14TH ANNUAL
MIT SLOAN
SPORTS ANALYTICS CONFERENCE
MARCH 6 - 7, 2020 | BOSTON, MA
presented by ESPN

The right place at the right time: Advanced off-ball metrics for exploiting an opponent's spatial weaknesses in soccer

Sergio Llana, Pau Madrero and Javier Fernández
FC Barcelona

1. Introduction

What should we do to win the next match? This is the most important question a coach can ask to a game analyst. The answer to this question is, however, much more complex to solve than to formulate. In the recent years, the fast-growing field of soccer analytics has been able to provide novel approaches for solving more specific aspects of the game such as evaluating the quality of observed shots and passes [1,2,3,4], estimating the expected value of potential unobserved actions [5], measuring space control and creation [6,7,8], and even assessing the mental pressure of players according to match demands [9]. While these approaches address isolated aspects of the game, most can serve as useful building blocks for solving the more demanding practical questions from coaches in a more comprehensive way. Here, we focus on a yet-little explored area in soccer analytics, opponent scouting, and more specifically, we address the question of how to identify opponent's defensive weaknesses and how to exploit them in order to gain a competitive advantage in the game.

We introduce the concept of off-ball advantage, a novel approach to identify when a player controls a valuable space in the field in such a way that passing to that player would create a considerable increase in the expected long-term outcome of the possession. While off-ball performance analysis has been recently addressed in soccer [4,7,10] as well as in many other team sports [11,12,13], most of what has been done to date does not relate player's movements with team's performance.

We move a step beyond from previous approaches on space control models [7,8] by associating spatial advantages with offensive and defensive team performance. Following our approach, these advantages can be identified frame-by-frame and in real-time. Additionally, we introduce the concept of effective value added, a statistic that attributes the long-term contribution of both on-ball and off-ball actions. In an individual level we are able to identify those players who contribute the most in the offensive production of a team, but also the ones responsible for defensive mistakes that allow goal opportunities for the opponent.

Diving deeper into the idea of exploiting spaces of value, we relate them to defensive actions. We introduce a novel method to assign defensive areas to each player in the defending team in order to establish defensive responsibilities. Based on this, we propose the idea of passes behind the back, a metric that allows us to explicitly relate an opponent's spatial weakness with specific players, making it easier to communicate to coaches where advantages are found and provide video clips to show how can these be exploited. We also show how the identification of opponent's formation block and pressure lines can be used to add the specific game context to this analysis.

The paper is structured in the following way: first, we describe previous work on space control and expected possession value that are used as building blocks of this work. Then, we show how can we leverage tracking data and observed on-ball events to learn a classification model for detecting potential off-ball advantages at any given time during a game. We then present a method for detecting teams' dynamic formations which leads to the identification of passes behind the back. Finally, we describe the technical details for calculating the effective value-added metric and present several practical applications that shows how the developments of this paper can be used to gain detailed insights for exploiting opponent teams' spatial weaknesses.

2. Background on Expected Possession Value

The concept of expected possession value (EPV) was first introduced in basketball, providing a methodology for estimating the instantaneous expectation of the outcome of any possession [14]. A recent adaption in soccer addresses the frame-by-frame evaluation of the outcome of the possession as well, but it also expands the previous work by incorporating necessary concepts for grasping the nuances of soccer possession such as the possibility of passes going to any location on the field, or the joint evaluation of both the expected outcome of successful and missed actions [5]. Figure 1 illustrates the frame-by-frame evolution of the expected outcome of the possession during a possession of FC Barcelona team (link to the video: https://bit.ly/2RTEfVd). We can observe the fluctuating nature of soccer possessions, the effect of ball-pressure by the opponent, but more importantly how the availability of passing and shooting options, as well as overcoming opponents formation lines increases the likelihood of observing a goal.

The model is built on top of a comprehensive representation of a possession state where space control [7], dynamic pressure lines, players' location and velocity, and relative angle and distance between players and the goal are considered. The most important contribution of this work is the idea of estimating the different components of EPV independently to finally joining them together in order to obtain the EPV value. The independent components estimation as well as the evaluation of risk and reward of passes on any location to the field provides a comprehensive framework for performing more specific advanced analysis of the impact of potential on-ball actions and the influence of players off-ball behavior. We leverage the passing component of this model by using the expected value of potential passes as a building block for assessing the space-value dynamics of each of the 22 players on the field.

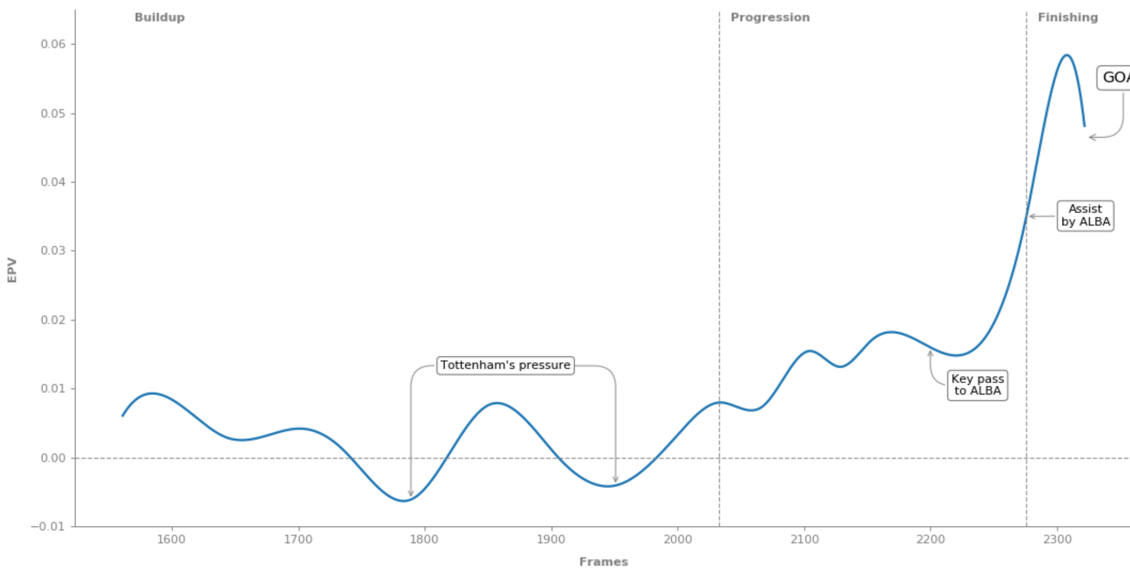

Figure 1. Evolution of the EPV throughout the first goal's attack in the TOT-FCB (UEFA Champions League, groups stage of 2018-2019 season). Note how the model is sensitive to the pressure on the ball carrier, decreasing the likelihood of scoring a goal. In addition, EPV increases significantly with the last two passes of the possession because they get to bypass several defenders and free space is exploited.

3. Exploiting spaces of value

We define the concept of off-ball advantage as those situations when a player is in a favorable disposition to receive a potential pass and, in case of receiving it, the player would likely improve the possession's value (see in Figure 2). The concept is similar to the idea of space creation [7], but off-ball advantages consider implicitly long-term rewards and the risk-reward trade-off thanks to use EPV as metric of value. In addition, as EPV yields a field-wide surface, we can evaluate potential passes to any location of the pitch and then, assess player's decision making.

Our objective is to detect these advantages automatically and we could do it by analyzing the increase in EPV after passing to a player, but how much relative increase does it imply that the receiver is in an off-ball advantage? Unfortunately, deciding the threshold to discern the true advantages from the noise is not straightforward because the EPV is likely to be non-linear. Its value is affected by several dynamic factors of the game, such as opponent's pressure, player movements or pitch control. For example, when a team is moving the ball close to its own goal and the opponent is pressing high up the pitch, the value of the possession will tend to be negative; therefore, almost any successful pass far from the pressure is expected to add a considerable amount of value. On the other hand, at the opposite last third of the pitch, when the team is close to scoring a goal and spaces are more reduced, it is expected to be much more difficult to be able to increase the value of the possession based on movement only.

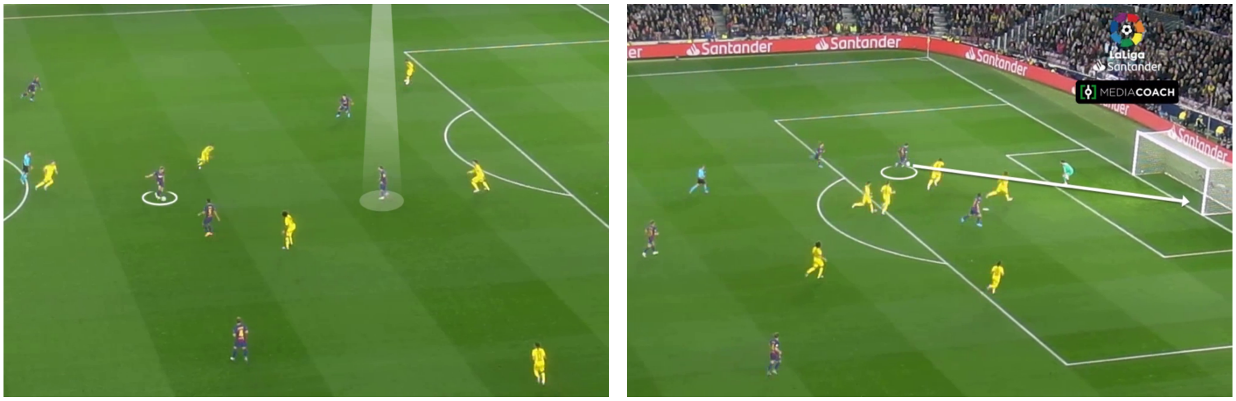

Figure 2. Example of an off-ball advantage by Messi against Borussia Dortmund in UEFA Champions League. Before De Jong executes the pass, Messi is in a space of value as there is no pressure on him. Once he had the ball, he increased the value of the possession, in this case achieving a goal.

Based on that, we have defined the concept of off-ball advantage by means of a probabilistic classifier that estimates whether a potential pass will become a short-term advantage. On average, the EPV added should be enough, but we will use the outcome of the following action to verify that the advantage was materialized. Note that an action includes the set of events performed by the receiver of the pass as Figure 3 shows.

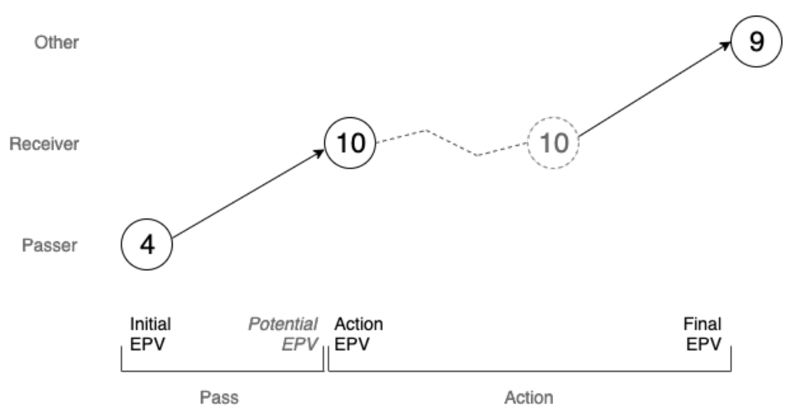

Figure 3. Example of the events to create the dataset. Player #4 passes the ball to player #10, who dribbles, drives the ball and eventually, passes to another player. We will say that #10 was in *off-ball advantage* by analyzing the potential value added before #4 passes and the real value added by the action. The real value is computed as the difference between the final EPV and the initial EPV.

3.1. Methods

For this analysis, we have used 65 matches played by FC Barcelona during the 2018-2019 and 2019-2020 seasons (both in La Liga and UEFA Champions League). The data is composed by event data and spatiotemporal tracking data, provided by *FootoVision* and *Metrica Sports*, with a frame rate of 25 fps. By selecting all passes and the receiver's following action, we formed a dataset of 46780 examples. Note that we have not included those passes with a negative potential value added, because, by definition, this would not correspond to a positional advantage.

We have decided to solve this binary classification problem by training a probabilistic classifier. It estimates the probability of a potential pass being an off-ball advantage, therefore its likelihood of increasing the value of the possession in short term. We have defined the following features:

- EPV before the execution of the pass.
- Distance and angle to the opponent's goal of the origin and destination of the pass.
- The potential value added by the pass, as the difference between the EPV we expected at the destination of the pass and the current EPV just before the execution of the pass.

We have labelled the dataset by defining the class of the problem as positive if the receiver of the pass accomplishes ending his action with a higher value than when the previous pass started. Note that, in case the action includes a pass, the action's final EPV is considered when the pass arrives to the destination. On the other hand, the class will be negative in those cases when the pass is not completed, the receiver losses the ball or his contribution to possession's value is negative.

After standardizing the dataset, we have trained and tuned a XGBoost [15] classifier using 10-fold cross validation. For reproducibility purposes, the best set of hyperparameters was 2000 estimators, learning rate of 0.01 and a depth of 4. In addition, we have used 10 as random state in and performed a stratified shuffle split in order to divide the datasets into train and test sets with a ratio of 70:30, preserving the percentage of positive and negative samples in both sets. Figure 4 shows that the model is correctly calibrated, and the ROC curve achieved.

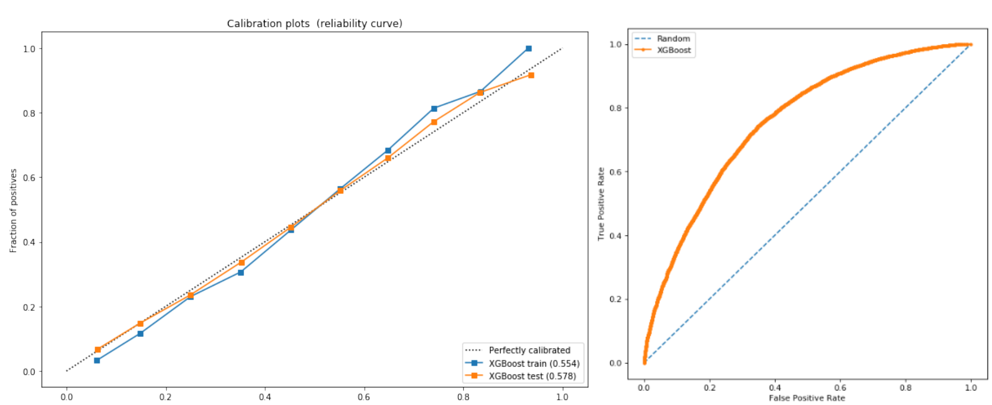

Figure 3. Calibration plot and ROC curve of the best estimator after training and validating several combinations of hyperparameters on a XGBoost classifier via grid search.

3.2. Where do players create positional advantages?

Context has a critical role whenever we want to analyze soccer because of its complexity and the number of factors that affect the behavior of both players and teams. These and other similar plots are created using the average positions of the players throughout the entire match, but there are several clearly separable situations in the game that should be analyzed in isolation in order to be able to extract useful insights and detect common patterns. As an example, player's location in regular play differ considerably from locations in set pieces, so if the latter windows of time are not excluded when computing an average location plot, the result would not be fair.

To this end, we have decided to structure the attacks in regular play in three phases, in a similar way as in [16]. An initial phase called buildup phase, when the team is in its own half of the pitch and the entire opponent team is between the ball and their own goal. Then, the progression phase is considered once the advanced defensive line of the opponent team is overcome and the team tries to get to last quarter of the pitch. Finally, the finishing phase includes those situations close to the opponent's box trying to score a goal. Note that these phases do not have to be sequential, with the possibility of direct jumps between non-consecutive phases. The three formation lines of each team are computed dynamically using the k-means clustering algorithm, with a fixed value of k=3.

With the proposed model in the previous subsection, we can automatically detect how many times and where in the pitch a player offers himself as a possible pass with potential value to another teammate. In order to avoid falling into the pitfalls of averaging situations with different contextual factors, we have selected examples filtered by attacking phase. Finally, the location of off-ball advantage in Figure 4 is plotted relative to the block and defensive lines of the opponents. Note that black crosses represent those advantages when the player received the ball and contributed positively to the possession's value.

The three horizontal lines correspond to the mean location of the defensive team's formation lines for all off-ball advantages of the player. On the other hand, the vertical lines refer to the mean limits of the defensive block. The positions of both the heatmap and the crosses are relative to those lines. So if an off-ball advantage is plotted between the vertical lines and between the first and second horizontal lines, it will not probably have happened exactly on the shown point of the pitch, but it will have happened inside the opponent's defensive block and between their first and second defensive lines.

After this explanation, we can observe how Philippe Coutinho and Leo Messi differ in the creation and exploitation of off-ball advantages ahead the ball carrier. The latter creates positional advantages between the middle and rear defensive lines mostly in the center of the block, even though his supposed position is right winger. Coutinho did produce a few less off-ball advantages and most of them were located inside the defensive block too, but less centric. This is a typical trait of FC Barcelona's wingers, to leave the wings free for both full backs. On the other hand, the off-ball advantages in which the player received the ball and achieved to add more value are distributed unequally. Messi tends to add value both inside and outside the opponent's block over a wider area.

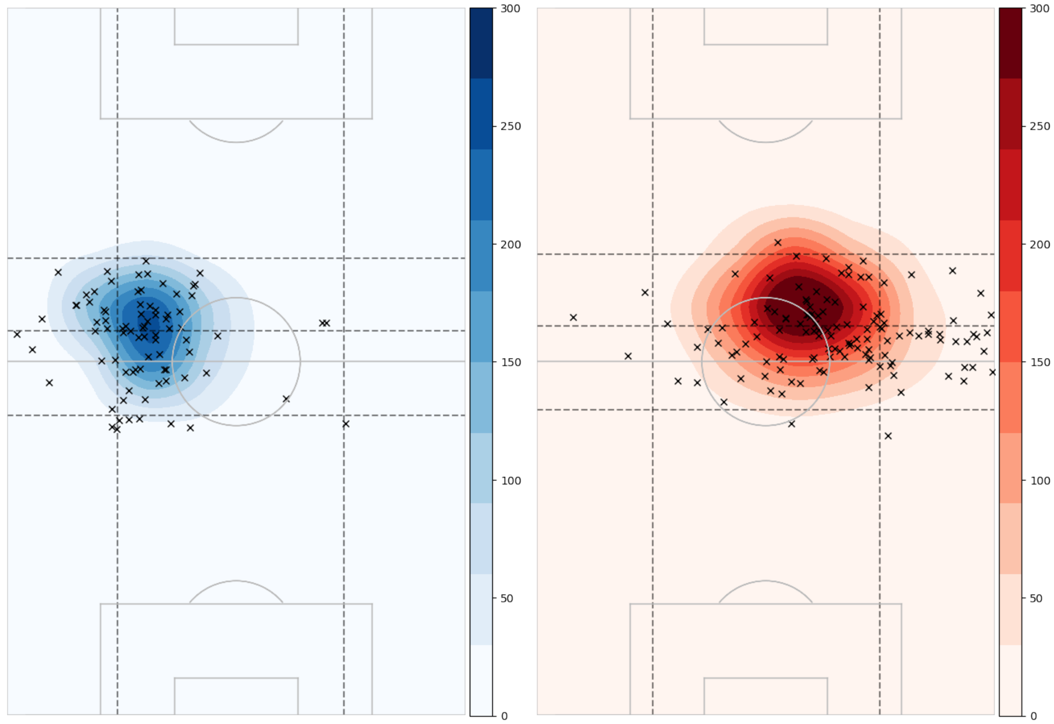

Figure 4. Relative locations to the opponent's defensive lines where Philippe Coutinho (on the left) and Lionel Messi (on the right) created off-ball advantages ahead the ball carrier in progression phases of the 65 matches used in the analysis. Black crosses represent those positional advantages that were exploited, and the player added value with the succeeding action.

Besides analyzing patterns of off-ball movements that add value to the possession, we have applied the presented model to discover usual relationships between pairs of players. Figure 5 shows the number of off-ball advantages created ahead the ball in finishing phases of the second match against Borussia Dortmund in UEFA Champions League 2019-2020.

Forwards are clearly the players that created more off-ball advantages, however there are differences between Leo Messi and Luis Suarez. The former stood out both in the creation and generation of those positional advantages, several of them in the wings outside the opponent block when Sergi Roberto had the ball. We can also observe how Sergio Busquets, compared to the other midfielders, was less times in positional advantage, as we are focusing on situations ahead the ball. His role in this phase was supporting the offensive players behind his position. Note that Antoine Griezmann and Ivan Rakitic did not play the full match, so that his numbers could be a bit lower than the ones of rest of the players.

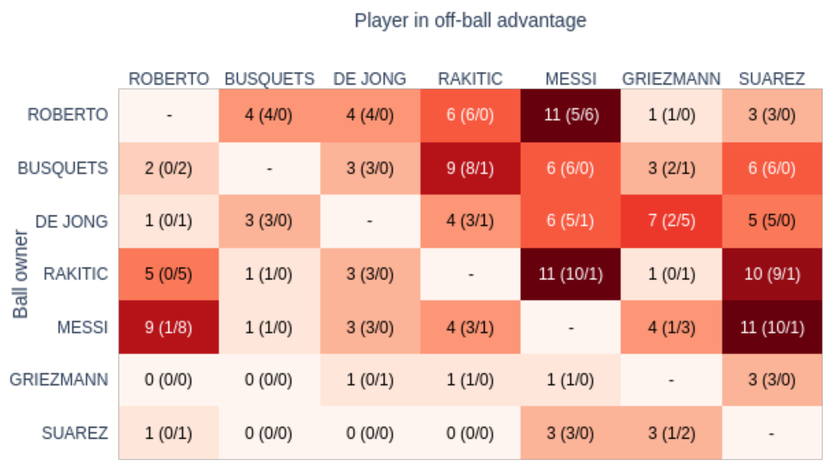

Figure 5. Relationships of FC Barcelona's players regarding off-ball advantages in the finishing phases of the match FCB-BVB (UEFA Champions League 2019-2020). The Y Axis contains the player who carries the ball when the respective player in the X axis was in off-ball advantage. The numbers in parentheses break up the positional advantages created inside and outside the defensive block.

3.3. Do all players resolve materialized off-ball advantages equally?

When a team is attacking, a constant movement of its players towards valuable spaces helps to disorder the opponent team, making it easier to get the ball closer to the opponent's goal to get an opportunity to score. But the team also needs that, whenever players get the ball, they move it wisely, adding value to the possession, and it goes without saying that Leo Messi is good at it.

In Figure 6, we have compared the distribution of potential EPV added by the potential pass when Messi and other forwards were in off-ball advantage versus the real EPV added by them in the positional advantages that were positively exploited and, therefore they added value with a successive action. All forwards were in disposition of receiving the ball with a similar potential value. However, if we focus on the distribution of the real value added by each player once they received the ball, Messi's curve is slightly translated towards higher values, so we could say that he took better advantage of his actions. Bear in mind that the higher the value added by a player, the rarer is to find this kind of actions.

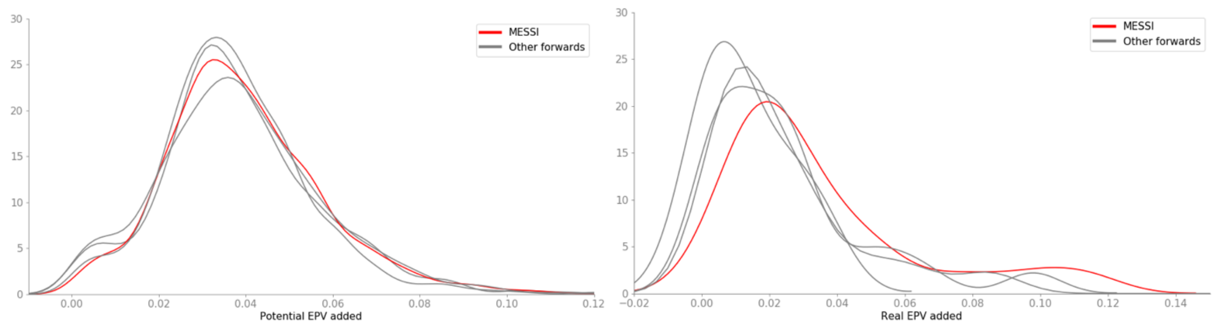

Figure 6. On the left, distribution of the potential EPV added by the pass that would connect the ball carrier with the player in off-ball advantage for several forwards. On the right, the distribution of the real EPV added by those forwards, when they were able to add value to the possession after creating a positional advantage. Leo Messi is highlighted in red whereas the rest of forwards are plotted in grey.

4. Where should players take advantage of free space?

For now, we have only discussed the offensive aspect of the creation and exploitation of spaces of value, but it could be looked from a defensive perspective too. The same way the team in possession of the ball tries to position its players in locations that maximize their chances of scoring a goal, the defending team will try to prevent those players from progressing with the ball. Whenever an off-ball advantage is exploited, there should be one defender responsible for allowing the completion of that pass. Following this line of thought, we have created a new defensive metric which consists on assigning the responsibility of a pass being produced by the opposition to a single defender.

Deciding exactly who is responsible for a defensive error is a complex task. It is normally not just a single action what produces the mistake, but an accumulation of events. We want to use this defensive metric to highlight trends on sequences of games of the same team, so once we have detected an anomaly, we can take a closer look at whether it is the player who is at fault, or the defensive organization of the team.

4.1. Methods

The main idea is to assign to each defender the responsibility for acting on any on-ball and off-ball action that happens on an area of the pitch. When a pass is received by an opponent inside a defender's zone, it will only be considered a defensive error if the receiver of the pass is closer to the goal than the defender. This way we can detect when a defender either allowed a pass to be received behind his back, or he was outside of his designated defensive area.

A first and naïve approach would be to establish static defensive areas determined by the initial structure of the team. However, that kind of approach would be expected to produce noisy and imprecise estimations of the players' defending areas due the constant movement of the defensive block. For example, a certain point in the middle of the pitch should belong to the defensive area of a central back when the team is pressing the opponent, whereas if the team is retreated, it would probably belong to the striker's area. Role changes would also affect this static approach since the defensive responsibilities and areas to cover would change.

Some of these problems have been tackled by using the mean positions of the defenders as seen in [17] and [18]. That would provide a more accurate idea over which spaces a player has occupied. But since our goal is to distribute the pitch space among players, focusing on their mean positions relative to their teammates as in [19] will probably be a more appropriate approach.

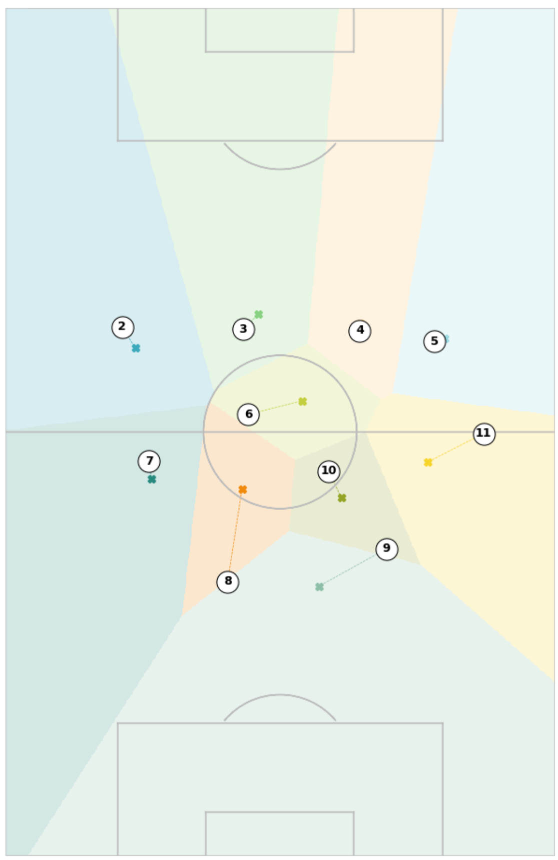

Figure 7. Defensive area assignment for a certain instant in a game. Each colored area corresponds to the defensive zone of action of a defender and the cross to his expected position. The dotted lines link the actual position of a player with its expected position. Note for example how player 8 is on the edge of his designated zone and far from his expected position to pressure the player with the ball. Player 6 has distanced himself from his expected position in order to cover the space left at the back of player 8.

We have computed the dynamic formations of the defending teams. These formations are built using the mean vectors between the players in windows of 2 minutes when the team was defending. Each pass from the attacking team will then have a corresponding defensive formation associated with it. Two different transformations need to be applied to the aforementioned formation before extracting the zones: it needs to be resized to match the spread of the defenders at the moment of the pass, and its centroid needs to be aligned with the current centroid of the defensive block. The result, as seen in Figure 7, will be a new layout of the players where the positions (colored crosses) symbolize their expected locations according to how the team was positioned while defending during the last few minutes.

As shown by the colored areas of Figure 7, we have used a Voronoi diagram to attribute the areas of the field to each player, selecting the expected locations of the players as the centroids. Which means that the defender who is responsible for acting on a certain pass, will be the one whose expected defensive location is closest to the position of the receiver of the pass. To determine whether a player is behind the defender's back or not, the distance of both players towards the goal is used. If the attacking player is closer, then he is receiving a pass behind the back. A plain example is shown in Figure 8, where Messi will receive a pass on Witsel's defensive area while being clearly closer to the goal than the Borussia's midfielder.

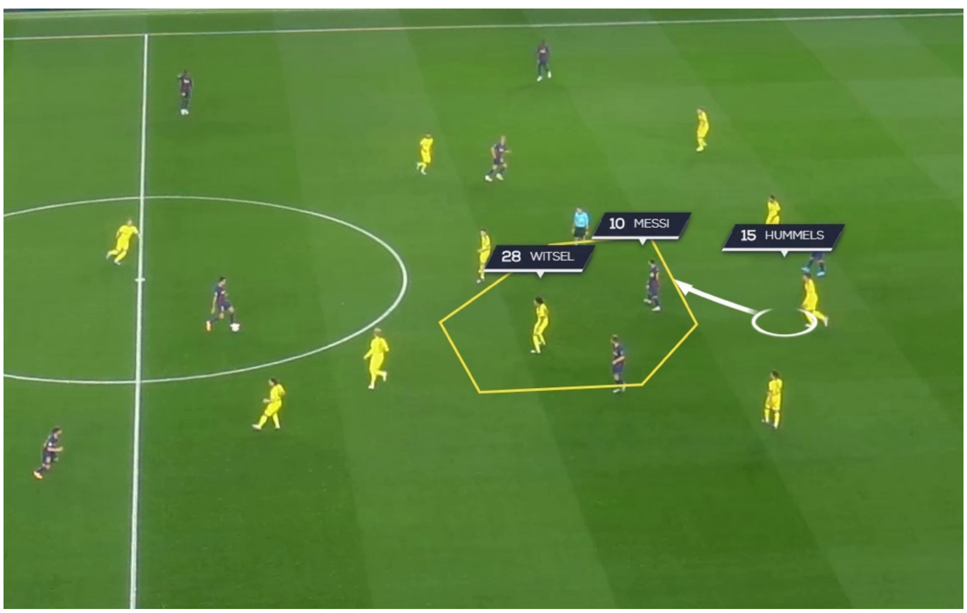

Figure 8. Pass from Busquets to Messi in the FCB-BVB (UEFA Champions League, 2019-2020 season). The yellow area represents Witsel's defensive area of action determined by our method. Note how Messi is positioned inside his area and behind his back. Hummels had to leave his own area to prevent Messi from moving towards the goal after receiving the ball.

4.2. What are the consequences of a pass behind the back?

The main application of this new defensive metric is to perform opponent scouting. It allows us to detect patterns in defensive movements that could help us understand how a team's defense behaves. It is also useful to define profiles of players in order to see whether they tend to leave their assigned defensive area and whether they are easily bypassed.

As an example of helpful information about the defensive structure of a future opponent, we have selected some matches against two teams. Team A plays with two central midfielders, whereas Team B plays with a single defensive midfielder. The goal was to try to find patterns on which of the two central backs tend to leave his own area whenever their opponents achieve passes behind the backs of the midfielders. In Table 1, we observe a tendency in Team A's central backs to cover the spaces behind their corresponding center midfielder. On a few occasions the center backs covered each other's center midfielders, what can be attributed to the disorganization caused by the dynamic nature of the game. We can draw a much stronger conclusion with Team B. The right center back tended to cover the space left by the defensive midfielder more times than the left center back.

Table 1. Number of passes received behind the back of midfielders when a center back was the closest player in front of the ball. The values in (a) have been retrieved from 3 games of team A from seasons 2018-2019 and 2019-2020. The values in (b) have been retrieved from 4 games of team B from season 2018-2019.

a) Team A playing with two center midfielders

	Left Center Midfielder	Right Center Midfielder
Left Center Back	14	7
Right Center Back	9	13

b) Team B playing with one defensive midfielder

	Defensive Midfielder
Left Center Back	3
Right Center Back	10

The detection of this new concept allows us to cut clips of video automatically, as they are the best way to show insights to both coaches and players. An example with three clips from the same match as in Figure 8 can be found on this link (https://bit.ly/2PIBeV2). The clips have been edited later using the *Metrica Play* software.

On the first clip of the video, we see a clear example of Barcelona's Sergi Roberto receiving a pass behind the back of the Borussia's full back Raphaël Guerreiro. The area marked by the yellow line corresponds again to the defender's area of action. It is clear to see that Roberto is positioned much closer to the goal than Guerreiro after his failed attempt to stop the pass, this leaves the Barcelona player with plenty of space to advance on the right wing.

The second clip corresponds to the same pass shown in Figure 8. Note how, when Hummels moves towards Messi, he leaves a valuable space behind his back which could be taken advantage of by Luis Suárez (Barcelona's striker). If we detect that Hummels tends to perform this movement on similar situations, we can tell the player receiving the ball (Messi in this example) to pass it to Suárez either directly or via another teammate. The final clip illustrates the fact that the pressure made by Borussia's Schultz to the player in possession of the ball (Messi) allows Rakitic to receive a pass behind the defender's back and to have more time and space to think on what to do next.

5. Attributing the long-term contribution

During the previous sections we have discussed how to detect the new concepts introduced. In order to communicate useful information about them, they will first need to be aggregated. We have contemplated three different kinds of aggregation, each one helpful to respond a different question.

The first two types are straightforward. One corresponds to the quantitative approach, where you want to answer to the question: how often does that happen? The second one focuses on a qualitative point of view. It tries to answer the question: when it happens, how much danger does it generate? In the case of the off-ball advantage, it would correspond to the increase of EPV when the pass is materialized. Regarding passes behind the back, it would be the increment of the EPV caused by the pass.

The third kind of aggregation aims to attribute the responsibility of the outcome of a possession. It tries to answer the question: how much do some player's actions contribute to the final offensive/defensive production of the team? It acts as a proxy of attribution of effective value. We want to know how significative the contribution was in relation to the final danger of the attack.

In practice, what we do is to assign a certain percentage of the final EPV of the possession to the action. This percentage is calculated by dividing the value added to the EPV by that action, by the sum of the equivalent values for all the on-ball actions of the possession. It should be noted that we discard those actions with negative value added from the sum, since we only want to focus on the actions which increased the team's probabilities of scoring a goal. We will call it effective value added.

This kind of aggregation becomes valuable for detecting defensive flaws, such as the allowance of passes behind the back. For example, if a defender tends to receive a lot of passes behind his back whose mean effective value added is low, it will mean that his team's defensive strategy allows him to be out of position. Being out of position is not always a bad thing, given that if you do not go to pressure the ball carrier, it might be harder to recover it. So, if the team can neutralize at long term the danger created on the possessions when it happens, then it will be part of the defensive strategy. On the other hand, if the mean effective value added of the passes behind the back of a player is high, getting behind his back will be a key to disrupt the whole defensive system.

5.1. Where is a team most hurt by passes behind the back?

Being able to attribute an effective value to a certain action allows us to look beyond player movements and to have a certain qualitative glimpse of reality. The meaning of a single effective value added can be difficult to understand due to its complexity, but the comparison between the aggregated values for players or teams allows us to depict certain patterns or irregularities. Carrying on with the opponent scouting, we can compare the mean effective value added by passes behind the different defenders' backs to see in which parts of its defensive formation a team can be most hurt, as seen in Figure 9.

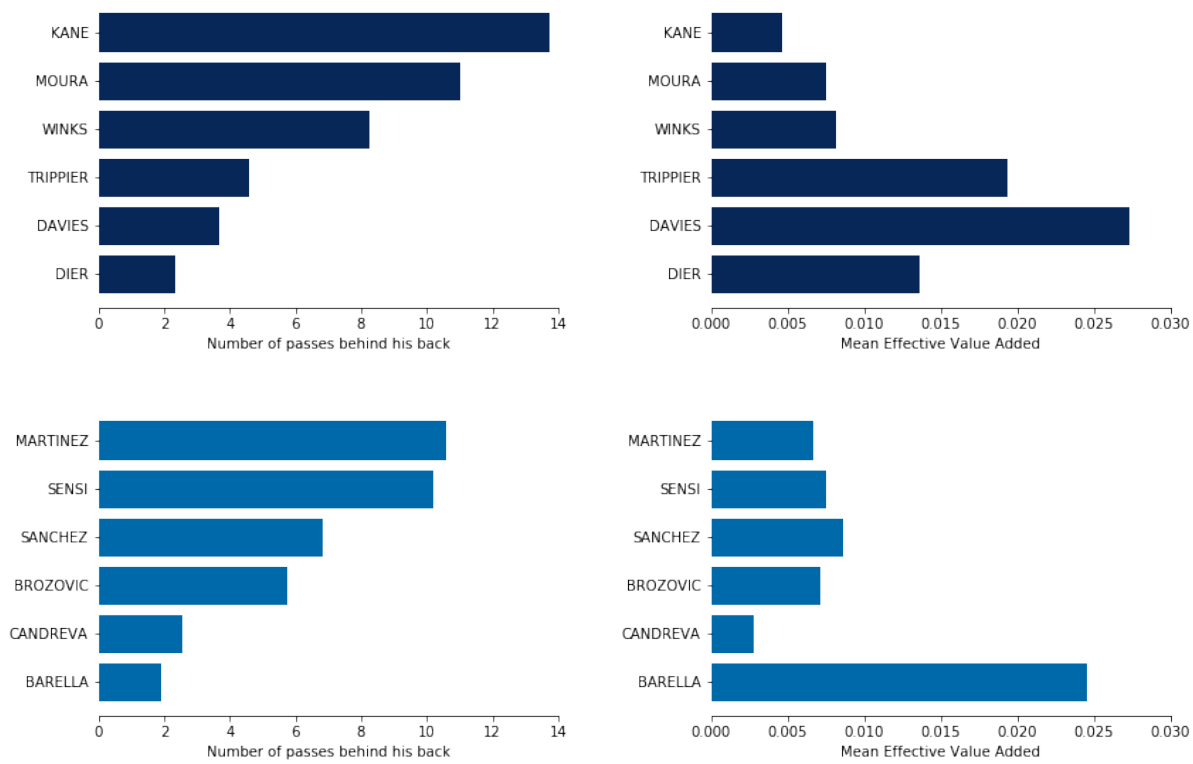

Figure 9. Overview of two defensive stats normalized per 90 minutes for two different UEFA Champions League group stage games. The plots at the top correspond to the TOT-FCB (2018-2019 season) and the ones at the bottom to the INT-FCB (2019-2020 season). The plots on the left show the number of passes received behind the back, whereas the ones on the right shows the mean effective value added of those passes.

The two plots on the left of Figure 9 show the number of times a FC Barcelona's player received a pass behind the back of a selection of defenders from both Tottenham (top) and Inter (bottom). Both plots present a common and obvious trend. Forwards are usually the players most easily bypassed, then the midfielders and finally the defenders.

The story changes when evaluating the mean effective value added by those passes for each defender on the plots on the right. During the game against Tottenham, Barcelona generated their most dangerous attacks by producing key passes behind the back of Tottenham's full backs Trippier and Davies. On the other hand, Inter generally managed to maintain a much more compact block, making it difficult for FC Barcelona to generate good goal scoring chances after making key passes behind the opponent's backs. The only exception was for Inter's midfielder Barella, who was bypassed twice and with a high effective value added.

6. Discussion

This paper presented a new approach for analyzing off-ball performance both at a team and individual level, which provides directly applicable tools to identify and exploit a team's defensive weakness, a key aspect for opponent scouting. Specifically, we proposed three novel ideas: (1) a probabilistic model to detect the concept of off-ball advantage, (2) a method to assign the responsibility of passes to a defender and (3) a way of attributing long-term contribution to players' actions.

First, with (1), we can assess how players take advantage of spaces of value via off-ball movements and whether they exploit these advantages in case of receiving the ball. Therefore, it allows to identify not only observed passes but also potential key passes that would increment possession's value. Secondly, with (2), we have introduced a new method to assign defensive areas to players at any frame by making use of the location of the opponent's block and adjusting it to any context. We have used these areas to attribute dangerous opponent passes to a player, especially those where the assigned defender was behind the ball and no one else was there to defend that space. This new concept allows to discover the opponent's vulnerabilities, and by applying (3), we can attribute the offensive production received by a team to a player in order to effectively communicate it to a coach.

References

[1] Decroos, T., Bransen, L., Van Haaren, J., & Davis, J. (2019). Actions speak louder than goals: Valuing player actions in soccer. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (pp. 1851-1861). ACM.
[2] Lucey, P., Bialkowski, A., Monfort, M., Carr, P., & Matthews, I. (2014). Quality vs quantity: Improved shot prediction in soccer using strategic features from spatiotemporal data. In Proceedings of MIT Sloan Sports Analytics Conference (pp. 1-9).
[3] Power, P., Ruiz, H., Wei, X., & Lucey, P. (2017). Not all passes are created equal: Objectively measuring the risk and reward of passes in soccer from tracking data. In Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1605-1613). ACM.
[4] Spearman, W. (2018). Beyond expected goals. In Proceedings of the MIT Sloan Sports Analytics Conference (pp. 1-17).

[5] Fernández, J., Bornn, L., & Cervone, D. (2019). Decomposing the immeasurable sport: A deep learning expected possession value framework for soccer. In MIT Sloan Sports Analytics Conference.
[6] Dick, U., Tavakol, M., & Brefeld, U. (2018). Rating Continuous Actions in Spatial Multi-Agent Problems.
[7] Fernandez, J., & Bornn, L. (2018). Wide Open Spaces: A statistical technique for measuring space creation in professional soccer. In MIT Sloan Sports Analytics Conference.
[8] Spearman, W., Basye, A., Dick, G., Hotovy, R., & Pop, P. (2017). Physics—Based Modeling of Pass Probabilities in Soccer. In Proceeding of the 11th MIT Sloan Sports Analytics Conference.
[9] Bransen, L., Robberechts, P., Van Haaren, J., & Davis, J. (2019). Choke or Shine? Quantifying Soccer Players' Abilities to Perform Under Mental Pressure. In 13th MIT Sloan Sports Analytics Conference.
[10] Le, H. M., Carr, P., Yue, Y., & Lucey, P. (2017). Data-driven ghosting using deep imitation learning.
[11] Lucey, P., Bialkowski, A., Carr, P., Yue, Y., & Matthews, I. (2014). How to get an open shot: Analyzing team movement in basketball using tracking data. In Proceedings of MIT Sloan Sports Analytics Conference.
[12] Miller, A. C., & Bornn, L. (2017). Possession sketches: Mapping NBA strategies. In MIT Sloan Sports Analytics Conference.
[13] Yu, D., Boucher, C., Bornn, L., & Javan, M. (2019). Playing Fast Not Loose: Evaluating team-level pace of play in ice hockey using spatio-temporal possession data. arXiv preprint arXiv:1902.02020.
[14] Cervone, D., D'Amour, A., Bornn, L., & Goldsberry, K. (2016). A multiresolution stochastic process model for predicting basketball possession outcomes. Journal of the American Statistical Association, 111(514), 585-599.
[15] Chen, T., & Guestrin, C. (2016, August). Xgboost: A scalable tree boosting system. In Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining (pp. 785-794). ACM.
[16] Aranda, R., González-Ródenas, J., López-Bondia, I., Aranda-Malavés, R., Tudela, A., & Anguera, M. T. (2019). "REOFUT" as an Observation Tool for Tactical Analysis on Offensive Performance in Soccer: Mixed Method Perspective. Frontiers in psychology, 10, 1476.
[17] Bialkowski, A., Lucey, P., Carr, P., Yue, Y., Sridharan, S., & Matthews, I. (2014). Large-scale analysis of soccer matches using spatiotemporal tracking data. In 2014 IEEE International Conference on Data Mining (pp. 725-730). IEEE.
[18] Narizuka, T., & Yamazaki, Y. (2018). Clustering algorithm for formations in football games. arXiv preprint arXiv:1805.07525.
[19] Shaw, L., & Glickman, M. (2019). Dynamic analysis of team strategy in professional football. In Barça Sports Analytics Summit.