

un-xPass: Measuring Soccer Player’s Creativity

Pieter Robberechts
KU Leuven, Dept. of Computer
Science; Leuven.AI
B3000 - Leuven, Belgium
pieter.robberchts@kuleuven.be

Maaïke Van Roy
KU Leuven, Dept. of Computer
Science; Leuven.AI
B3000 - Leuven, Belgium
maaïke.vanroy@kuleuven.be

Jesse Davis
KU Leuven, Dept. of Computer
Science; Leuven.AI
B3000 - Leuven, Belgium
jesse.davis@kuleuven.be

ABSTRACT

Creativity is highly valued in soccer players. It contributes to exciting and unpredictable play, which can help teams to overcome defensive strategies and create scoring opportunities. Consequently, evaluating the creative abilities of players is an important aspect of the player recruitment process. However, there is currently no clear way to measure creativity in soccer. It is not captured by the typical result-based performance indicators, as being creative entails going beyond just doing something useful, to accomplishing something useful but in a unique or atypical way. Therefore in this paper, we define a novel metric to quantify the level of creativity involved in a player’s passes. Our Creative Decision Rating (CDR) utilizes machine learning techniques to assess two important factors: the originality of a pass, and its value in terms of increasing the team’s chances of scoring a goal. We validated our metric on StatsBomb 360 contextual event stream data of the 2021/22 English Premier League season and show through a number of use cases that it provides another angle on a player’s skill, complementing existing player evaluation metrics. Overall, our metric provides a concise method for capturing and quantifying the creativity of soccer players and could have important implications for player recruitment and talent development in the sport.

CCS CONCEPTS

• **Applied computing**; • **Computing methodologies** → **Model development and analysis**; **Machine learning**;

KEYWORDS

Sports Analytics, Soccer, Football, Creativity, Spatiotemporal Statistics

ACM Reference Format:

Pieter Robberechts, Maaïke Van Roy, and Jesse Davis. 2023. un-xPass: Measuring Soccer Player’s Creativity. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD ’23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3580305.3599924>

1 INTRODUCTION

Quantitative analysis plays a growing role in the recruitment process of soccer players [2, 7, 34]. The wide availability of detailed data and extensive research has led to the creation of new metrics, which offer an objective and more complete picture of a player’s abilities

and potential. Nowadays, teams are increasingly using these metrics to help assess a player’s technical [8, 28], physical [19, 26], tactical [9, 10] and mental [3] skills.

Another skill that is highly valued by coaches, scouts and fans is creativity, which in soccer refers to a player’s ability to come up with new and imaginative solutions to problems posed by the opposing team. Because the presence of creative players within a team enhances its ability to break down tight defenses and create goal-scoring opportunities [16], creativity has been emphasized as an important factor of success by both researchers [23] and practitioners [5]. Also among fans, the best soccer players are often described as being “creative”. For example, Kevin De Bruyne is widely regarded as a master of creativity.

Despite being highly valued, no previous research has looked at statistical models to objectively quantify creativity. That is because creativity is generally seen as an intangible quality, something which cannot be analyzed through statistics. In psychology, creativity is commonly defined as “the ability to produce work that is both novel (i.e., unexpected, original) and appropriate (i.e., useful)” [23]. Hence, creativity involves not only accomplishing something useful, but doing it in a unique or atypical way. For example, a cutback pass almost always puts a player in a good scoring position. Yet, it is typically not perceived as being creative as it is generally the most straightforward option when the game state allows it.

In this paper, we operationalize this psychological definition in the context of soccer by positing that creative actions should (1) differ from the typical action that the vast majority of players would select in a given game state, and (2) have more promising results than this typical choice. Based on this, we propose a novel Creative Decision Rating (CDR) metric for passes. The CDR assigns a single value per pass which is the difference between the expected values of the chosen pass option and the predicted typical pass. In order to compute the CDR, we need three estimates: the likelihood of each possible pass destination, the long-term reward of each passing option and the success probability of each passing option.

We use machine learning methods to learn models for each of these components from StatsBomb 360 event stream data. This data describes all on-the-ball actions that were performed during a game, augmented with snapshots of player locations at the time of each annotated event.¹ Although models for each of the three aforementioned components have been previously developed for traditional event stream data and spatiotemporal tracking data, the hybrid 360 data poses unique challenges. We will compare two

KDD ’23, August 6–10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD ’23)*, August 6–10, 2023, Long Beach, CA, USA, <https://doi.org/10.1145/3580305.3599924>.

¹ See <https://statsbomb.com/what-we-do/soccer-data/360-2/> for a detailed description of the data. A subset of the data used in this research is publicly available at <https://github.com/statsbomb/open-data>.

approaches for learning each of the separate models: a gradient-boosted trees model based on handcrafted features and a deep learning CNN model based on the SoccerMap architecture [12].

To validate our metric, we compute the CDR for the 2021/22 English Premier League season and show that Kevin De Bruyne is indeed the most creative player, followed by Tariq Lamptey and Trent Alexander-Arnold. When looking at pairs of players, we found that the interactions between Mohammed Salah and Sadio Mané exhibited the highest creativity. Additionally, we found that the average level of creativity is not affected by the game state (i.e., time remaining and goal difference), but that the variance of creativity increases as time progresses.

To summarize, this paper makes the following contributions:

- (1) We propose a novel metric, the CDR, for capturing the complex notion of creativity in soccer;
- (2) We compare deep learning and feature-based approaches for estimating (i) the likelihood of each possible pass destination, (ii) the long-term reward of each passing option, and (iii) the success probability of each passing option from snapshots of player positioning in StatsBomb 360 data;
- (3) We present a number of use cases and highlight the most interesting findings derived from our metric;
- (4) We publish a Python package (available at <https://github.com/ML-KULEuven/un-xPass>) that implements the CDR metric and enables the construction of the underlying machine learning models.

The remainder of this paper is organized as follows: Section 2 presents our metric for measuring the creativity of a player’s pass selection. Next, Section 3 presents the experimental setup, and Section 4 presents a detailed description of the approaches used for learning the machine learning components that underlie our metric. Section 5 provides insights into how our metric performs in practice and presents our most interesting findings. Section 6 discusses related work. Finally, Section 7 concludes the paper and outlines potential directions for future research.

2 MEASURING CREATIVE PASSING

Our definition of creativity requires quantifying two things about actions:

- (1) **Usefulness**, which is tied to how an action affects a team’s chance of scoring a goal
- (2) **Originality**, which is tied to the pass that the vast majority of players would select in a given game state

After we describe how to model these, we show how both can be combined to measure creativity. Our approach is based on event stream data, which describes on-the-ball actions such as passes, dribbles and shots observed during the match. Each action is described by a tuple of attributes a_i containing the action type, its time of occurrence in the match, the origin and destination location, the player who performs the action, and the body part used to execute the action.

2.1 Valuing the usefulness of actions

When considering event stream data, a soccer match can be viewed as a sequence of n consecutive actions a_1, a_2, \dots, a_n . Each action

a_i with outcome $o_i \in \{\text{success, fail}\}$ moves the game from state $S_{i-1} = \{a_1, \dots, a_{i-1}\}$ to state $S_i = \{a_1, \dots, a_{i-1}, a_i\}$. Consequently, a logical approach to capture the usefulness of actions is by measuring the difference in quality between the pre-action game state S_{i-1} and post-action game state S_i . In recent years, several performance metrics have been introduced based on this idea (e.g., xT [30], VAEP [8], g+ [17], PV [21], OBV [31]). At a high level, they all quantify the usefulness (U) of actions according to the following equation:

$$U(A = a, O = o|S) = Q(S|A = a, O = o) - Q(S),$$

where Q captures the value or quality of a particular game state. Generally, this quality is expressed in terms of the likelihood of scoring or conceding a goal. Thus, U captures the change in scoring probability following the execution of a particular action.

This typical approach to valuing an action’s usefulness is based on a holistic viewpoint. However, at the atomic level, each action comes in three parts:

Action selection. First, a player must select the best possible action from a set of options. This involves choosing which of several types of actions (e.g., pass, dribble, shot) to take and how to exactly implement (“execute”) the action (e.g., which corner to target on a shot), within the scope of the current game state.

Action execution. Second, the player must attempt to successfully carry out that action. Each player has different technical abilities though [14], and this will affect the action’s implementation.

Action resolution. Third, a player’s decisions and execution are translated into results. Typically, these are determined by teammates (e.g., the player receiving a pass) or opponents (e.g., the goalkeeper stopping a shot).

We posit that creativity is part of the action selection phase and thus try to make an abstract away the result.² Therefore, we compute the expected usefulness of an action as the weighted sum of the value of both outcomes:

$$\mathbb{E}[U(A = a|S)] = \sum_{o \in \{o^+, o^-\}} P(O = o|A = a, S) \cdot U(A = a, O = o),$$

where $P(O = o^+|A = a, S)$ is the probability that action a succeeds in game state S , and $P(O = o^-|A = a, S)$ is the probability that it fails.

2.2 Valuing the originality of actions

The second aspect of creativity concerns originality with respect to the choice of actions. Generally speaking, in soccer, there are three possible high-level actions a player can perform with the ball: a player can try to shoot at the goal, pass it to another teammate, or drive with the ball up the pitch. In this paper, we restrict the space of possible actions A to passes, since these are the type of action that is mostly linked to creativity.

More concretely, we characterize a pass $a_i \in A$ by its origin and destination. The origin corresponds to the ball’s current location and is included in the current game state. The destinations can

²Obviously, a player needs to pair creativity with technical abilities to make an impact. We discuss this in Section 5.2.

be defined in terms of player identities (i.e., passing to a specific player) or in terms of locations on the field. To be able to distinguish between simple passes to a specific player’s location and (often more creative) through balls that a player should run into, we opt for the second approach in this research. Thus, let L be the set of all the possible locations on a soccer field. Then, we can define D_i to be the selected pass destination location of pass a_i and $P(D = l|S)$ to be a transition probability model for passes to any location $l \in L$. Finally, the originality of a player’s pass selection decision can be defined as the complement $1 - P(D = l|S)$.

2.3 Valuing the creativity of actions

By comparing the expected usefulness of the chosen pass with that of other typical passes that the game state allowed, we can now assess whether a player is performing unexpected actions that lead to more promising results, thereby indicating a notion of creativity. Our intuition is that there is seldomly more than one obvious high-value pass option in any game state and that professional soccer players are all experts at selecting such options. Therefore, it is generally sufficient to compare the selected pass with the predicted most likely pass option $l_{\text{typical}} = \arg \max_{k \in L} P(D = k|S)$.

Combining these insights, we define the creative decision rating (CDR) of a pass to a location $D = l$ in a game state S as:

$$CDR(D = l|S) = \mathbb{E}[U(D = l|S)] - \mathbb{E}[U(D = l_{\text{typical}}|S)].$$

Since the chosen action and the most typical action have the same pre-action game state, this reduces to

$$\begin{aligned} &= \mathbb{E}[U(D = l|S)] - \mathbb{E}[U(D = l_{\text{typical}}|S)] \\ &= (\mathbb{E}[Q(S|D = l)] - \mathbb{E}[Q(S)]) - (\mathbb{E}[Q(S|D = l_{\text{typical}})] - \mathbb{E}[Q(S)]) \\ &= \mathbb{E}[Q(S|D = l)] - \mathbb{E}[Q(S|D = l_{\text{typical}})] \\ &= \sum_{o \in \{o^+, o^-\}} P(O = o|D = l, S) \cdot Q(S|D = l, O = o) \\ &\quad - \sum_{o \in \{o^+, o^-\}} P(O = o|D = l_{\text{typical}}, S) \cdot Q(S|D = l_{\text{typical}}, O = o) \end{aligned}$$

This leads us to the task of estimating three components that produce a single estimation of pass creativity when combined:

- (1) Pass selection $P(D = l|S)$: an estimate of the likelihood of a pass being made to every other location on the field.
- (2) Pass success $P(O = o|D = l, S)$: an estimate of pass success probability for a pass to every other location on the field.
- (3) Pass value $Q(S|D = l, O = o, S)$: an estimate of the game state value following a successful or unsuccessful pass to every other location on the field.

Each component can be estimated utilizing a standard supervised machine learning pipeline, where given some input features describing the game state, we can train a model to yield a probability between 0 and 1 for each location on the pitch. The next two sections provide a detailed description of our approach to train these models.

3 EXPERIMENTAL SETUP

In this section, we describe the dataset used and experimental settings for the inference and evaluation of each model component.

3.1 Data

We build our datasets based on StatsBomb 360 event stream data. This type of data is extracted from broadcast video and consists of the regular human-annotated event stream data enhanced with snapshots of player positioning. We work with the SPADL representation [8] of this event stream. The 360 snapshots are recorded at the time of each on-the-ball action and include the location and relationship to the ball carrier (i.e., teammate or opponent) of all players visible in the video.

Following our model design, we focus exclusively on passes. Additionally, we discard passes not executed with the foot, passes originating from dead-ball situations (i.e., we discard corners, free-kicks, goal-kicks, kick-offs, and throw-ins), and passes where the origin or destination location falls outside the 360 snapshot. By applying these criteria, we create a training dataset of 118,758 passes from the data of the 2020 European Championship, and the top-10 teams in the 2020/21 seasons of the English Premier League. A random sample of 20% of these passes is used as a validation set for model selection. A dataset of 93,631 passes extracted from the top-10 teams in the 2021/22 season of the English Premier League is set apart for evaluating the models and developing the use cases.

3.2 Model choice and settings

A key design choice revolves around how to represent the game state. One option is to manually create meaningful features using expert knowledge [8, 28]. This approach provides interpretability and leverages our understanding and insights about the game. Alternatively, it is possible to use a purely data-driven approach to learn a good representation of the game situation [11, 13, 18]. The former is typically used for modeling event stream data, while the latter is common for modeling full spatial tracking data. Since it is unclear which approach fits best for the hybrid 360 event stream data, we compare two classes of models: a deep learning model based on the SoccerMap [12] architecture and an XGBoost [6] model based on hand-crafted features.

The SoccerMap architecture is a fully convolutional network-based architecture, inspired by the deep jet nonlinear feature hierarchy [20]. Given a tensor-based representation of the ball’s and players’ current locations, the network can be trained to produce location-wise predictions (i.e., a 104×68 probability surface covering the full extent of a soccer field) with regard to future actions and results in a purely data-driven manner. The combination of layers at different resolutions allows for capturing relevant information about the current game state at both local and global levels, enabling the model to make local predictions that respect global structure. We train the SoccerMap models using the PyTorch Lightning framework with the adaptive moment estimation (ADAM) algorithm. We perform a grid search on the learning rate ($1e-3$, $1e-4$, $1e-5$, $1e-6$), and batch size parameters (16, 32, 64). We use early stopping with patience set to 10 epochs and a delta of $1e-3$ for the pass success probability model, and $1e-5$ for the pass selection and pass value models.

For all XGBoost models, we applied the Tree-structured Parzen Estimator (TPE) algorithm for optimizing the max tree depth ([1, 9]), learning rate ([$1e-2$, 0.25]), L1 regularization ([$1e-8$, 100]), L2

regularization ($[1e-8, 100]$), and the minimum loss reduction required to make a further partition on a leaf node ($[1e-8, 1]$). We use early stopping with patience set to 100 boosting rounds. To obtain a probability surface that covers the full extent of a soccer field, we simulate a pass to each of the $l \times h$ grid cells and compute the corresponding feature representation. The prediction for each simulated pass is then mapped to the corresponding grid cell. For computational reasons, we use a coarser 26×17 grid, which we then upscale using bilinear interpolation.

4 INFERENCE OF MODEL COMPONENTS

In this section, we provide a detailed description and empirical evaluation of the machine learning components required to compute the creative decision rating of an action. In our evaluation of these components, we address the following four research questions:

- Q1 Different representations of the input data are generally used for analyzing event stream data (i.e., features based on domain knowledge) and tracking data (i.e., data-driven deep learning). Which of these approaches is best for learning from StatsBomb 360 event stream data?
- Q2 Estimating a pass selection probability surface requires reasoning about the possible future locations of all other players. Is the incomplete player location context available in the StatsBomb 360 data sufficient to obtain a reasonably accurate estimate?
- Q3 The end location of a pass is an important feature when estimating the success probability of a pass. How should one deal with the unknown (intended) end location of failed passes?
- Q4 Previous work has suggested the use of expected goals (xG) values instead of binary goal / no goal labels to estimate the probability of a goal but both approaches have not been compared.

4.1 Pass selection

As discussed in Section 2, our goal is to estimate the probability of making a pass to every other location on the field (e.g., Figure 1), rather than to specific players. Estimating such a full pass selection probability surface requires a model that can make fine-grained local predictions that are calibrated globally over the entire soccer field, which is difficult to achieve with an XGBoost model. Therefore, we use SoccerMap to implement this component. The target output is a sparse matrix where a value of 1 is given for every observed pass in its corresponding destination location. As the input, we use the following nine channels which are adapted from [13]:

- Channel A & B: Two sparse matrices with the locations of the players in the attacking and defending team, respectively.
- Channel C & D: Two dense matrices with the distance to the ball and the goal for every location.
- Channel E & F: Two dense matrices with the sine and cosine of the angle between every location and the ball location.
- Channel G: One dense matrix with the angle between every location and the goal.
- Channel H & I: Two sparse matrices with the two components of the velocity vector of the ball, derived from the timestamps and ball location in the event data during the




Figure 1: Pass selection probability surface on a logarithmic scale for an example game situation where purple colors represent more likely pass destinations. Blue and red circles represent the player’s location of the attacking and defending team respectively. The blue team is in possession and plays left to right. The white arrow represents the selected pass.

two preceding actions. These channels give an indication of the direction in which the ball is moving.

Due to the lack of other approaches for estimating the full pass selection surface, we compare our results against approaches that predict the most likely receiver. First, we establish a naive baseline that consistently predicts the nearest teammate as the most likely receiver. Second, we train an XGBoost ranking classifier to predict the likelihood that a given player is effectively the receiver of a possible pass based on the following set of handcrafted features for each pass option in the 360 snapshot: origin and destination location, pass distance, pass angle, angle to goal at origin and destination, distance to the nearest defender at the destination, and distance of the nearest defender to a straight line between origin and destination. Per pass, one data point receives a positive label, the effective receiver, and all the other data points are labeled zero.

Table 1 presents the results for the baseline models and the SoccerMap model on the test data. To allow for a comparison between the location-based SoccerMap model and the receiver-based baseline models, we compute the most likely receiver from the prediction surface of the SoccerMap model as the teammate closest to the most likely pass destination and compute the accuracy as the percentage of passes for which the receiver was predicted correctly. The SoccerMap model is only slightly less accurate than the feature-based model. Although, it should be noted that mapping the location-based predictions to players introduces an additional error, in particular for passes that are played at a distance from the player’s current location. Hence, the difference in accuracy between the XGBoost and SoccerMap models is likely smaller in reality than what the evaluation reflects. Based on these results and a visual inspection of the estimated pass selection probability surfaces (e.g., Figure 1), we conclude that we can sufficiently accurately predict the likelihood of each possible pass destination.

4.2 Pass success probability

For estimating pass success, we define a binomially distributed outcome, according to the definition of success used by StatsBomb.

Table 1: The average loss and accuracy of three alternative implementations for the pass selection component: (1) a naive baseline predicting the nearest teammate as the most likely receiver, (2) an XGBoost model based on handcrafted features and (3) a deep learning model based on the SoccerMap architecture.

Model / Feature set	LogLoss	Accuracy
Closest teammate		0.395
XGBoost		0.538
SoccerMap	6.277	0.513

That is, passes that reach a teammate on-side are labeled as “successful”. All passes that go out of bounds or that are intercepted are labeled as “failed”.

One important caveat with respect to modeling pass success is that the intended target location of a pass is only known for successful passes. The end location of failed passes is recorded as the location where the ball was intercepted or went out of bounds. Hence, it is impossible to construct an accurate feature representation for attempted passes. Previous work has addressed this issue in two ways: (1) by ignoring the problem and assuming that most passes will be intercepted near their intended destination [12], and (2) by estimating the intended receiver based on the direction of the pass and the positions of potential receivers [1, 27]. Approaches to estimate the intended receiver range from simple distance-based rules to advanced physics-based approaches that model the ball trajectory and player movement [1]. As the 360 snapshots lack information about the velocity of players and the ball, we follow the approach of Power et al. [27] and estimate the intended receiver as the one closest to where the ball was intercepted and with the smallest angle to the line of the pass:

$$\text{Expected receiver} = \frac{\text{Min Distance}}{\text{Distance}} \times \frac{\text{Min Angle}}{\text{Angle}} \quad (1)$$

If the predicted receiver is positioned outside the field’s boundaries, we clip its coordinates to the field’s nearest boundary. If no player is within 20° of the pass line, we assume that the intended receiver is not included in the snapshot and we proceed with the observed end location. Figure 2 illustrates this procedure.

Estimating the intended receiver has some obvious limitations though. First, the estimate will be inaccurate when the intended receiver is not included in the 360 snapshot, when the ball is blocked early in its flight path, when the ball is deflected, or when two players are close to each other. In addition, the end location of attempted passes is replaced by the (x,y) coordinates of the expected receiver at the time of the pass. In reality, a pass is often played in front of the receiver to run onto.

We implemented three simple location-based gradient-boosted trees models to evaluate whether the advantages of identifying the intended receiver outweigh the limitations: (i) receiver-agnostic: a baseline model in which we use only the origin of the pass; (ii) observed end location: we use both the observed origin and end location of the pass; and (iii) intended end location: we replace the end location of the pass by the coordinates of the most likely receiver. Instead of the raw coordinates, we use the distance and

angle to the goal and the distance to the sideline as features in each of these models. Additionally, we describe the relationship between the start and end location of a pass by the pass distance (total and along both axes) and angle.

Our results show that using the intended end location provides little to no advantage. First, we can observe that using the observed end location is more accurate (Table 2). Although, we hypothesize this is because the observed end location often gives away the outcome of a pass. For example, passes with an observed end location on the sideline or within a 1 meter radius of the origin are most likely failed passes. More importantly, using the intended end location produces inaccurate estimates for through balls, because these are mapped to player locations when they fail. For example in Figure 3, the model using the intended end location predicts the prior probability for a large part of the opponent’s penalty box, since there are no examples with a potential receiver on that location in the training data. The main downside of using the observed end location is that locations near the pass origin are assigned very low success probabilities, which is probably due to blocked passes.

Given these results, we build further upon the observed end location and extend the location-based feature set with attributes available in the event stream:

- (1) The ball height (ground, shoulder level, above shoulder level);
- (2) The speed (distance covered / time) during the two preceding actions in the possession sequence;
- (3) The time the passer was in ball possession before attempting the pass.

Furthermore, we craft features from the 360 snapshots that describe the situation around the passer, receiver and ball trajectory:

- (4) The distance of the nearest defender to the passer and receiver;
- (5) The distance of the nearest defender to a straight line connecting the pass location and the receiver’s location;




Figure 2: The intended receiver (green circle) is identified as the teammate closest to where the ball was intercepted or went out of bounds (red circle) that has the smallest angle to the line of the pass. Only players within 20° of the pass line (gray area) are considered as potential intended receivers.




Figure 3: Pass success probability surface for a pass starting at the location of the blue circle to all other locations using a receiver-agnostic model, a model trained on the observed end location of each pass, and a model trained on the estimated intended end locations.

- (6) The number of opposing players in the passing path, where the path is defined as a triangular corridor between the pass origin and the receiver’s location with a base of 1 meter at the receiver’s location.

These features are inspired by recent work of Szczepański and McHale [33], Anzer and Bauer [1], Goes et al. [15] and Power et al. [27].

Next to the feature-based models, we again learn a SoccerMap deep learning model. As the input, we use channels A through G as defined in the previous section. Finally, we establish a naive baseline that assigns the average pass completion to all passes. As shown in Table 2, the XGBoost model using the full set of handcrafted features generally performs best.

4.3 Pass value

The final two components correspond to the expected usefulness of completed and failed passes, which reduces to estimating the value of a game state. Several approaches have been introduced to capture the value of game states [8, 13, 29], but generally they estimate the likelihood of scoring (i.e., the offensive value) or conceding (i.e., the defensive value) a goal in the near future. We train two separate models for these offensive and defensive components of game state value; however, we use an equivalent architecture for both cases. This allows us to inspect the model’s predictions at a higher level of granularity. The total game state value is then obtained as the difference between both models’ estimates.

Table 2: The performance of three alternative implementations for the pass success probability component: (1) a naive baseline predicting the average pass completion, (2) an XGBoost model based on handcrafted features and (3) a deep learning model based on the SoccerMap architecture.

Model / Feature set	Prec.	Rec.	F1	AUC	Brier
Avg pass completion	0.858	1.000	0.924	0.500	0.122
XGBoost					
Receiver agnostic	0.871	0.988	0.926	0.704	0.110
Intended end location	0.901	0.980	0.939	0.849	0.131
Observed end location	0.916	0.975	0.945	0.886	0.075
+ event attributes	0.925	0.970	0.947	0.916	0.068
+ 360-based features	0.936	0.966	0.951	0.939	0.062
SoccerMap	0.926	0.975	0.950	0.930	0.065

We define the offensive value as the probability of scoring a goal in the next 10 actions; and correspondingly the defensive value as the probability of conceding a goal in the next 10 actions [8]. Since goals are rare and provide a weak learning signal, previous work has proposed the use of xG values to train possession value models [17, 32]. Therefore, we make use of StatsBomb’s xG values to determine the likelihood of scoring from any shot in the next 10 actions. If this sequence of 10 actions contains multiple shots, we combine their xG value as

$$xG_{\text{seq}} = 1 - \prod_{\text{shot} \in \text{seq}} 1 - xG_{\text{shot}}.$$

This corresponds to taking the complement of the probability that the defending team does not allow a goal from a sequence of shots. Our results (Table 3) show that a model trained on xG values indeed provides more accurate estimates than a model trained on a binary goal / no goal variable.

Table 3: The performance of five alternative implementations for the offensive pass value component: (1) the location-based expected threat (xT) framework, (2) the event data-based VAEP framework trained on binary goal/no-goal labels and (3) trained on xG values, (4) the VAEP framework with features extracted from 360 snapshots trained on xG values, and (5) a deep learning model based on the SoccerMap architecture. All models are evaluated on completed passes only using binary goal/no-goal labels.

Model / Feature set	AUC	Brier	LogLoss
xT	0.741	0.012	0.062
VAEP	0.757	0.011	0.056
VAEP-xG	0.766	0.011	0.056
VAEP360	0.776	0.011	0.055
SoccerMap	0.743	0.011	0.064

The key difference between various game state value models is the amount of contextual information they take into account. The most basic approaches only consider the location of the ball [29, 30, 36], while intermediate approaches leverage more contextual information that can be extracted from traditional event stream data [8, 17, 21, 31], and the most advanced approaches extract a detailed spatial representation from tracking data accounting for the locations of all players [13]. However, to the best of our knowledge, no public models exist for estimating the game state value based on 360 data,³ which can be seen as a limited version of the full spatiotemporal tracking data as it provides only a partial view of the player’s locations and lacks information about the velocity and acceleration of players and the ball. Again, we experiment with a set of XGBoost models based on handcrafted features and a deep learning model based on the SoccerMap architecture. As baselines, we consider the location-based Expected Threat (xT) [30]⁴ and event data-based VAEP [8] frameworks.

³Moreover, StatsBomb’s proprietary On Ball Value (OBV) metric does not use the 360 snapshots yet.

⁴We use the xT grid provided at https://karun.in/blog/data/open_xt_12x8_v1.json

Table 4: The top-10 players who completed at least 250 passes in the 2021/22 English Premier League season in terms of our Creative Decision Rating (CDR), normalized for minutes played.

	Player	Team	CDR90
1	Kevin De Bruyne	Manchester City	0.0834
2	Tariq Lamptey	Brighton	0.0817
3	Trent Alexander-Arnold	Liverpool	0.0584
4	Raphinha	Leeds United	0.0563
5	Hakim Ziyech	Chelsea	0.0527
6	Martin Ødegaard	Arsenal	0.0480
7	Lucas Moura	Tottenham Hotspur	0.0473
8	Harry Kane	Tottenham Hotspur	0.0472
9	Bukayo Saka	Arsenal	0.0415
10	Mason Mount	Chelsea	0.0413

For the XGBoost models, we extend VAEP’s game state representation with features extracted from the 360 data that capture the number of outplayed players and the ball interceptability. The number of outplayed players is computed using a simplified version of Impect’s Packing Rate: a defender is considered to be packed if he is positioned between the ball and the goal before a forward pass, but would be further from the goal than the ball after the pass. For interceptability, we use the number of defenders in a 3 meter and 5 meter radius around the pass’ end location. For the SoccerMap models we use the seven input channels described earlier (A- G), as well as two dense matrices with the number of players of the team in possession and the opposing team that would be packed after a pass to each location.

We need the expected value of both completed and failed passes. Since the outcome of the pass is a feature in VAEP’s game state representation, a value for completed and failed passes can be obtained by modifying the feature value. For the SoccerMap models, we learn separate models on completed and failed passes. For simplicity, we only report the performance of all offensive game state value models on completed passes in Table 3.

5 USE CASES

We now present a number of observations that result from computing the CDR metric for the 2021/22 Premier League season. First, we present a ranking of the most creative players and quantify how their creativity pairs with their technical abilities. Second, we look at the creativity of interactions between pairs of players. Third, we look for a relationship between a player’s position and creativity. Fourth, we look at the effect of the game state on creativity.

5.1 Most creative players

Table 4 shows the top-10 players in terms of Creative Decision Rating per 90 minutes (CDR90) in our dataset. To obtain a robust ranking, we only include players who performed at least 250 passes according to the criteria defined in Section 3. Kevin De Bruyne tops the ranking, closely followed by Tariq Lamptey. Liverpool’s right-back Trent Alexander-Arnold completes the top three.

There is a natural tension between the creativity and quantity of a player’s passes. This is illustrated in Figure 4, which shows the number of passes that players execute on average per 90 minutes (quantity) in function of the average creativity of these actions. The reason is twofold. First, players who perform more passes generally play in a more defensive position that requires less creativity to progress the ball. For example, Harry Kane has a very high average CDR, but attempts relatively few passes as a striker. Second, if a player performs a high number of passes, then it is harder for each pass to have a high value. However, as shown by the dotted isoline, players like Kevin De Bruyne, Tariq Lamptey, Trent Alexander-Arnold, Raphinha and Hakim Ziyech pair a high creativity with a large number of passes.

5.2 Combining creativity and technical skills

Our creativity metric rewards players who attempt atypical passes, regardless of the result of the pass. While this enables a fine-grained evaluation of a player’s creativity, it omits an important piece of a player’s performance evaluation. To assess whether a player pairs vision with the technical abilities to successfully complete the passes he attempts, we apply the execution rating metric proposed by Bransen et al. [3]. Their metric measures the technical execution quality of a pass as the difference between the observed outcome of the pass (e.g., did the cross reach a teammate) and the predicted probability that the pass would be successful based on the context under which it was performed. Formally, the Execution Rating (ER) is defined as

$$ER(O = o, D = l|S) = [o^+] - P(O = o|D = l, S), \quad (2)$$

where $[o^+]$ takes the value of one if the pass succeeds and is zero otherwise, and $P(O = o|D = l, S)$ is given by the action success predictor from the previous section. Intuitively, the metric rewards

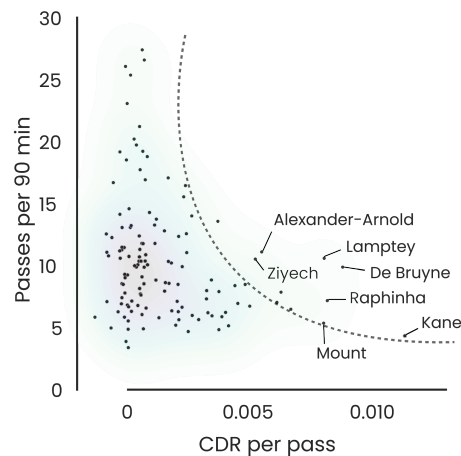


Figure 4: Scatter plot contrasting a player’s average Creative Decision Rating (CDR) with the average number of passes performed per 90 minutes. The dotted isoline shows the gap between the top-ranked players and the rest. Only players who completed at least 250 passes in the 2021/22 English Premier League season are included.

players who successfully perform difficult passes and punishes players who flub an easy pass.

Figure 5 compares the average CDR and ER of passes, grouped by player. In general, we found no strong relationship between the creative and technical abilities of players. The execution rating is mainly determined by a player’s position, with defensive midfielders scoring the highest and strikers scoring the lowest. We assume this is because the consequences of losing possession are less detrimental for attacking players. Hence, they can afford more mistakes.

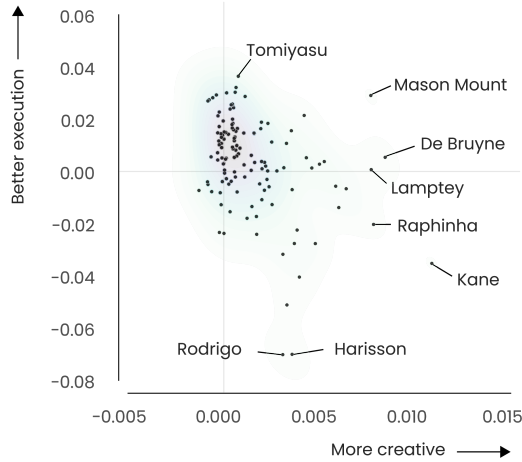


Figure 5: The relation between average creative decision rating and execution rating for each player in the 2021/22 English Premier League season.

5.3 Most creative duos

A limitation of our creativity metric is that it gives all credit to the player who gives the pass. Yet, it is often the player on the receiving end who enables the pass by slipping into a pocket of space. Therefore, Table 5 looks at the pairs of players who exhibited the highest creativity with mutual passes. Unsurprisingly, the pairs that rank highest have spent many minutes together on the pitch and are fully attuned to each other. Mohamed Salah and Sadio Mané have led Liverpool’s attack for five seasons, while Harry Kane and Heung-Min Son are a sterling duo at Tottenham Hotspur.

5.4 The link between creativity and position

Being creative is most often associated with the role of a playmaker or “number 10”, who typically operates from the position of an attacking midfielder. Figure 6 shows that players who operate from this position indeed score high on our creativity metric. However, we found that wingers are on average the most creative players in the 2021/22 Premier League. This position has become more common for offensive playmakers to carry out in recent years. For example, Messi typically operates from a wide offensive position. Another trend is to use a wingback as a playmaker. This is illustrated by Tariq Lamptey and Trent Alexander-Arnold, who both rank among the top creative players according to our metric.

Table 5: Player pairs with the highest Creative Decision Rating (CDR) per interaction in the 2021/22 English Premier League season. Only pairs with at least 50 interactions are included.

	Players	Team	CDR
1	Mohamed Salah Sadio Mané	Liverpool	0.0387
2	Harry Kane Heung-Min Son	Tottenham Hotspur	0.0202
3	Leandro Trossard Neal Maupay	Brighton	0.0132
4	Bernardo Silva Kevin De Bruyne	Manchester City	0.0126
5	Kevin De Bruyne Philip Foden	Manchester City	0.0125
6	Marc Cucurella Neal Maupay	Brighton	0.0114
7	Raphinha Rodrigo	Leeds United	0.0114
8	João Cancelo İlkay Gündoğan	Manchester City	0.0096
9	Bruno Fernandes Cristiano Ronaldo	Manchester United	0.0095
10	Jordan Henderson Sadio Mané	Liverpool	0.0079

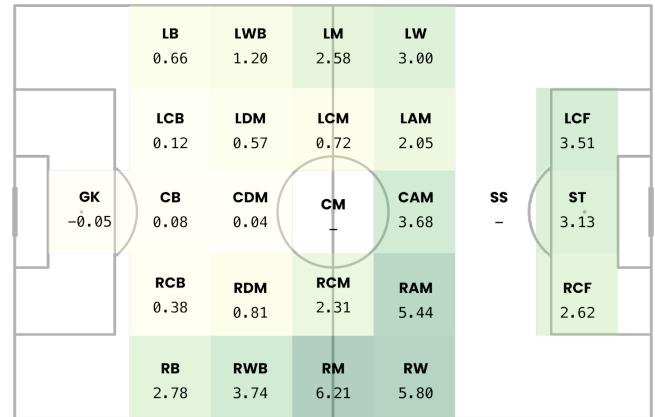


Figure 6: The average creative decision rating per pass, grouped by a player’s position on the field. Darker green colors reflect higher creativity. Insufficient data was available for the CM and SS positions. For a definition of each position, we refer to StatsBomb’s data specification guide.

Interestingly, we found a large imbalance between the left and right wings, with the right wing being significantly more creative. Further research on other leagues should be carried out to determine whether this could be generalized or should be attributed to the player selection in the 2021/22 Premier League season.

5.5 The link between creativity and game state

Both the time remaining and the goal difference have no clear effect on the average CDR (Figure 7). However, the variance in creativity increases as time progresses, particularly in injury time. Possibly, this could be attributed to fatigue and teams taking more defensive risks near the end of the game. This creates extra space, which could be exploited with creative passing. However, as players get more tired, this also gives rise to more “missed” creative opportunities.

6 RELATED WORK

Creativity and its importance in team sports have primarily been investigated in the sports science literature [5, 16, 23]. Specifically for soccer, Kempe and Memmert [16] showed that successful teams use more highly creative actions to score goals. However, the evaluation of creativity in these studies is mostly conducted via psychological assessments [25, 35] or experts rating each action of a player [22, 24]. Obviously, this approach does not scale.

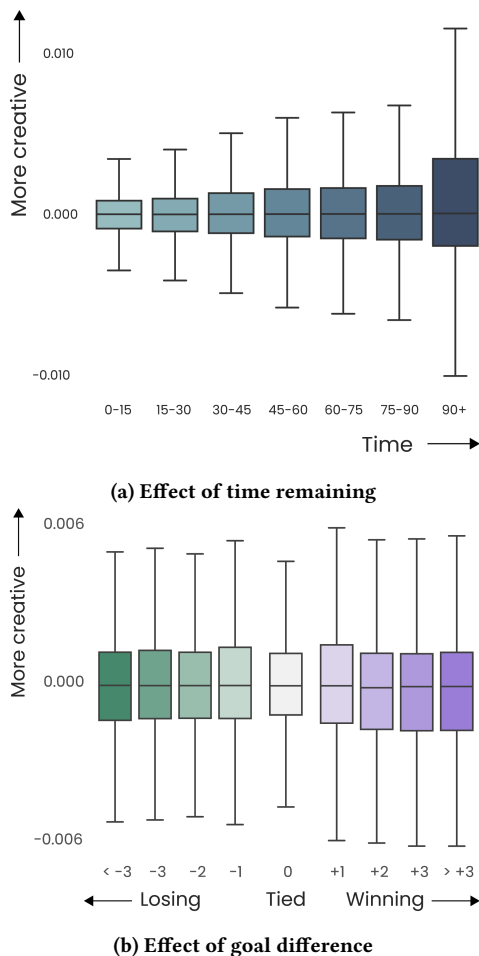


Figure 7: The distribution of the creative decision rating grouped by 15-minute time intervals (top) and goal difference (bottom). A negative goal difference corresponds to the team in possession being behind.

Creativity can be seen as an aspect of decision-making, which is typically evaluated as a trade-off between an action’s risks and rewards [4, 15, 27]. For example, Goes et al. [15] and Power et al. [27] evaluate passes based on how likely they are to be successfully completed to a teammate (i.e., risk), and also by how much they would increase the chance of something good happening (e.g., scoring) if successful (i.e., reward). While our creativity metric builds on these risk-reward frameworks, neither risk nor reward aligns with what one would intuitively label as capturing creativity. Capturing creativity requires adding a notion of originality and making an abstraction of the actual result of the action.

7 CONCLUSIONS

The creativity metric introduced in this paper is a first step towards capturing the intricate concept of creative abilities in soccer players. By utilizing this metric, it becomes feasible to compare players’ creativity both in general and across various scenarios. This can provide clubs and analysts with valuable information during the scouting process.

The metric primarily focuses on quantifying creativity based on the deviation from typical pass choices and the expected value associated with those choices. While it can capture atypical or unexpected actions, it does not explicitly account for the underlying reasons behind those actions, such as habit or personal preferences. This limitation is prevalent in the broader field of soccer analytics, where the lack of player-specific models hampers the ability to disentangle individual player characteristics from team dynamics and tactics. To address this limitation and gain a better understanding of the interplay between habit and creativity, future research could explore methods to incorporate player-specific models. Furthermore, our future plans include extending the creativity metric to other action types, particularly shots and dribbles. Currently, the metric tends to undervalue the creativity of passes in game states where a shot or dribble would be the more typical choice. By extending the pass selection model to an action selection model, this limitation can be overcome, enabling the evaluation of creativity across different action types. Finally, we intend to explore alternative approaches for evaluating the value of typical actions. Although the current method of selecting the most likely pass works well in practice, it may yield inaccurate results when multiple passes are equally probable.

ACKNOWLEDGMENTS

This research received funding from the Flemish Government (AI Research Program), the Research Foundation – Flanders (EOS No. 30992574) and the KU Leuven Research Fund (iBOF/21/075). We thank StatsBomb for providing the data used in this research.

REFERENCES

- [1] Gabriel Anzer and Pascal Bauer. 2022. Expected passes. *Data Mining and Knowledge Discovery* 36, 1 (Jan. 2022), 295–317. <https://doi.org/10.1007/s10618-021-00810-3>
- [2] Ryan Beal, Timothy J. Norman, and Sarvapali D. Ramchurn. 2019. Artificial intelligence for team sports: a survey. *The Knowledge Engineering Review* 34 (2019), e28. <https://doi.org/10.1017/S0269888919000225>
- [3] Lotte Bransen, Pieter Robberechts, Jan Van Haaren, and Jesse Davis. 2019. Choke or Shine? Quantifying Soccer Players’ Abilities to Perform Under Mental Pressure. In *MIT Sloan Sports Analytics Conference*. 1–25.

- [4] Borja Burriel and Javier M. Buldú. 2021. The quest for the right pass: Quantifying player's decision making. In *Proceedings of the StatsBomb Innovation in Football Conference*. StatsBomb, London, United Kingdom, 1–22. <http://statsbomb.com/wp-content/uploads/2021/11/Javier-M-Buldu.pdf>
- [5] Jean-Michel Bénézet and Hansruedi Hasler. 2018. Youth Football Training Manual. <https://digitalhub.fifa.com/m/1b3da6976c9290aa/original/mxpozhrv2gishmxrllpf-pdf.pdf>
- [6] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco, California, USA) (KDD '16). ACM, New York, NY, USA, 785–794. <https://doi.org/10.1145/2939672.2939785>
- [7] Tom Decroos. 2020. Soccer analytics meets artificial intelligence: learning value and style from soccer event stream data. Ph. D. Dissertation. KU Leuven.
- [8] Tom Decroos, Lotte Bransen, Jan Van Haaren, and Jesse Davis. 2019. Actions speak louder than goals: Valuing player actions in soccer. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery and data mining* (KDD '19). ACM, New York, NY, USA, 1851–1861. <https://doi.org/10.1145/3292500.3330758> Place: Anchorage, AK, USA.
- [9] Tom Decroos and Jesse Davis. 2020. Player vectors: Characterizing soccer players' playing style from match event streams. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Springer International Publishing, Cham, 569–584.
- [10] Tom Decroos, Maaïke Van Roy, and Jesse Davis. 2021. SoccerMix: representing soccer actions with mixture models. In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track*. Springer International Publishing, Cham, 459–474.
- [11] Uwe Dick and Ulf Brefeld. 2019. Learning to rate player positioning in soccer. *Big data* 7, 1 (2019), 71–82. <https://doi.org/10.1089/big.2018.0054>
- [12] Javier Fernández and Luke Bornn. 2021. SoccerMap: A Deep Learning Architecture for Visually-Interpretable Analysis in Soccer. In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track*, Yuxiao Dong, Georgiana Ifrim, Dunja Mladenić, Craig Saunders, and Sofie Van Hoecke (Eds.). Springer International Publishing, Cham, 491–506.
- [13] Javier Fernández, Luke Bornn, and Daniel Cervone. 2021. A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. *Machine Learning* 110, 6 (2021), 1389–1427.
- [14] Leon Forcher, Leander Forcher, Sascha Härtel, Darko Jekauc, Hagen Wäsche, Alexander Woll, Timo Gross, and Stefan Altmann. 2022. Does Technical Match Performance in Professional Soccer Depend on the Positional Role or the Individuality of the Player? *Frontiers in Psychology* 13 (2022), 2847. <https://doi.org/10.3389/fpsyg.2022.813206>
- [15] Floris Goes, Edgar Schwarz, Marije Elferink-Gemser, Koen Lemmink, and Michel Brink. 2022. A risk-reward assessment of passing decisions: comparison between positional roles using tracking data from professional men's soccer. *Science and Medicine in Football* 6, 3 (July 2022), 372–380. <https://doi.org/10.1080/24733938.2021.1944660>
- [16] Matthias Kempe and Daniel Memmert. 2018. "Good, better, creative": the influence of creativity on goal scoring in elite soccer. *Journal of Sports Sciences* 36, 21 (Nov. 2018), 2419–2423. <https://doi.org/10.1080/02640414.2018.1459153>
- [17] Matthias Kullowatz. 2020. Goals added: Deep Dive Methodology. American Soccer Analysis. Retrieved May 22, 2023 from <https://www.americansocceranalysis.com/home/2020/5/4/goals-added-deep-dive-methodology>
- [18] Guiliang Liu, Yudong Luo, Oliver Schulte, and Tarak Kharrat. 2020. Deep soccer analytics: learning an action-value function for evaluating soccer players. *Data Mining and Knowledge Discovery* 34, 5 (2020), 1531–1559.
- [19] Sergio Llana, Borja Burriel, Pau Madrero, and Javier Fernández. 2022. Is it worth the effort? Understanding and contextualizing physical metrics in soccer. arXiv:2204.02313 [stat.ML]
- [20] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, New York, NY, USA, 3431–3440. <https://doi.org/10.1109/CVPR.2015.7298965>
- [21] Nils Mackay. 2019. *Introducing a possession value framework*. Stats Perform. Retrieved May 22, 2023 from <https://www.statsperform.com/resource/introducing-a-possession-value-framework>
- [22] Daniel Memmert. 2010. Testing of tactical performance in youth elite soccer. *Journal of Sports Science & Medicine* 9, 2 (2010), 199.
- [23] Daniel Memmert. 2015. *Teaching tactical creativity in sport: Research and practice*. Routledge, Abingdon, UK. <https://doi.org/10.4324/9781315797618>
- [24] Daniel Memmert and Klaus Roth. 2007. The effects of non-specific and specific concepts on tactical creativity in team ball sports. *Journal of Sports Sciences* 25, 12 (Oct. 2007), 1423–1432. <https://doi.org/10.1080/02640410601129755>
- [25] Thomas Morris. 2000. Psychological characteristics and talent identification in soccer. *Journal of Sports Sciences* 18, 9 (Jan. 2000), 715–726. <https://doi.org/10.1080/02640410050120096>
- [26] José M. Oliva-Lozano and José M. Muyor. 2022. Understanding the FIFA quality performance reports for electronic performance and tracking systems: from science to practice. *Science and Medicine in Football* 6, 3 (2022), 398–403. <https://doi.org/10.1080/24733938.2021.1984557> PMID: 35862165.
- [27] Paul Power, Hector Ruiz, Xinyu Wei, and Patrick Lucey. 2017. 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 (KDD '17)*. Association for Computing Machinery, New York, NY, USA, 1605–1613. <https://doi.org/10.1145/3097983.3098051> event-place: Halifax, NS, Canada.
- [28] Pieter Robberechts and Jesse Davis. 2020. How Data Availability Affects the Ability to Learn Good xG Models. In *Workshop on Machine Learning and Data Mining for Sports Analytics at ECML/PKDD 2020 (MLSA)*. Springer International Publishing, Cham, 17–27.
- [29] Sarah Rudd. 2011. A Framework for Tactical Analysis and Individual Offensive Production Assessment in Soccer Using Markov Chains. In *New England Symposium on Statistics in Sports*. <http://nessis.org/nessis11/rudd.pdf>
- [30] Karun Singh. 2019. *Introducing Expected Threat (xT)*. Retrieved May 22, 2023 from <https://karun.in/blog/expected-threat.html>
- [31] StatsBomb. 2020. *Introducing On-Ball Value*. Retrieved May 22, 2023 from <https://statsbomb.com/articles/soccer/introducing-on-ball-value-obv/>
- [32] StatsBomb. 2021. *On Ball Value*. Technical Report.
- [33] Łukasz Szczępański and Ian McHale. 2016. Beyond completion rate: evaluating the passing ability of footballers. *Journal of the Royal Statistical Society. Series A (Statistics in Society)* 179, 2 (2016), 513–533. <http://www.jstor.org/stable/4396554>
- [34] Karl Tuyls, Shayegan Omidshafiei, Paul Muller, Zhe Wang, Jerome Connor, Daniel Hennes, Ian Graham, William Spearman, Tim Waskett, Dafydd Steel, et al. 2021. Game Plan: What AI can do for Football, and What Football can do for AI. *Journal of Artificial Intelligence Research* 71 (2021), 41–88.
- [35] Torbjörn Vestberg, Roland Gustafson, Liselotte Maurex, Martin Ingvar, and Predrag Petrovic. 2012. Executive Functions Predict the Success of Top-Soccer Players. *PLoS one* 7, 4 (April 2012), e34731. <https://doi.org/10.1371/journal.pone.0034731>
- [36] Derek Yam. 2019. *Attacking Contributions: Markov Models for Football*. StatsBomb. Retrieved May 22, 2023 from <https://statsbomb.com/articles/soccer/attacking-contributions-markov-models-for-football/>