

“Not All Passes Are Created Equal:” Objectively Measuring the Risk and Reward of Passes in Soccer from Tracking Data

Paul Power
STATS
Leeds, U.K
paul.power@stats.com

Xinyu Wei
STATS
Sydney, Australia
xwei@stats.com

Hector Ruiz
STATS
Barcelona, Spain
hruiz@stats.com

Patrick Lucey
STATS
Chicago, USA
plucey@stats.com

ABSTRACT

In soccer, the most frequent event that occurs is a pass. For a trained eye, there are a myriad of adjectives which could describe this event (e.g., “majestic pass”, “conservative” to “poor-ball”). However, as these events are needed to be coded live and in real-time (most often by human annotators), the current method of grading passes is restricted to the binary labels 0 (unsuccessful) or 1 (successful). Obviously, this is sub-optimal because the quality of a pass needs to be measured on a continuous spectrum (i.e., $0 \rightarrow 100\%$) and not a binary value. Additionally, a pass can be measured across multiple dimensions, namely: i) *risk* – the likelihood of executing a pass in a given situation, and ii) *reward* – the likelihood of a pass creating a chance. In this paper, we show how we estimate both the risk and reward of a pass across two seasons of tracking data captured from a recent professional soccer league with state-of-the-art performance, then showcase various use cases of our deployed passing system.

CCS CONCEPTS

•Computer systems organization → Embedded systems; Redundancy; Robotics; •Networks → Network reliability;

KEYWORDS

Unsupervised; Clustering; Soccer; Passing; Tracking Data;

1 INTRODUCTION

Pep Guardiola once stated, “You have to pass the ball with a clear intention, with the aim of making it into the opposition’s goal. It’s not about passing for the sake of it” [12]. For attack minded managers like Guardiola - the ultimate *reward* of playing is to create and exploit dangerous situations by effectively passing the ball to create imbalances between the attack and defense. For other




Figure 1: Current measures in soccer analytics assign the same amount of credit for the pass of Matic to Fabregas (left), and Matic to Costa (right), even though the latter is more likely to lead to a chance on goal (higher reward). In this paper, we show how we objectively measure both risk and reward from data as well as showing applications using these measures.

managers, due to their playing roster and/or strategic mindset, their goal is to minimize *risk* when they have possession and pray on the mistakes of their opponents [13].

Even though both measuring the risk and reward of a pass in soccer would be clearly useful, existing soccer analytic measures currently capture neither. For example in Figure 1, we show a snapshot of play where the player with the ball (Matic) passes it to Fabregas (left) and Costa (right). Looking at both situations it is obvious that the pass to Costa is more likely to lead to a shooting chance (i.e., higher reward) but is inherently more risky and requires more skill to execute than the pass to Fabregas. However, current passing measures assign both passes with the same weighting (1 for successfully making the pass and 0 for not). This is a problem in terms of team and player analysis, as the current defacto passing-metric is pass completion. As such, even though the pass to Costa maybe better for the team, the high likelihood of the pass being intercepted would normally result in the player and team’s metric suffering in addition to not adequately reflecting the game situation.

A better alternative would be to capture the likelihood of a player executing a pass (something that we describe as the risk of a pass) as well as the likelihood that the pass will create a chance on goal (something that we describe as the reward of a pass). Figure 2 showcases an example where such a metric could be useful in analyzing the quality of a pass given the difficulty of each passing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](http://permissions.acm.org).

KDD'17, August 13–17, 2017, Halifax, NS, Canada.

© 2017 ACM. 978-1-4503-4887-4/17/08...\$15.00

DOI: <http://dx.doi.org/10.1145/3097983.3098051>

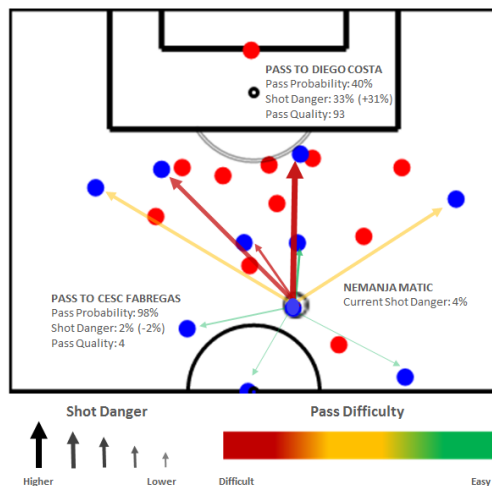


Figure 2: Shows the pass with the least amount of “risk” (i.e., likelihood of successfully executing a pass) – Matic to to Fabregas, as well as the pass with the highest “reward” (i.e., likelihood of a pass creating a chance on goal) – Matic to Costa.

option and the probability of the pass leading to a shot in a given situation. We can now clearly capture that the pass to Costa is both more difficult to complete (40% vs 98%) and also increases the probability of a shot occurring compared to the pass to Fabregas (+31% vs -2%). In this paper, we showcase how we can estimate both the risk and reward of a pass in an objective manner.

At a high-level, there are two potential ways which we could first estimate the risk of a pass: i) *human-labeled* – where a human expert assigns a quality rating to each pass (e.g. bad/average/good or 0-10) and we train a classifier to emulate that expert (or host of experts), or ii) *event-outcome-labeled* – where we optimize an objective function which uses the outcome of pass event (e.g., 0 = unsuccessful pass, 1 = successful pass).

While both approaches are valid, the event-outcome labeled approach is preferred as the human-labeled approach is still inherently subjective (even at the expert level, consensus on what a good pass is, is still ambiguous). This is an important point, as if we are deploying this at scale across many teams the idea of “the oracle of passing” is unsatisfying. A more satisfying explanation is to say that our passing quality measure correlates with the likelihood of executing a pass. Additionally, this approach allows us to start to measure the reward of the pass by correlating pass reward with other concrete events such as shot on goal.

In this paper, we show how we can objectively measure the risk and reward of a pass using player tracking and ball event data, as well as applications of our deployed system across a host of use cases.

2 RELATED WORK

Although there is yet to be a large publicly available dataset for researchers to compare ideas and methods on, some data has become available through vendors, which has facilitated with recent advances in soccer analysis. The majority of these works have

solely focused on using ball-event data, again because this has been the most common form of data. As such, some really good analysis of passes and team-play have occurred. Notable works including Lucey et al., [8, 9], which used an aggregated method of describing passes to identify the characteristic playing patterns of teams. Gyramati et al. [4] and Wang et al. [15] conducted similar analysis. More recently Brooks et al., [2] describe a novel player ranking system based entirely on the value of passes completed.

In terms of directly measuring the passing quality through the use of a supervised learning methodology, Horton et al., used domain experts to create ground truth labels of good, OK and bad for passes in order to train a supervised learning model to classify passing quality [6]. The authors took inspiration from Taki and Hasegawa [14] work which looked to measure the dominant region of a player in order to model the probability of a players region intercepting that of the ball. The authors highlighted the issue of having conflicting agreement between the experts creating the training labels. The same level of agreement was also captured through the model as the experts however, the paper helped to highlight the issue of relying upon expert opinion to create “the oracle of passing”. Link et al., [7] hand crafted a set of features using player and ball spatial-temporal data; pressure, density, zone and control to measure the dangerousness of a pass in the attacking 3rd of the pitch. McHale et al., [10] used a generalized additive mixed model to measure the probability of completing a pass whilst attempting to control for the random effect of the team and player. Gyramati et al., [5] recently proposed a “QPass” method to quantify the quality of a pass. Outside of soccer, Cervone et al., [3] proposed a model to capture an ‘expected possession value’ to measure how a players decision increased or decreased the teams chance of scoring a point in basketball.

While these papers have helped to advance the techniques used to measure the quality and effect of a pass, nearly all of these methods do not take into account both the game and team context – mainly due to the absence of tracking data which captures these important contextual cues. We believe our paper advances the understanding of the game because of this added information.

3 MEASURING PASS RISK AND REWARD

3.1 Data Inventory

In this paper, we used ball event data and player tracking data from the English Premier League games between 2014/15-2015/16 seasons totaling 726 matches. Each match contains the trajectories of each player including; X and Y location sampled at 10hz; time stamp; player name; team name and match name. Event data files include; the event name (e.g. pass, shot, tackle); ball X and Y position for the origin and destination; ball time stamp; team and player in possession; match name and identity of the opposing player involved in any duel situation such as a tackle or header.

In terms of passes within the sample for the two full seasons, 571,287 passes were attempted with 468,265 being successful with an average of 380.46 being attempted and 320.91 completed (84.35%) per game. Table 1 provides a full summary of types of passes. As expected, it can be seen that backward and sideways passing are much more likely to be executed than a forwards pass (95.78%, 93.03% vs 72.28%). Additionally, as we get closer to the goal, it can

Pass Type	Total	Successful	Percentage
All Passes	580,980	492,337	84.74%
Backwards	133,763	127,889	95.61%
Sideways	209,190	194,234	92.85%
Forwards	238,027	170,214	71.51%
Short	163,193	133,947	82.07%
Medium	313,693	266,695	85.02%
Long	104,094	67,623	64.94%
First Time	172,163	140,449	81.58%
Final Third	86,365	49,793	57.65%
Penalty Area Entry	47,453	17,175	36.19%

Table 1: Summary of high level derived events for 2013-14 to 2015-16 seasons

be seen that passing gets riskier with passes in the forward third being successfully made 58% of the time, and passes into the penalty area being made 37% of the time. From this, it is obvious that we need to capture the “context” of the pass to best measure the risk of each pass.

3.2 Defining Pass Risk

In this paper, we define the “risk” of a pass as *the likelihood that the player will successfully make the pass given a player has possession of the ball and has the ability to pass it in a given situation*. Based on this definition, the risk can therefore be estimated utilizing a standard supervised learning pipeline, where given some input features describing the game situation, we can train a classifier to yield a probability between 0 and 1.

To see the value of different feature representations, we first established the following baselines: i) *Naive*: we assign the average pass completion to all passes (i.e., 85%), ii) *Ball-Information*: we use the starting and end position of the pass, and iii) *Tracking/Feature-Crafted*: In addition to the ball-information, we utilize the player trajectory data to craft soccer specific features which we define as *Micro Features*. These features aim to capture critical coaching points that coaches feel influence the chance of a pass being completed and model the relationship between the main protagonists involved during the pass. These features include:

- (1) *Speed of the player in possession and the intended receiver*: This feature is thought to have an impact on a player’s control over the ball when they are passing and trying to receive a pass.
- (2) *Speed of the nearest defender toward the passer and the receiver*: This provides an indication of the level of pressure that is being applied to the passer and receiver. Pressure is a key factor in forcing player’s either into an error or to move the ball away from dangerous areas.
- (3) *Distance of nearest defender to the passer and receiver*: This acts as a proxy for whether a player is marked by a defender, which would make the skill of executing and receiving a pass more difficult due to the attentions of a defender. The closer the defender is the more skill required to execute the pass and keep possession.




Figure 3: Example of Expected Receiver Membership: P1 has two options to pass to P2 and P3, however the defender (D1) has intercepted the pass. P1 in this case is identified as target player as he is closest to where the ball was intercepted and has the smallest angle to the line of the pass.

- (4) *Nearest defender angle to the passing line*: As well as pressing the passer and receiver, the quality of the passing channel between the two player’s is a key factor in successful passes.
- (5) *First time pass*: We consider the effect of playing a pass on the first touch or not. This is considered a major skill that differentiates between elite and average passers.
- (6) *Time from regaining possession*: A pass coming quickly after winning the ball in the attacking third is more likely to see the opposition in an unorganized state compared to pass after a long possession in your own half.

The features we crafted build on the work of McHale et al. An important additional feature we included was that of *intended receiver*. This is key to identify in order to narrow down the potential search space of all possible passes. Previously, McHale et al [10] create a simple distance membership to see which attacking player was closest to the intercepted or incomplete pass. To improve this initial model we include the angle of the potential receiver to the passing line. The expected receiver membership can therefore be determined as:

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

The player with the highest membership (0 - 1) was identified as the most likely player to receive the pass (Figure 3). As a result the end x, y location was replaced with the expected receivers coordinates. It is recognized that there are several limitations with this simple model in situations where two players are close to each other or the ball is intercepted early in its flight path. In addition a further limitation of the model is that we have not predicted where the ball may have been passed to and just used the absolute position of the ‘expected receiver’. For example a pass may be intended to played in front of the attacker to run onto.

In terms of classifiers used, we deployed a logistic regressor. The reason why we chose a logistic regressor over a non linear black box model, such as a random decision forest, is we wanted the model to be interpretable. This is an import feature of our model as it is vital that a coach can understand what is driving

the passing performance of a player. As such a coefficient can be thought of as a coaching point. By using a linear model, we enforce the linearization of the data to be done in the feature space (see next subsection). To train the classifiers we used 352,466 examples, and tuned the parameters on our evaluation set which consisted of 114,257 examples. We tested the various feature sets on 114,257 examples. In terms of labels, all successfully executed passes were labeled as positive examples, and all unsuccessful passes were the negative examples.

3.3 Defining Pass Reward

We define the “reward” of a pass as *given a player has possession of the ball and has the ability to pass it in a given situation, pass reward is estimating the likelihood that the pass made will result in a shot within the next 10 seconds*. We used a similar supervised learning approach but instead of using the labels of successful and unsuccessful passes, we used the labels of whether a shot occurred within 10 seconds of the pass or not. We chose a 10 second window due it being previously used to analyze shot outcome in recent literature [11] and is heuristically a period of time used by coaches in practice.

As shots occur very sparsely, an obvious drawback is the imbalance in number of positive/negative examples in training the classifier. In our training set we had 7427 positive and 136180 negative examples. In our test set we had a split of 3062 positive and 58484 negative examples. The performance of the classifiers are shown in the right columns of Table 2, and it can be seen that using the tracking data obtains the best performance.

4 INCORPORATING CONTEXT

4.1 Learning Context Directly from Data

As shown in Table 1 there is tremendous variation in the passing execution rates depending on the context on where the pass is taken (e.g., backwards pass ($\approx 96\%$) vs pen-area entry ($\approx 37\%$)). Even though the features crafted from the tracking data capture the context at a micro-level, having high-level contextual information at the team level may improve prediction – as well as provide useful contextual information for the coaches.

As soccer is a highly strategic sport, capturing the strategic features of teams is important in contextualizing analysis. As such we can think of our approach to creating a passing dictionary as a set of leaf nodes within a tree structure (Figure 5). In the previous section we have done a reasonable job of capturing the microfeatures from the raw data. We need to now capture the

tactical features (i.e., game-state) as well as the formation features (i.e., team-structure).

As we are only interested in analyzing passes, we focused solely on open-play scenarios. In soccer coaching analysis, three distinct game-states are used for analysis: i) build-up, ii) counter-attack, iii) unstructured-play. Not only does this approach enable the contextualization of analysis, it also *linearizes* the data which should improve our overall estimations of risk and reward. To test out this method, we compared our previous models described in the previous section to a model which incorporated these contextual features. Initially, we only tested using the tactical features (tracking + tactics). As can be seen in the bottom half of Table 2 we show that the overall prediction of risk and reward improved.

To obtain further contextual information, we then captured formation features - which is essentially getting an indication of the location and spread of the players within the team structure. To do this, we employed a formation clustering method described by Bialkowski et al., [1].

By aligning players to a specific role it allows use to find similar situations by using the XY positions of all players to describe the formation structure for both the attacking and defending teams. This is important as we can now capture if a pass is being played between the lines of midfielders and defenders or if a pass is attempting to break the final line of defenders. As we have aligned each player to a role we can teach the model the spatial relationship for all players in relation to the goal by calculating the polar coordinates of all player to the goal being attacked. Finally as we want to understand how different roles impact the chances of completing passes we calculate the polar coordinates of the defenders to the passer and receiver.

A key tactical influence on the type of pass attempted and therefore the risk of completing a pass is the defensive block. A defensive block can be split into high-block, medium-block, and low-block. In order to generate a label for this we performed kmeans clustering on the aligned player and ball data. Employing this approach yielded better risk and reward estimation (tracking + formation).

Examples of our tactical and formation features can be show in Figure 4. In the first example, we show a typical counter-attack where the ball has been regained by the blue team’s goal-keeper and quickly moved forward before the defense can reorganize. We can see by the player traces that the possession started in a high block situation, we also see that due to the defenders being out of position the average passing risk was low while the average pass reward was high due to their progressive nature. In the second example, shows a possession in the build-up phase in a mid-block where the red team has passed the ball laterally to try and unlock a highly organized defense. Due to the lateral nature of the passes the average pass danger is lower. However, because the number of defenders in front of the ball is higher than our counter attack example, the skill level required to complete a more penetrative pass is larger.

Features	Pass Risk		Pass Reward	
	Log-Loss	RMSE	Log-Loss	RMSE
Naive	0.4317	0.3621	0.1977	0.2174
Ball-Only	0.3623	0.3306	0.1771	0.2185
Tracking	0.3268	0.3194	0.1566	0.2045
Tracking + Tactics	0.2918	0.2960	0.1560	0.2420
Tracking + Formation	0.2125	0.2438	0.1391	0.1939

Table 2: Results of the various features used to describe pass risk.

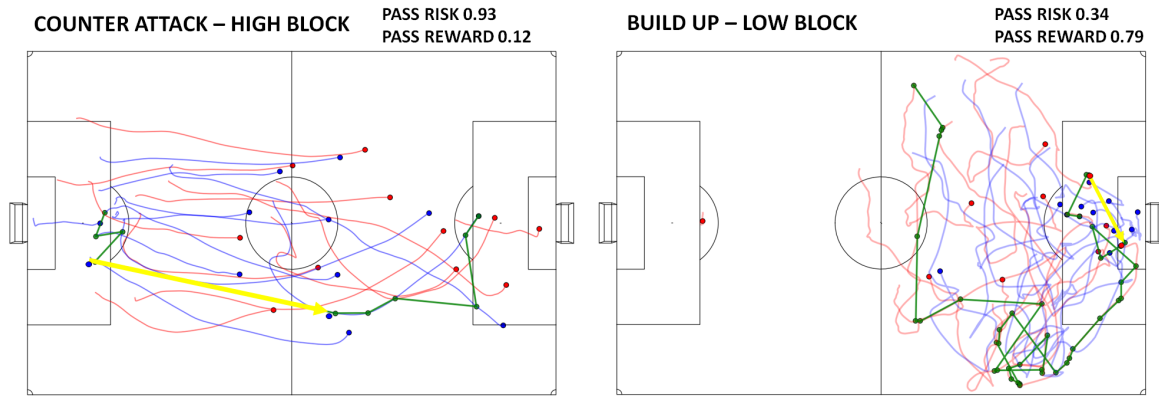


Figure 4: Examples of possessions highlighting different match contexts: (Left) Counter-Attack - High Block, and (Right) Build-Up - Low Block. The yellow line represents the Pass Risk and Pass Reward for each situation.

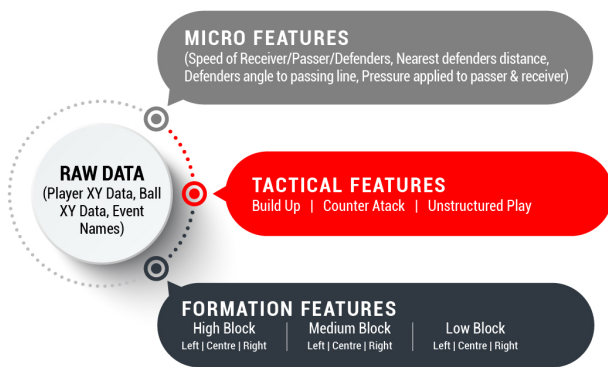


Figure 5: Our tree based method to add contextual features for each model.

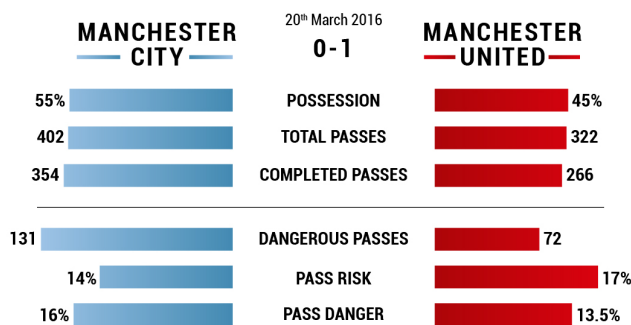


Figure 6: Above the line: An example of the standard stats provided shows a basic summary of passing stats. Below the line: An example of advanced passing stats.

5 APPLICATION I: GAME ANALYSIS

5.1 Match Analysis

Existing match summaries fail to provide a true insight into the strategic concepts used by teams and how well they were able to execute them. Take Figure 6 for example, which illustrates the typical post-game summary of passing performance of two teams. Using the basic statistics, we can see that Manchester City dominated possession (55% vs 45%), and had more passes (402 vs 266) and a higher completion rate (88% vs 82%). However, Manchester United actually won the game 1-0. This begs the question, *were Manchester United lucky or were they more effective in their possession compared to City?* By using our pass risk and reward models we get a more revealing picture.

The two obvious measures to include are that of average **pass risk** and **pass danger** – the latter is a synonym for pass reward, with higher values corresponding with more passes the team has made that leads to potential shots on goal. Additionally, we define **dangerous pass** which, is a pass that is in the top 25th percentile of passes with the highest reward. This threshold is determined from the training/evaluation set and not within game (i.e., it is a fixed threshold for all games). Using these values, we can see that although Manchester City played nearly twice as many dangerous passes than Manchester United (131 v 72), their passes were generally less risky (14% v 17%) and more dangerous (16% v 13.5%). From these measures, we can now get a sense of how the game is being played.

5.2 Specific Play Analysis

The ability to play the *critical pass* that unlocks a defense is one of the most highly sort after skills in soccer. Currently, these passes are manually defined during the game by a human judge, which is highly subjective and variable. Sometimes these passes are extremely obvious (e.g. a pass that leads directly to a shot) but at other times they could be the third or fourth previous pass that was the critical moment in the move. By assessing the reward of each pass during a play, we can objectively assess who is responsible for changing the attacking momentum in a possession. Figure 8 provides an example of a such a play. In the time line it

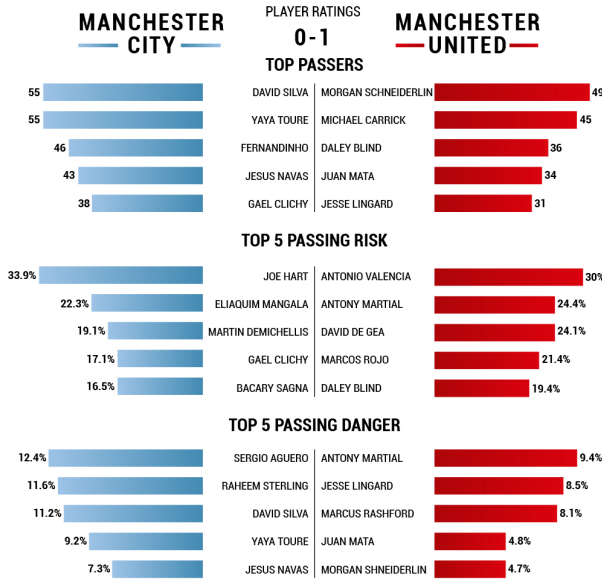


Figure 7: Top 5 players with the most risky and rewarding (dangerous) passes.




Figure 8: Example of a play with timeline underneath that describes when the critical pass was made during the move – the critical pass (or pass with the highest reward) was the one Fabregas played to Willian 15 seconds before the actual last pass was made.

can be seen that even though Willian makes the final assist (*key pass*), we can see that Fabregas played the *critical pass* during the move, increasing the probability of a shot occurring from 0.13 to 0.21. Willian's ability to maintain the ball in a dangerous area is

vital but we can now identify and credit Fabregas as playing the most critical pass in the move. As such we are now able to not just capture the outcome of a possession but the process that led to end outcome. By combining the passing risk and passing reward models we can now learn a new dictionary of objectively measured dangerous or critical passes.

6 APPLICATION II: SEASON-WIDE PLAYER ANALYSIS

6.1 Ranking the Riskiest Players

By modeling the risk associated with completing a pass given a specific context, we can measure the average skill required to complete a pass – our pass risk prediction estimates what the average player will do in that situation. The intuition is that a player with more talent would be more likely to execute a risky pass compared to a player with less skill. To do this we create a new statistic called **Passing Plus Minus (PPM)**. Given that we have a pass risk prediction for each pass, y_{risk} , we can simply calculate the PPM as:

$$\text{Passing Plus/Minus} = \sum_{s=1}^S (1 - y_{risk}^s) - \sum_{u=1}^U (y_{risk}^u - 1) \quad (2)$$

where S and U are the number of successful and unsuccessful passes. For example, if a player completes a pass with $y_{risk}^s = 0.9$ they are awarded a credit of +0.1 but if the pass is unsuccessful they are penalized -0.9. PPM is simply the difference of the aggregated credits and penalties and is normalized per 90 minutes (i.e., the length of a game) to enable comparisons between players. This is important if a player has played only a portion of the game. PPM allows coaches/analysts to quickly assess which players are completing more passes than an average player (*positive score*) and who is completing less passes than an average player (*negative score*). A player with a score of 0 can be thought of as an average player.

The second new metric we introduce is **Difficult Pass Completion (DP%)** which measures both how many high risk passes a player makes and completes. A difficult pass is defined as a pass whose probability is in the 75th percentile of the most high risk passes (the least likely to be completed). We calculate DP% as follows:

$$\text{Difficult Pass Completion} = \frac{\sum_{i=1}^n i = \text{DPS}}{\sum_{i=1}^n i = \text{DPA}} \quad (3)$$

where DPS is the number of difficult passes completed and DPA is the number of difficult passes attempted. This measure shows which players are able to attempt and complete the most risky/difficult passes.

Figure 9 (left) shows the distribution of passing ability for all outfield players who played over 1800 minutes (20 games) in the English Premier League for the 2015-2016 season. The X-axis shows PPM, with the Y-axis showing DP%. We can see that there is a positive correlation between the two metrics which intuitively makes sense as players who are completing more passes than expected should be completing more difficult passes than average. Unsurprisingly creative players such as Cesc Fabregas (+1.19/90%), Mesut Ozil (+0.96/95%), and Eden Hazard (+0.79/65%) feature as players

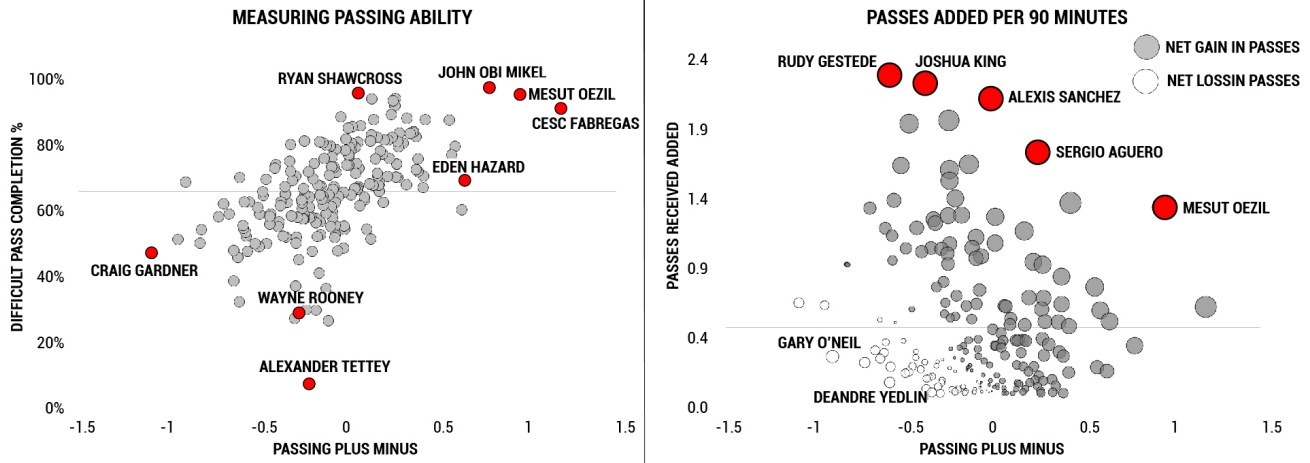


Figure 9: Left: A comparison between a players ability to complete more passes than expected (X-axis) and execute the most difficult passes (Y-axis). Right: A comparison between a players ability to complete more passes than expected (X-axis) and receive the more difficult passes than expected (Y-axis). Axis line represents the the league average.

with the highest PPM and DP%. The worst passer in the league that season was Craig Gardiner who actual cost his team -1.09 passes per game and had one of the lowest DP% (47%) rates. A surprising discovery is that Wayne Rooney (-0.27/45%) is one of the lowest rated players in both metrics. Wayne Rooney has been one of England’s and Manchester United’s best players in recent history however, the media has been reporting that Rooney’s abilities have been in decline. Our data supports this opinion, demonstrating that Rooney not only completed fewer passes than an expected but also, has the twelfth lowest DP% in the league.

6.2 Ranking of Best Players Receiving Passes

In order to complete a high risk pass the passer needs to have someone equally skilled to receive the pass. To measure this we create a new metric called **Passes Received Added (PRA)**. We only consider the passes that have been received by a player that fall within the 75th percentile of passing risk (the most difficult to complete). As such Passes Received Added can be calculated as:

$$\text{Passes Received Added} = 1 - XD_{pr} \tag{4}$$

where XD_{pr} is the probability of a pass being completed. For example, if a player receives a pass with a 0.4 probability of being completed they are awarded 0.6. If the pass was weighted at 0.2 the receiver would be awarded 0.8. We then sum these scores and standardize them to a per 90 minute value. A player who has a score of 0, receives as many dangerous passes as expected while a player with a score greater than 0 receives more dangerous passes than expected. Figure 9 (right) shows PRA on the Y-axis compared to Passing Plus Minus on the X-axis. Again, only outfield players who played more than 1800 minutes are included in the sample. We now see a combination of tall target players such as Rudi Gusted and Joshua King and small quick center forwards who play in highly skilled teams such as Alexi Sanchez and Sergio Aguero rank at

the top of this metric. By combining PMM and PRA we get **Total Passes Added (TPA)**:

$$\text{Total Passes Added} = \text{PPM} + \text{PRA} \tag{5}$$

TPA shows which players help their team keep or lose the ball more than an average player. Players with a negative score lose the ball more often, while players with a positive score help keep the ball more than an average player. TPA represents the z value (bubble size) for figure 9 with gray circles representing players who add passes and white circles showing players who lose passes. By combining these three metrics We can see not just which players add and lose passes but how they do this. For example we can see that Sergio Aguero and Mesut Oezil complete and receive more passes than average, while Rudi Gested and Joshua King complete fewer passes than an average player but receive more passes than average. Conversely we see that Gary O’Neil and DeAndre Yedlin both fail to complete and receive more passes than an average player resulting in them having a negative effect on their team.

6.3 Ranking of Reward Passes

The ability to execute difficult passes is a key quality to assess in players, however we also want to be able to assess the impact the pass will have in creating danger for the opposition. Our pass reward model measures the likelihood of a shot being created in the next 10 seconds and allows us to assess which passes are the most dangerous. To capture the highest reward passes we take the 75th percentile of dangerous passes (> 6% likelihood). We therefore define a **Dangerous Pass (DP)** as an attempted pass that has a great than 6% chance of leading to a shot in the next 10 seconds. To create the Dangerous Passes metric we sum all the passes attempted by a player within this threshold. We then standardize DP to 90 minutes to again allow easy comparison.

As with our PRA metric we can now assess which players not only make dangerous passes but can also receiver them. To measure




Figure 10: A comparison between players who receive and make dangerous passes. Bubble size equals total dangerous passes and received. Players below the line make more dangerous passes. Players above receive more dangerous passes.

this we create a new metric called **Receiving Dangerous Passes (RDP)**. As with DP we only consider the passes a player receives that have a greater than 6% likelihood of leading to a shot in the next 10 seconds. As such we just simply count the number of these dangerous passes received and normalize for 90 minutes. We are now able to assess which players create and exploit the most dangerous passes.

Figure 10 shows the distribution of all outfield players who played more than 1800 minutes for the season. We see a different ordering compared to the players who had high scores for PPM and TPA. Based on the danger of a pass (X axis) David Silva and Mesut Ozil are the most dangerous players in possession of the ball. As we would expect, attacking players make up the majority of players executing high reward passes, however we also see Alexander Kolarov, a full back, feature high in the metric. Another interesting discovery is that Wesley Hoolahan is ranked 19th for dangerous passes made. This is interesting as Norwich City were relegated in this season yet had one of the most dangerous players in the league. When looking at players who receive the most dangerous passes (Y axis), we also see Alexander Kolarov features again indicating his ability to not only make dangerous passes but to receive them as well. This is critical information for an opposing coach who may not automatically assess a full back such as Kolarov as being a main threat to stifle.

In addition to analyzing future opponents, we can also use these new metrics to better assess a player a team may be looking to recruit. For example a team with a low budget may not have been aware how effective Wesley Hoolahan was in playing dangerous passes and could consequently 'beat the market' in signing such a player from a team who has dropped out of the league for a lower transfer fee.

7 APPLICATION III: TEAM-BASED ANALYSIS

With teams playing two games a week a critical element of the analysis process for a coach is to quickly find patterns a team will use and understand how dangerous these patterns are. Given that we can now assign a risk and reward rating for each pass and capture the tactical context of the pass, we can now measure the passing style of a team. To do this, we segmented the the most dangerous passes (greater than 6% likelihood of leading to a shot), and applied k-means clustering on the XY coordinates of the origin and destination of the pass. It was decided that 16 clusters was optimal based on the coaches identifying that the clusters identified provided the optimal amount of information without being overloaded with information. Figure 11 (right) shows the 16 cluster centroids with the color reflecting the average reward.

We can immediately see that the most dangerous passes occur around the edge of the penalty area (cluster 1,2 and 3). While these passes have the highest reward they are also have the highest risk, requiring high levels of skill (figure 11). Interestingly, when examining the centroids, we can see that the average position of the pass types are not symmetrical with cluster 4 being more dangerous than cluster 9 (19% vs 10%) for example. We can also see that the least riskiest pass (cluster 10) has an average risk of 3% and an average reward of 12% which is higher than six other clusters. Passes from clusters including 8, 10 and 16 may be a better option to use for less skilled teams as there is a lower risk but still relatively high reward.

Figure 11 visualizes the passing tendencies of all teams in the league via a Hinton diagram. Each column is standardized to the league with the size of the square indicating a team makes more passes while the intensity of the color shows how dangerous a team is on average when using that pass. In this case the redder a square the more dangerous a team is. Immediately, it can be seen which teams dominate possession with the like of Manchester City and Arsenal having consistently larger squares. While certain teams may have more possession we can also see if they are effective in using these passes. Arsenal for example have the most passes in clusters 2 and 3 yet Leicester City (Cluster 2) and West Ham United (Cluster 3) are more dangerous. Looking at Leicester City in more detail we can see that they had a highly effective mix of combining high risk and high reward passes (Clusters 1 and 2) and low risk medium reward passes (Cluster 6 and 12).

The ability to quickly find how an opponent plays and where they are most dangerous is a key break through in the analysis of passing.

8 SUMMARY

In this paper, we presented an objective method of estimating the risk and reward of all passes using a supervised learning approach. We showed that adding contextual features improved the prediction performance in addition to giving semantic information to each pass. We then showcased four applications that covered how these tools can be used to describe individual match and play analysis, in addition to player and team analysis. Not only can it be used to describe the offensive behavior of players and teams, it can also be used to describe the defensive behavior as well.

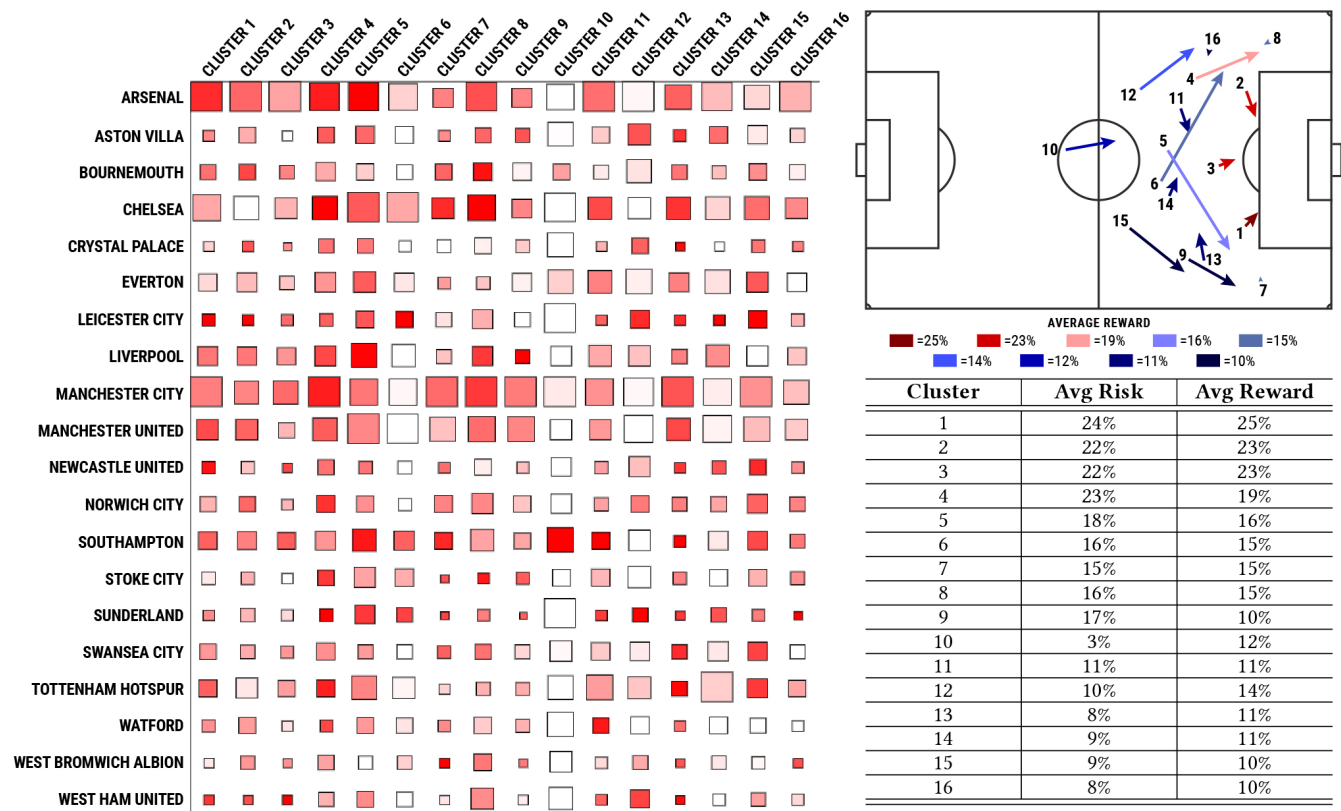


Figure 11: Hinton Diagram depicting how often team use a type of pass and the average reward of each pass. The larger the square the more number of passes a team uses, and the more intense the color corresponds to the higher average reward of that pass.

REFERENCES

[1] A. Bialkowski, P. Lucey, P. Carr, Y. Yue, S. Sridharan, and year=2014 I. Matthews, booktitle=ICDM. Large-scale analysis of soccer matches using spatiotemporal tracking data.

[2] J. Brooks, M. Kerr, and J. Gutag. 2016. Developing a Data-Driven Player Ranking in Soccer Using Predictive Model Weights. In *KDD*.

[3] Daniel Cervone, Alex DfiAmour, Luke Bornn, and Kirk Goldsberry. 2016. A Multiresolution Stochastic Process Model for Predicting Basketball Possession Outcomes. *J. Amer. Statist. Assoc.* 111, 514 (2016), 585–599.

[4] L. Gyarmati and X. Anguera. 2015. Automatic Extraction of the Passing Strategies of Soccer Teams. In *KDD Workshop on Large-Scale Sports Analytics*.

[5] L. Gyarmati and R. Stanojevic. 2016. QPass: a Merit-based Evaluation of Soccer Passes. In *KDD Workshop on Large-Scale Sports Analytics*.

[6] M. Horton, J. Gudmundsson, S. Chawla, and J. Estephan. 2014. Classification of passes in football matches using spatiotemporal data. In *KDD Workshop on Large-Scale Sports Analytics*.

[7] D. Link, S. Lang, and P. Seidenschwarz. 2016. Real Time Quantification of Dangerousness in Football Using Spatiotemporal Tracking Data. *PLoS ONE* (2016).

[8] P. Lucey, A. Bialkowski, P. Carr, E. Foote, and I. Matthews. 2012. Characterizing Multi-Agent Team Behavior from Partial Team Tracings: Evidence from the English Premier League.

[9] P. Lucey, D. Oliver, P. Carr, J. Roth, and I. Matthews. 2013. Assessing team strategy using spatiotemporal data. In *KDD*.

[10] I McHale and L Szczepanski. 2015. Beyond completion rate: evaluating passing ability of footballers. *Royal Statistical Society, Series A.* (2015).

[11] M. Monfort P. Carr P. Lucey, A. Bialkowski and I. Matthews. 2015. Quality vs Quantity: Improved shot prediction in soccer using strategic features from spatiotemporal data. In *MIT Sloan Sports Analytics Conference*.

[12] Marti Perarnau. 2014. *Pep Confidential: Inside Pep Guardiola’s First Season at Bayern Munich*. Birlinn.

[13] J. Skyes and N. Paine. 2016. How One Man’s Bad Math Helped Ruin Decades Of English Soccer. (2016). <https://fivethirtyeight.com/features/how-one-mans-bad-math-helped-ruin-decades-of-english-soccer/>

[14] T. Taki and J. Hasegawa. 2000. Visualization of dominant region in team games and its application to teamwork analysis. In *Computer Graphics International*. 227–235.

[15] Qing Wang, Hengshu Zhu, Wei Hu, Zhiyong Shen, and Yuan Yao. 2015. Discerning Tactical Patterns for Professional Soccer Teams: An Enhanced Topic Model with Applications. In *KDD*.