Consequently, a pass can only be blocked (according to our definition) if an opposing player is located in the passing direction and could reach a pass within 0.4 seconds — assuming the average speed of a pass this roughly translates to a 5 meter radius of the passing origin. In all cases where this criteria is not fulfilled, we set the probability of the pass being blocked to zero. For the remaining passes (6% of them were blocked) we trained a XGBoost model to estimate the likelihood that a pass will be blocked. We used several features introduced in Table 2 (marked in column B) like the proximity of the nearest opposing player within the passing direction ($+/-$ 90 degrees), the location on the pitch (i.e. x/y-coordinates), the time of possession and the intended distance. The final blocking model was trained on 312, 413 passes (9, 372 blocked) with a split into 60% training, 20% validation and 20% test data. The outcome of the prediction is presented in Table 4. For comparison it includes a naive model using the average block probability (6.5%) as a baseline (row 2). The final hyperparameters of the blocking model can be found in Appendix B (Table 8).

6 Manual validation

We described statistical evaluations for each component of the entire approach in the respective chapters (i.e. Tables 3, 4). However, to further validate our results, we performed three separate expert-based validation studies of the following components:

(1) Synchronization of passing events with positional data
(2) Detection of the intended receiver for unsuccessful passes
(3) Outcome of the final xPass model

For each of the validation studies three different football experts looked at $3,600$ passes from 10 different games, with one game shared amongst all three, to gather the inter-rater reliability.

In order to evaluate (1) in the context of passes, the football experts were presented with identified time stamps, and they were tasked to annotate, whether the timestamps are correct. Overall they identified 99.1% of the timestamps as correct and had a pairwise inter-rater reliability of 99.3% However, since this approach is very binary (and potentially biased), we conducted a separate thorough evaluation study of the pass synchronization, described in Appendix A.

Since we can only systematically assess the accuracy of the intended receiver identification for successful passes (see Sect. 2), in (2) the subjects were tasked to identify the intended recipient of unsuccessful passes. Of the $1,307$ unsuccessful passes, our prediction agrees with the human labels in 72.0% of the cases and the inter-rater reliability is 96.2%.

The third and most relevant validation study evaluates, how well we can judge the difficulty of a pass (3). This is especially relevant because our final xPass values result from a combination of different machine learning models, each with their own inaccuracies. Therefore, the final outcome was evaluated manually by football experts. Estimating pass probabilities is a very challenging task for humans (even for football experts). To circumvent this issue, we provide experts with sets of two passes and let them assess which of the two is more difficult. Comparing passes with very similar xPass values is likely not a very reliable ground truth, and comparing passes with large xPass differences should be a trivial task with a high accordance between experts and our model. Therefore, in order to minimize the human-labeling effort, we group pairs of passes in three different categories based on their absolute xPass differences:

– Small difficulty difference ($< 10\%$),
– Medium difficulty difference ($10 — 30\%$),
– Large difficulty difference ($> 30\%$).

Per match we select 300 pass comparisons and the majority of them (90%) in the second category, 7% in the first category and 3% in the last category.[7]

Limited by the inter-labeler accordance, especially in critical situations, Table 5 shows that our model achieves satisfactory results. To investigate how much the addition of the blocking model helps the predictions, we further compute the accuracy of the model without a superimposed blocking model. This simpler model has a lower accuracy of 71.1% over the entire data set.

---

[7] The pass comparisons were randomly selected with the above described distribution, so the final numbers are subject to randomness and slightly deviate from the target distribution.

Table 5 The average pairwise accuracy are depicted and the number of pairs in a given subset are in the brackets

	Evaluation	Labeler accordance	xPass Model
1	Small difficulty difference	0.786 (70)	0.640 (286)
2	Medium difficulty difference	0.812 (787)	0.715 (3175)
3	Large difficulty difference	1.000 (30)	0.983 (118)
4	All	0.786 (887)	0.718 (3,579)

7 Discussion

A general limitation of our approach is its sensitivity to positional data accuracy. While the quality of tracking data has been increasing continuously over the past decade, the accuracy of ball tracking has not been properly validated in the literature yet (Anzer and Bauer 2021). The spatio-temporal synchronization of positional and event data — typically acquired through independent systems — presents crucial improvement for the analysis of passes. By training various models on different feature sets, we show how much each additional set increases the model's quality. Spearman et al. (2017) also pointed out the necessity of this synchronization step, but did not provide any details, nor an evaluation of their implemented approach. By adopting the methodology from Anzer and Bauer (2021) (synchronization of shot events) to passes, we use a reproducible approach (independent of the event-/tracking-data provider) and evaluate its accuracy manually in two independent experiments (Sect. 6 (1) and Appendix A).

Both, the player-movement model (Sect. 3.2) and the ball trajectory model (Sect. 3.1) draw heavily from previously published work (Brefeld et al. 2019; Spearman et al. 2017). We combine both to estimate the target of a pass and made only minor adjustments in order to improve the prediction accuracy on our data set. One thing we found regarding the movement model, is that the tighter the speed interval, the more the shape of the resulting hull is circular instead of elliptical, contrasting the findings of Brefeld et al. (2019), that finds oval shapes while using broader speed ranges. This could imply that movement ranges for particular initial speeds are circular, but when using a wide range of initial speeds, the total observed range is a combination of the movement circles along the movement direction, thus taking an elliptical shape.

Similar, as in Spearman et al. (2017) we ignore wind, rotation of the ball, and the Magnus force in the ball trajectory estimation. Our approach struggles to identify the intended receiver, when the underlying pass attempt fails completely. Fortunately, this case happens very rarely in the highest professional environments.

Implementing a separate blocking model guarantees that we have an unbiased estimation of pass probabilities. Furthermore, the manual validation (Sect. 6) shows, that it also more accurately coincides with expert assessments regarding the pass difficulty. The relatively low predictive power of the blocking model is likely caused by the nature and quality of the tracking data. The players' x/y-coordinates merely describe their center of gravity and, especially at the moment of the pass, centimeters may decide whether a pass is blocked or not. Therefore, as long as so-called limb-

tracking (recordings of players' entire bodies) does not become more widely available, it will remain hard to estimate if an opponent can extend their leg to block a pass.

Probabilistic metrics (e.g. expected goals, xPass) are hard to manually evaluate, since even experts cannot estimate a ground truth percentage reliably. For this reason we developed an evaluation study design delivering a useful ground truth while maintaining a high inter-labeler reliability. In previous research the quality of pass difficulty models was purely measured by the accuracy of the binary pass success classification. Our work goes one step further through a manual validation study with football experts, allowing us to also evaluate (1) the synchronisation of positional and event data, (2) the receiver estimation for unsuccessful passes, and (3) the pass difficulty. While in (2) we achieved, a reasonable accuracy of 72.0%, the experts showed a very high inter-rater reliability of 96.2%. This can in part be explained, by the fact that they were given video sequences of the passes extending far further than the 0.4 seconds our estimation uses. When examining the cases with the differences, we found that this is mostly caused by long balls (e.g. goal kicks, half-field crosses) where multiple players could be the target, but only one of them gets involved in an aerial duel. The human observers then chose the teammate that lost the aerial duel. But for the purpose of our work, in these cases the possible target players are very close to each other, meaning that the feature calculation and, therefore, their xPass values are very similar. Apart from that erroneous ball tracking data can lead to wrong target predictions (e.g. when the ball has a sharp cut, often called "elbow", in its trajectory after 0.4 seconds, without being touched). The much lower accuracy of the intended target identification achieved by Li et al. (2019) (for successful passes: 27.87%) and Vercruyssen et al. (2016) (for successful passes: 50.00%; for all passes: 41.00%) shows how difficult of a problem this generally is. However, this comparison is not completely fair, since they only consider information at the time of the pass, while we use the first 0.4 seconds after the pass as well.

We are able to increase the accuracy of successful pass estimation on a team level (Spearman et al. (2017): 80.5%, our approach: 91.5%) as well as for the task of predicting the receiving player (Spearman et al. (2017): 67.9%, our approach 89.9%). Power et al. (2017) presents a pass prediction with a root mean square error (RMSE) of 0.2483 which is slightly improved by our approach (RMSE: 0.2428). While our approach uses hand-crafted features, Stöckl et al. (2021) show in their work that using Graph Neural Networks one can forgo extensive feature crafting and achieve similar accuracy results for a variety of football related machine learning tasks. As one of their applications they compare how well they can predict if a pass will be completed using a GNN, without (hardly) any feature crafting, to a simple xPass model based on standard features and find both models to achieve a similar accuracy of 0.86 and 0.85. While their accuracy is below the one our model achieves (0.92), their work still shows, that GNN's are capable of quickly working with unstructured football data, and yet achieve a relatively high accuracy.

Overall, the major benefit of an expected pass model, is that it enables more granular analysis of passing behaviour than would be possible with simple pass completion rate metrics. It can be used to quantify players' and teams' performances (Spearman et al. 2017; Power et al. 2017), by looking at their risk profiles or their efficiency. For example, the players over-performing their expected completion percentages in the

Table 6 Top 10 xPass over-performers in Bundesliga Season 2020/2021 from matchday 1–15 (at least 200 passes)

	Player	Passing Performance	Average xPass
1	Kingsley Coman	5,3 %	85,44%
2	Raphaël Guerreiro	4,9 %	86,92%
3	Max Kruse	4,7 %	82,32%
4	Christopher Nkunku	4,5 %	83,49%
5	Sebastian Rudy	4,4 %	80,69%
6	Ritsu Doan	4,3 %	76,48%
7	Daniel Caligiuri	4,1 %	77,64%
8	Josip Brekalo	4,1 %	81,33%
9	Rafael Czichos	4,1 %	84,40%
10	Joshua Kimmich	4,1 %	88,84%

Performance is defined as the difference between completion rate and average xPass values

Bundesliga season 2020/2021 (up to matchday 15) the most are shown in Table 6. The column "Passing Performance" indicates how much a player's actual completion rate exceeds his average xPass values.

Another application is to evaluate possible pass options, and with that a player's decision making skills. At any given time while a player is in possession we can calculate success probabilities for hypothetical passes to teammates as shown in Fig. 4.[8] This is done by using a combination of the full model (Table 3, row 1) and the blocking model (Table 4, row 1). To find the ideal exit angle and velocity we perform a grid based search over "sensible" combinations and maximise for the xPass values. The resulting probabilities are shown for each teammate. We can see that Robin Koch (#2 of the blue team) chose one of the hardest pass options with a completion probability of 52.7%. The additional hypothetical passes — shown as lines with the respective success probabilities — come with some limitations: First, the probabilities are based on a data set, where players actively opted for a pass and since we can assume a certain amount of rationality in the decision making, values for hypothetical passes might be skewed as a consequence. Second, we assume that passes can be played at any time in any direction, without the need to properly set up before, which obviously warps reality. For instance, in the displayed situation the passing player decides to play a diagonal ball across the pitch to #7. The pass option of another long diagonal ball to number #3 — on the right side of the pitch — would require some preparation allowing opponents, especially #3 and #19 (of the red team), to get into a better defending positions. Furthermore, and this holds true in general for our xPass model, we simply compute the probability that a pass is successful, i.e. arrives at the intended target. It does not tell us if after the first touch, the teammate can hold the ball or looses it immediately thereafter.

---

[8] The video footage of the situation can be found here https://dfb-my.sharepoint.com/:v:/g/personal/pascal_bauer_dfb_de/EWIWkaF8Gp5CjPCdQRs5KXsB6Rt0LKHoKomXUFogNsR2Wg?e=u6EziX. The ball trajectory of the chosen pass is displayed in Fig. 1.
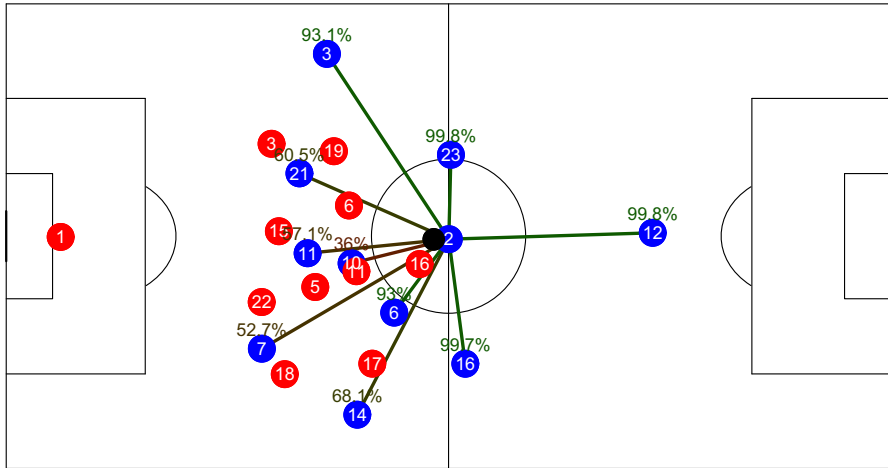
Fig. 4 Pass probabilities of hypothetical passes. This is the same situation as displayed and described in Fig. 3. The video sequence can be found here: https://dfb-my.sharepoint.com/:v:/g/personal/pascal_bauer_dfb_de/EWIWkaF8Gp5CjPCdQRs5KXsB6Rt0LKHoKomXUFogNsR2Wg?e=u6EziX

All together, we present a novel methodology for quantifying football's most relevant actions while addressing some of the shortcomings of previously published work and compare the results with existing literature. Our metric can be used to scout players outperforming their expected completion rates, identify and target weak spots in opposing teams, or show players alternative passing options they may have missed. To even better evaluate the decision making of a player, one would need to combine our risk model with a reward model (e.g. Steiner et al. (2019); Goes et al. (2019); Fernandez et al. (2018)) to not only assess a player's risk profile, but also whether they are making the best possible decision.

Declarations

Ethics and Reproducability By informing all participating players, all tracking is compliant to the general data protection regulation (GDPR)(https://gdpr-info.eu/, accessed 07/20/20.). An ethics approval for wider research program using the respective data is authorized by the ethics committee of the Faculty of Economics and Social Sciences at the University of Tübingen. The data are property of the DFL e.V. / DFB e.V. and cannot be shared public. However, interested researchers can request samples of data under non-disclosure agreement constraints at the respective institutions. With the description of the respective tracking

Table 7 Time shift of data synchronisation against manual label

Time Shift	00.00	00.04	00.08	00.12 − 00.48	00.52 − 02.00
Dispersion	41.6%	38.3%	8.6%	15.6%	4.1%

With our frequency of 25 Hz one frame equals 00.04 seconds

vendors and systems, peers working in the football industry can reproduce the results by using any kind of professional football data.

A The synchronization of pass events

Our algorithm to synchronize passes is based on the work introduced by Anzer and Bauer (2021) to synchronize shots. Since passes occur much more frequently (915 per game) than shots (23 per game), we need to adjust the algorithm to ensure we identify the right one. For that we add one additional rule: if the algorithm finds multiple pass moments in the considered time window by the passing player, we require that the actual receiver of the pass (if there is one), must be within a 2 m distance of the ball within 5 s after the pass moment.

For one Bundesliga match (FC Bayern München vs Borussia Dortmund, matchday 24 of the 2020/2021 season) football experts gathered frame-accurate timestamps from watching the video footage for every pass (1, 088 in total). We then compared these timestamps to the ones our synchronization algorithm produced. Overall only twelve passes showed a deviation of more than two seconds. After further inspection we found that seven of them were wrongly annotated in the manual collection process, and in the other five our algorithm identified a wrong timestamp, due to either faulty tracking data (1), blocked passes (2) or identifying the wrong of two options (2). Table 7 shows how much the time stamps differ for the remaining 1, 076 passes. As we can see in about 80% of the passes the synchronization finds either the same frame or the one next to it as the human.

B Details on XGBoost expected pass and blocking model

Table 8 shows the selected hyperparameters of our final xPass model (Table 3, row 1) and the hyperparameters of the blocking model (Table 4, row 1). Additionally, Figure 5 shows the feature importance of the xPass model using SHAP-values[9] (Lundberg et al.

---

[9] The abbreviation **SHAP** stands for **SH**apley **A**dditive ex**P**lanation.

Table 8 Final choice of hyperparameters for the xPass model (Table 3, row 1) and the blocking model (Table 4, row 1)

|	Hyperparameter	Description	Range	XPass	Blocking
1	Learning rate	Controls the step size used per update	[0, 1]	0.098	0.049
2	Max depth	Limits the depth of the tree	[0, 10]	8	6
3	Subsample	Controls number samples applied to the tree	(0, 1]	0.593	0.695
4	Min child weight	Controls instance weight of a node	[0, 10]	1.641	9.45
5	Class balancer	Controls the balance of negative and positive weights (Number of negative cases / Number of positive cases)	(0, ∞)	1	1
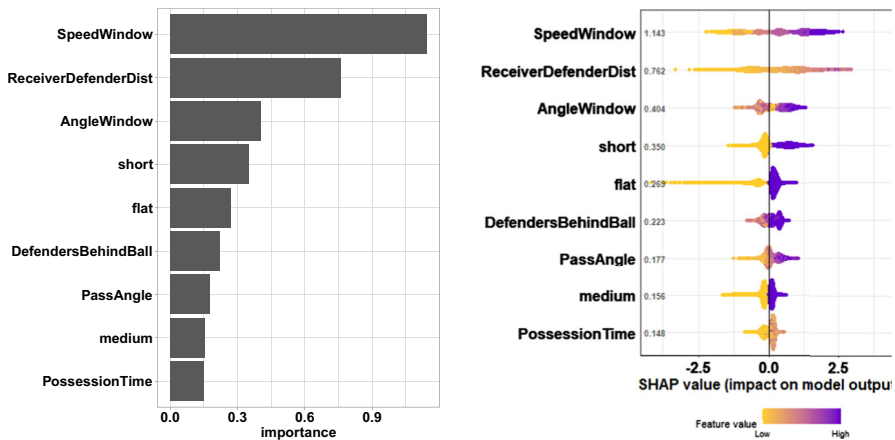
Fig. 5 Feature influence to the xPass model (Table 3, row 1) based on SHAP-values

2017). The left plot shows the absolute, overall influence of the respective features on our prediction. As discussed in Sect. 4, the speed window (*SpeedWindow*) in which the pass can be received by the intended target has the highest influence. The right plot in Fig. 5, where each colored dot (from yellow to violet) indicates the contribution of the feature to the model, shows that this relation is almost linear. The larger the speed window (i.e. violet plots), the higher the xPass value. This illustration also shows the influence of the binary features like *flat, medium* or *short*. The continuous feature, time of possession shows a dispersion of the dots similar to the binary features. This is caused by the fact, that direct played passes (i.e. 'one touch') are implicitly separated from all other passes by the model.

References

Andrienko G et al. (2017) Visual analysis of pressure in football. In: Data Mining and Knowledge Discovery 31.6, pp. 1793–1839. issn: 1573756X. https://doi.org/10.1007/s10618-017-0513-2 (cit. on pp. 6, 7)

Andrienko G et al. (2019) Constructing Spaces and Times for Tactical Analysis in Football. In: IEEE Transactions on Visualization and Computer Graphics 27.4, pp. 2280–2297. https://doi.org/10.1109/TVCG.2019.2952129. https://ieeexplore.ieee.org/document/8894420 (cit. on p. 1)

Anzer G, Bauer P (2021) A Goal Scoring Probability Model based on Synchronized Positional and Event Data. In: Frontiers in Sports and Active Learning (Special Issue: Using Artificial Intelligence to Enhance Sport Performance) 3.0, pp. 1–18. https://doi.org/10.3389/fspor.2021.624475. (cit. on pp. 2, 3, 9, 13)

Anzer G, Bauer P, Brefeld U (2021) The origins of goals in the German Bundesliga. J Sports Sci. https://doi.org/10.1080/02640414.2021.1943981

Arbués SA et al. (2020) Using player.s body-orientation to model pass feasibility in soccer. https://arxiv.org/abs/2004.07209

Asai T et al. (2007) Fundamental aerodynamics of the soccer ball. In: Sports Engineering 10.2, pp. 101–109. issn:1369-7072. https://doi.org/10.1007/bf02844207 (cit. on pp. 2, 3)

Battaglia PW et al. (2018) Relational inductive biases, deep learning, and graph networks (cit. on p. 2)

Bauer P, Anzer G (2021) Data-driven detection of counterpressing in professional football a supervised machine learning task based on synchronized positional and event data with expert-based feature extraction. Data Mining and Knowledge Discovery. https://doi.org/10.1007/s10618-021-00763-7

Bradley PS et al. (2013) The effect of high and low percentage ball possession on physical and technical profiles in English FA Premier League soccer matches. In: Journal of Sports Sciences 31.12, pp. 1261–1270. issn: 02640414. https://doi.org/10.1080/02640414.2013.786185 (cit. on p. 1)

Bransen L, Haaren JV (2019) Measuring football players' on-the-ball contributions from passes during games. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 11330 LNAI, pp. 3.15. issn: 16113349. https://doi.org/10.1007/978-3-030-17274-91 (cit. on p. 1)

Brefeld U, Lasek J, Mair S (2019) Probabilistic movement models and zones of control. In: Machine Learning 108.1, pp. 127.147. issn: 15730565. https://doi.org/10.1007/s10994-018-5725-1. (cit. on pp. 2, 4, 9)

Brier GW (1950) Verification of Forecasts Expressed in Terms of Probability. In: Monthly Weather Review 78.1, pp. 1.3. issn:0027-0644. https://doi.org/10.1175/1520-0493(1950)078<0001:vofeit>2.0.co;2 (cit. on p. 7)

Brooks J, Matthew K, John G (2016). Developing a data-driven player ranking in soccer using predictive model weights. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 49.55. https://doi.org/10.1145/2939672.2939695 (cit. on p. 1)

Chawla S et al. (2017) Classification of passes in football matches using spatiotemporal data. In: ACM Transactions on Spatial Algorithms and Systems 3.2. issn: 23740361. https://doi.org/10.1145/3105576 (cit. on p. 2)

Chen T, Guestrin C (2016) XGBoost: A scalable tree boosting system. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining 13-17, pp. 785.794. https://doi.org/10.1145/2939672.2939785 (cit. on p. 6)

Fernandez J, Bornn L (2018) Wide Open Spaces : A statistical technique for measuring space creation in professional soccer. In: MIT Sloan Sports Analytics Conference, Boston (USA), pp. 1.19 (cit. on pp. 2, 4, 10)

Fernández J, Bornn L, Cervone D (2020) A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. https://arxiv.org/abs/2011.09426

Francisco PA et al. (2020) Seeing in to the future : using self-propelled particle models to aid player decision-making in soccer. In: MIT Sloan Sports Analytics Conference. Boston (USA). pp. 1–23 (cit. on p. 2)

Goes F et al. (2021) A risk-reward assessment of passing decisions: comparison between positional roles using tracking data from professional men's soccer. In: Science and Medicine in Football 00.00, pp. 1.9. issn: 2473-3938. https://doi.org/10.1080/24733938.2021.1944660.(cit. on p. 2)

Goes FR et al. (2019) Not Every Pass Can Be an Assist: A Data-Driven Model to Measure Pass Effectiveness in Professional Soccer Matches. In: Big Data 7.1, pp. 57.70. issn: 2167647X. https://doi.org/10.1089/big.2018.0067. (cit. on pp. 2, 10)

Gómez JLI et al (2019) Landscapes of passing opportunities in Football . where they are and for how long are available ? In: Barça sports analytics summit February, pp. 1.14 (cit. on p. 2)

Hubáček O, Šourek G, Železný F (2018) Deep learning from spatial relations for soccer pass prediction. In: CEUR Workshop Proceedings 2284, pp. 162.169. issn: 16130073. https://dtai.cs.kuleuven.be/events/MLSA18/papers/hubacek_mlsa18.pdf (cit. on pp. 2, 6)

Król M et al. (2017) Pass Completion Rate and Match Outcome at the World Cup in Brazil in 2014. In: Polish Journal of Sport and Tourism 24.1, pp. 30.34. issn: 2082-8799. https://doi.org/10.1515/pjst-2017-0004 (cit. on p. 1)

Li H, Zhang Z (2019) Predicting the receivers of football passes. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 11330 LNAI, pp. 167.177. issn: 16113349. https://doi.org/10.1007/978-3-030-17274-9_15 (cit. on pp. 2, 9)

Linke D, Link D, Lames M (2018) Validation of electronic performance and tracking systems EPTS under field conditions. In: PLoS ONE 13.7, pp. 1.20. issn: 19326203. https://doi.org/10.1371/journal.pone.0199519 (cit. on p. 3)

Linke D, Link D, Lames M (2020) Football-specific validity of TRACAB's optical video tracking systems. In: PLoS ONE 15.3, pp. 1.17. issn: 19326203. https://doi.org/10.1371/journal.pone.0230179 (cit. on p. 3)

Lundberg SM, Lee SI (2017) Consistent feature attribution for tree ensembles. https://arxiv.org/abs/1706.06060

Mchale IG, Lukasz S (2014) A mixed effects model for identifying goal scoring ability of footballers. In: Journal of the Royal Statistical Society. Series A: Statistics in Society 177.2, pp. 397.417. issn: 09641998. https://doi.org/10.1111/rssa.12015 (cit. on p. 6)

McHale IG, Relton SD (2018) Identifying key players in soccer teams using network analysis and pass difficulty. In: European Journal of Operational Research 268.1, pp. 339.347. issn: 03772217. https://doi.org/10.1016/j.ejor.2018.01.018 (cit. on p. 1)

Nazareth JL (2004) An Optimization Primer. In: An Optimization Primer, pp. 2.5. https://doi.org/10.1007/978-1-4684-9388-7 (cit. on p. 6)

Oggiano L, Satran L (2010) Aerodynamics of modern soccer balls. In: Procedia Engineering 2.2, pp. 2473. 2479. issn: 18777058. https://doi.org/10.1016/j.proeng.2010.04.018. (cit. on pp. 2, 3)

Pappalardo L et al. (2019) A public data set of spatio-temporal match events in soccer competitions. In: Scientific Data 6.1, pp. 1.15. issn: 20524463. https://doi.org/10.1038/s41597-019-0247-7.(cit. on p. 3)

Pettersen SA et al. (2014) Soccer video and player position dataset. In: Proceedings of the 5th ACM Multimedia Systems Conference, MMSys 2014 (Singapore, March 2014), pp. 18.23. https://doi.org/10.1145/2557642.2563677 (cit. on p. 3)

Power P et al (2017) Not all passes are created equal: Objectively measuring the risk and reward of passes in soccer from tracking data. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining Part F1296:1605–1613. https://doi.org/10.1145/3097983.3098051 (cit. on pp. 1, 2, 6, 9, 10)

Reep C, Benjamin B (1968) Skill and Chance in Association Football Author. In: Journal of the Royal Statistical Society 131.4, pp. 581–585. issn: 14698005. https://www.jstor.org/stable/2343726?seq=1 (cit. on p. 1)

Rein R, Raabe D, Memmert D (2017) Which pass is better? Novel approaches to assess passing effectiveness in elite soccer. In: Human Movement Science 55.July, pp. 172–181. issn: 18727646. https://doi.org/10.1016/j.humov.2017.07.010. (cit. on p. 2)

Spearman W et al. (2017) Physics-Based Modeling of Pass Probabilities in Soccer. In: MIT Sloan Sports Analytics Conference, Boston (USA), pp. 1–14. https://www.researchgate.net/profile/William-Spearman/publication/315166647_Physics-Based_Modeling_of_Pass_Probabilities_in_Soccer/links/58cbfca2aca272335513b33c/Physics-Based-Modeling-of-Pass-Probabilities-in-Soccer.pdf (cit. on pp. 2–4, 6, 9, 10)

Stein M et al. (2017) How to Make Sense of Team Sport Data: From Acquisition to Data Modeling and Research Aspects. In: Data 2.1, p. 2. issn: 2306-5729. https://doi.org/10.3390/data2010002 (cit. on p. 1)

Steiner S et al. (2019). Outplaying opponents–a differential perspective on passes using position data. In: German Journal of Exercise and Sport Research 49.2, pp. 140–149. issn: 25093150. https://doi.org/10.1007/s12662-019-00579-0 (cit. on pp. 2, 7, 10)

Stöckl M et al. (2021) Making Offensive Play Predictable - Using a Graph Convolutional Network to Understand Defensive Performance in Soccer. In: MIT Sloan Sports Analytics Conference, Boston (USA), pp. 1–19 (cit. on pp. 2, 9)

Szczepański L, Mchale I (2016) Beyond completion rate: Evaluating the passing ability of footballers. In: Journal of the Royal Statistical Society. Series A: Statistics in Society 179.2, pp. 513–533. issn: 1467985X. https://doi.org/10.1111/rssa.12115 (cit. on p. 1)

Vercruyssen V, Raedt LD, Davis J (2016) Qualitative spatial reasoning for soccer pass prediction. In: CEUR Workshop Proceedings 1842. issn: 16130073 (cit. on pp. 2, 9)

Williams LRT (2000) Coincidence timing of a soccer pass: Effects of stimulus velocity and movement distance. In: Perceptual and Motor Skills 91.1, pp. 39–52. issn:00315125. https://doi.org/10.2466/pms.2000.91.1.39 (cit. on p. 1)