

Making Offensive Play Predictable - Using a Graph Convolutional Network to Understand Defensive Performance in Soccer

Michael Stöckl, Thomas Seidl, Daniel Marley & Paul Power | Stats Perform

1. Introduction

1.1 Measuring defensive quality in soccer

The art of good defending is to prevent something from happening before it has even happened. Virgil Van Dijk is considered one of the best defenders in world soccer as he has the ability to prevent a pass being made to an open attacker to shoot by forcing the ball carrier to pass somewhere else less dangerous. However, while we know this is great defending, in today's stats, Van Dijk would not receive any acknowledgement. A defender's contribution is simply measured by the number of tackles or interceptions they make. But what if we were able to measure actions that have been prevented before they were made?

The aim of a defense and a defender is to make offensive play predictable. For example, Jürgen Klopp's Liverpool, press the opposition with the aim of forcing them to give the ball away in specific areas of the pitch by limiting the number of passing options available in dangerous areas. If the art of good defending is to make play predictable, then it should be measurable. Given enough data, we should be able to predict where a player will pass the ball, the likelihood of that pass being completed and whether this pass will result in a scoring opportunity. It therefore stands that we should be able to measure if a defender forces an attacker to change their mind or to prevent an attacker from even becoming an option.

Figure 1 shows a situation from a match between Liverpool vs Bayern Munich in the 2018/19 UEFA Champions League that leads to Mané (red 10) scoring. Our model identifies that Milner (red 7) is the primary target for Van Dijk (red 4) in the first instance. However, due to the combination of Gnabry (blue 22) closing down Milner, Lewandowski (blue 9) closing down Van Dijk and Mané making an *active run*, behind the defence, Mané becomes both the most likely receiver and a high threat for scoring. This demonstrates our ability to model how players decision making is influenced and how a situation can move from low threat to high threat by the off-ball actions of attackers and defenders. [\[LINK TO VIDEO\]](#).

In this paper we present a novel Graph Convolutional Neural Network (GNN) which is able to deal with highly unstructured and variable tracking data to make predictions in real time. This allows us to accurately model defensive behaviour and its effect on attacking behaviour, i.e., *preventing actions before they have occurred*.

To do this we trained the following models:

- **xReceiver:** Predicts the likelihood of every player becoming the pass receiver at any moment within a player possession.
- **xThreat:** Predicts the probability of a shot occurring in the next 10 seconds if a pass was played to an attacker.
- **xPass:** Predicts how likely a pass would be completed to each attacker off the ball at any moment within a player possession.

and introduce new defensive concepts:

- **Player Availability:** Using the outputs from xReceiver and xPass we infer how available every attacker is off the ball at each frame.
- **Defensive Impact:** We are able to detect high level defensive concepts such as ball and man orientated defending, defensive position play and off ball runs.
- **Disruption Maps:** Global visual representations of defending teams' ability to disrupt the oppositions attacking strategy.

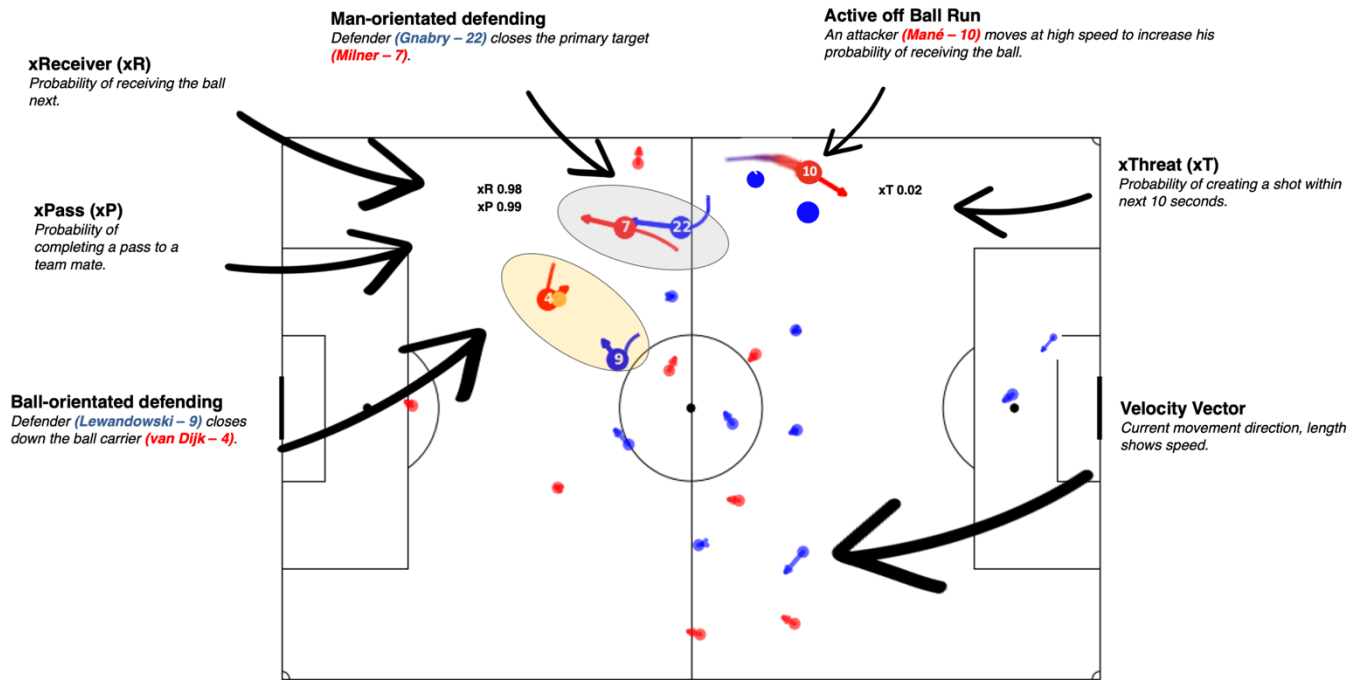


Figure 1: We train three models (xPass, xReceiver & xThreat) to better understand defensive and offensive off the ball behaviour, such as **man-orientated defending**, **ball-orientated defending** and **active off ball runs**. This lets us better understand how Sadio Mané scored a 1-0 lead vs Bayern Munich during Liverpool’s way to the UCL final in 2017/18. [LINK TO VIDEO](#)



Related Work

1.2.1 Dealing with unstructured data

Tracking data is highly unstructured and can be difficult to model due to most machine learning techniques requiring tabular datasets where features are inserted in a specific order. To solve this ordering issue Lucey et al. [1] presented the concept of aligning players to a formation template [1, 2]. However, this method has several limitations. Firstly, in soccer, teams use different formations so a player at role 10 for a team using 433 would be very different to a team playing 352. Therefore, it is difficult to compare predictions between these players. In addition, this method is reliant upon teams having the same number of players on the pitch (11 per team) meaning different models have to be learned when players have been sent off for example. With regards to real time inference a further limitation is the speed at which players need to be aligned before calculating features for inference.

Fernandez [3] and Brefeld [4] use Convolutional Neural Networks (CNNs) on an image representation of tracking data to circumvent the ordering/alignment issue while predicting probabilistic pitch control surfaces. Converting tracking data directly to images is suboptimal as one gives up a very low-dimensional data set and converts it into a high-dimensional sparse representation. Tracking data has an irregular structure due to missing players and a lack of clear scheme to order players in a sequence or frame [19]. Both techniques are also time consuming and could cause timing issues for feature generation in real-time applications. Instead of using an image-based representation we used Graph Neural Networks (GNNs) which 1) neglects the need for ordering features, 2) can cope with varying number of players on the pitch and 3) learns local and higher scale features directly from the tracking data. Horton applied a set-learning framework to model passing in football which is similar to a simple graph (no edges) [18].

1.2.2 Evaluating and Predicting Future Actions in Sports

Concepts such as xThreat and xPass are not new with previous research using tracking data to predict the likelihood of a pass being complete or a goal being scored after a specific action [2, 4, 6]. In addition to measuring the value of an action, the concept of valuing a players' off ball position, has also been investigated [2, 3]. Spearman [9] also developed a model to evaluate off-ball scoring opportunities in soccer. These models create a surface area combining xThreat and xPass values to understand how dangerous a team's or player's current possession is and also what space they control.

Wei [7, 8] modelled the probability of where the next action will go in tennis and soccer. Franks [10] predicted defensive match ups and measured the influence of defenders on the offenses shooting performance. Ghosting [11, 12] hallucinated where a team of defenders will move to next based on where the attackers have moved, and the ball is moving to. This potentially provides useful tools to assess the defensive strategy of teams by evaluating the difference in Expected Goal (xG) or Possession Value (PV) compared to a global baseline.

2. Method

2.1 Data

To train and validate the three models (xTransition, xThreat, and xReceiver) we used 1,200 games of tracking data from multiple seasons of Top 5 European football leagues sampled at 10Hz per second. Tracking data consists of (x, y) positions for each player and the ball, the team and player ids, time, half, and event id at each frame. In total, the dataset consisted of one million passes which was split into a 90/10 train and test set.

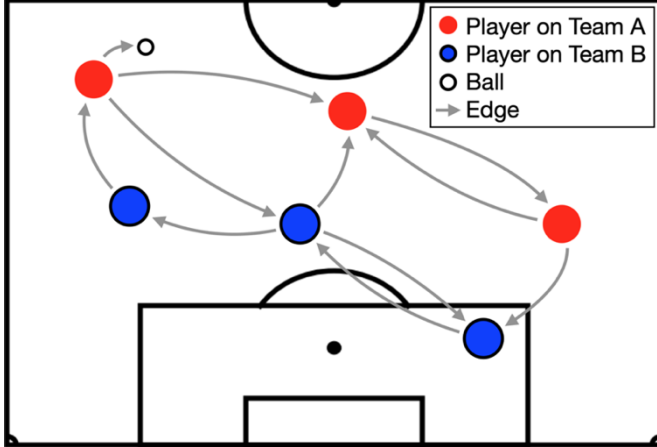


Figure 2: Sketch of the graph representation used for the tracking data. Individual players and the ball are shown as nodes in the graph with directed edges connecting them.

Individual players and the ball are shown as nodes in the graph connected by directed edges.

Edges are weighted to allow the model to learn which nodes are on which team.

For learning a xThreat model only frames relating to the moment of passing events were considered. For the xTransition and xReceiver models we included tracking data from not only the individual passes but also tracking data from one half-second (5 frames) and one second (10 frames) before the pass. Including these two additional moments prior to each pass event allowed us to achieve a semantic regularization during training preventing the model to overfit to the pass moment where players' movements already indicate where the ball will be played to some degree.

2.2 Graph Convolutional Network

To represent the tracking data in a well-defined structure that avoids ordering issues, we used a graph. A graph $G(V, E, U)$ is defined by nodes V , edges E , and global features U . In our representation, as shown in Figure 2, the nodes represent the player and ball tracking data, and the edges contain information about the relationship between the nodes. No global features were included in this approach. The edges e_{ij} are directed and connect a *sending* node v_i to a *receiving* node v_j .

To learn the relationship between the graph input and outputs, we used a GNN. Specifically, we apply the spatial GNN approach that includes separate operations, known as blocks, on the edges and nodes of the graph [13]. An edge block is defined by a neural network that takes inputs from the edge features, sending node features, receiving node features and outputs a new edge embedding.

Similarly, a node block is defined by a neural network that takes inputs from the node features, aggregated sending edge features, aggregated receiving edge features and outputs a new node embedding. A permutation invariant function is required to aggregate the sending and receiving edge features, e.g., the mean or sum of those features. We used a similar GNN architecture for each of the three models: an edge block followed by a node block. Each block has a multilayer perceptron (MLP) with three layers and the number of units in each layer varies between each task.

Using an edge block in the GNN allowed the network to *learn* the relationships between different nodes for each task. This flexibility is not present in standard approaches such as MLPs or convolutions, where the connections between inputs are defined.

Each model outputs a prediction for each player from the final node block. The prediction on each node is the likelihood of a player receiving the next pass (xReceiver), the pass will be completed (xPass), and if there will be a shot on goal within the next 10 seconds (xThreat).



2.3 Features

All three GNN models use the same set of input features. Node features include player XY position, speed, acceleration, angle of motion, distance and angle to the attacking goal, distance to the ball carrier, difference in the angle of motion to the ball carrier, and a flag that indicates whether the player is the ball carrier. The edge features include a flag defining the relationship between the two nodes (teammate 2, or opponent 1), the distance between the two players, and the difference in the angle of motion.

2.4 Training

We trained a unique GNN for each of the xTransition, xThreat, and xReceiver models¹. Across the different tasks, the nodes are weighted in the loss calculation to stabilize the training. The nodes representing the defensive players were masked out for all models. During the training of the xTransition and xThreat model, the only players considered were those that received the pass or were the intended receivers. For the xReceiver model the (intended) receiver of a pass has a weight of 1.0 and all other teammates of the ball carrier have a weight of 0.1 to balance the signal to background.

2.5 Model Training Results

We compared the trained graph models xReceiver, xPass and xThreat against respective baseline models we trained earlier. The baseline models were MLPs that were based on many handcrafted features considering the players' motion characteristics (speed, acceleration, moving direction), relationships between the players shown by differences in the motion features and we also considered the ball information (where it is played with which speed) as seen in [11, 12].

The loss and accuracy of all three GNN models were better than or the same as the metrics of the respective baseline model. Whereas the accuracy was calculated the same for all models, the logloss is not comparable for the xReceiver models. As described above, the logloss of the GNN xReceiver contains all teammates of the passer, even though most of them with small weights*. However, the logloss for the baseline xReceiver is calculated only considering the player a pass was played to. In total, the GNN models are better or at least of the same quality as the baseline models although much fewer handcrafted features were used since the graph was able to learn the spatial relationships between the players on the pitch. Interestingly the baseline xReceiver model had the speed of pass, angle of pass and distance to receiver calculated for the pass moment. However, these features were not included in the GNN. This indicates that the GNN is able to learn the complex dynamic characteristics of football based on the player motion features at the node level combined with the inter-player motion dynamics and team identity at the edge. This is what allows us to use a simpler feature representation allowing faster end to end prediction (0.05s). See Appendix A1 for details.

Table 1: Training metrics of the three GNN models

	Accuracy	logloss
baseline xPass	0.85	0.29
GNN xPass	0.86	0.28
baseline xReceiver	0.73	0.70
GNN xReceiver	0.83	0.07*
baseline xThreat	0.95	0.16
GNN xThreat	0.95	0.16

¹ We use a custom, pure-TensorFlow package based on the DeepMind graph_nets library [https://github.com/deepmind/graph_nets] to train our GNNs.

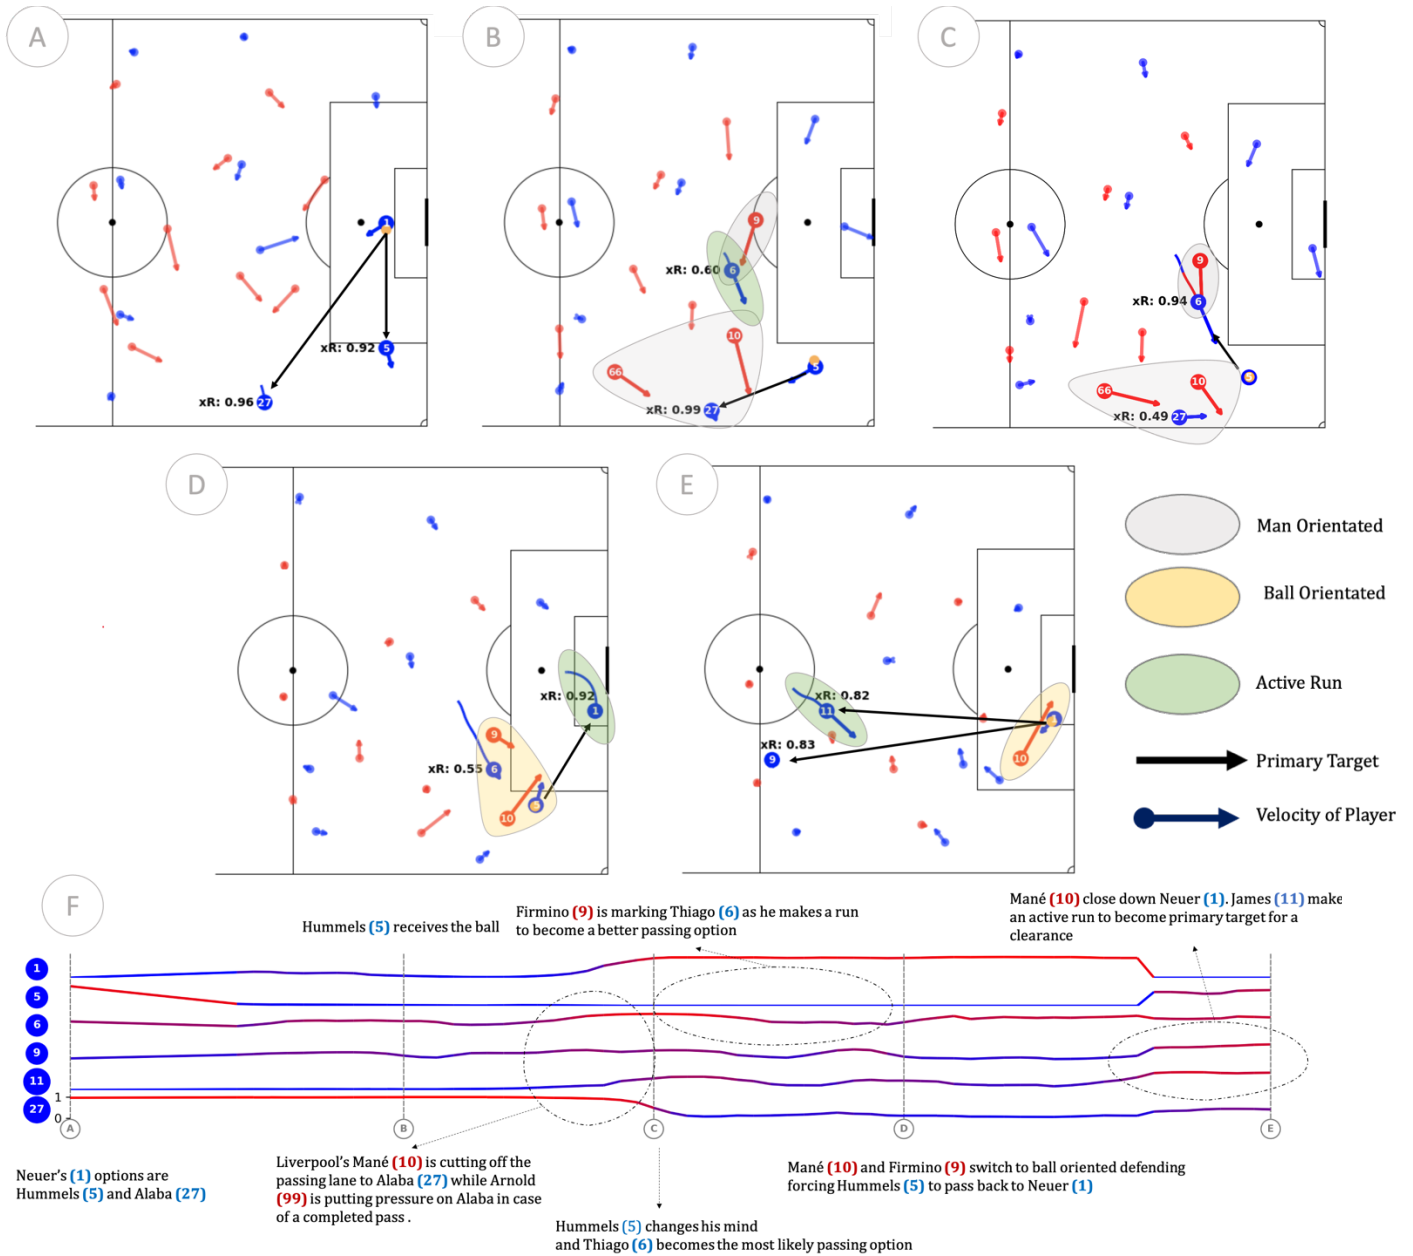


Figure 3: Example of how we are able to capture emergent behaviour. **A-E** show specific moments within a possession of Bayern Munich (blue team). **F** shows frame level estimates for xReceiver (low values: blue, high values: red) for a subset of offensive players throughout the possession. Players are identified by jersey number. Grey ellipsoids show collective man orientated defending whereas orange ellipsoids show ball orientated defending from Liverpool (red team). Green ellipsoids show active runs from the blue team. [LINK TO VIDEO](#)



3. Measuring the Unmeasurable

We introduce a new defensive toolbox, *Defensive Impact*, which analytically determines how much the defending team disrupts the opposition's play. We could measure this by looking at where teams intercept the ball, however as discussed in Section 1 defending is about *preventing what hasn't happened yet*. We demonstrate that the xReceiver model not only accurately predicts who will receive the ball next at the moment of a pass, but when applied at the frame level, captures human decision making. We find the model can predict who will receive the ball in advance of a pass being made and also (more impressively) when a player changes his mind and more importantly, why. Our xThreat and xPass models allow us to value not just what did happen but what could have happened or more accurately **what was prevented** (Figure 3 [link to video](#)).

In the following section we will demonstrate how we created the features that power Defensive Impact and the new insights that are now possible to be generated using an example game between Lazio vs Juventus from 2018/19 Serie A season.

3.1 Disruption Maps

While we are able to create very granular insights with Defensive Impact, the ability to find summary insights is equally important. To provide compressed representations of our insights we introduce *Disruption Maps*. A Disruption Map is a weighted 2d distribution that shows where a team, positively or negatively, disrupted the opposition's off ball options. To calculate a Disruption Map, we first generate a 'spatial identity' (Figure 4 top left) for each team for their xReceiver, xThreat and xPass model output. This acts as a global representation of a team (see Appendix A2 for outputs for all teams in 2018/19 Serie A season). We then calculate the same surface at a game level (Figure 4 top middle) and subtract the two surfaces to create the Disruption Map (Figure 4 top right). This final image reveals where the opposition disrupts (or not in some cases) the opposition's normal strategy/flow². We can now use these maps to determine which team had the biggest impact on their opposition's attacking style and efficiency. We are also able to break this down at a player level to see who was target both on and off the ball.

3.2 Lazio's "Stellungsspiel" - Defensive Position Play

Despite Lazio being beaten 1-2, Juventus' Giorgio Chiellini stated post game "It was the worst Juventus performance of the season for the first 60 minutes" [14]. We used our Disruption Maps to understand what Lazio did to affect Juventus so much and which players they targeted Figure 4 bottom shows Juventus' Disruption Maps for the game. We can clearly see Lazio considerably decreased the xThreat from their left side of the penalty box (Ronaldo's side), with a 10% decrease in Juventus' xThreat compared to all other Juventus games in that season. In addition, we can also see how Lazio increased the risk of completing a pass to any option in the wide channels between the touchline and six-yard box and in the "second penalty area" outside the penalty area by 20%. These are the most dangerous zones where shots and goals are created and the spaces that Juventus' three primary attackers operate. With regards to where players were likely to be a target for a pass, we can see that xReceiver was higher than average around Lazio's penalty box however, as discussed, these areas had a significantly higher risk of competition and more importantly significantly lower probability of leading to goal even if they were completed.

² Disruption Maps estimate tactical deviation from an average style, where differences can be attributed to opposition and/or changes of your own team's tactic for a specific game. To differentiate between the two is out of the scope of this paper.

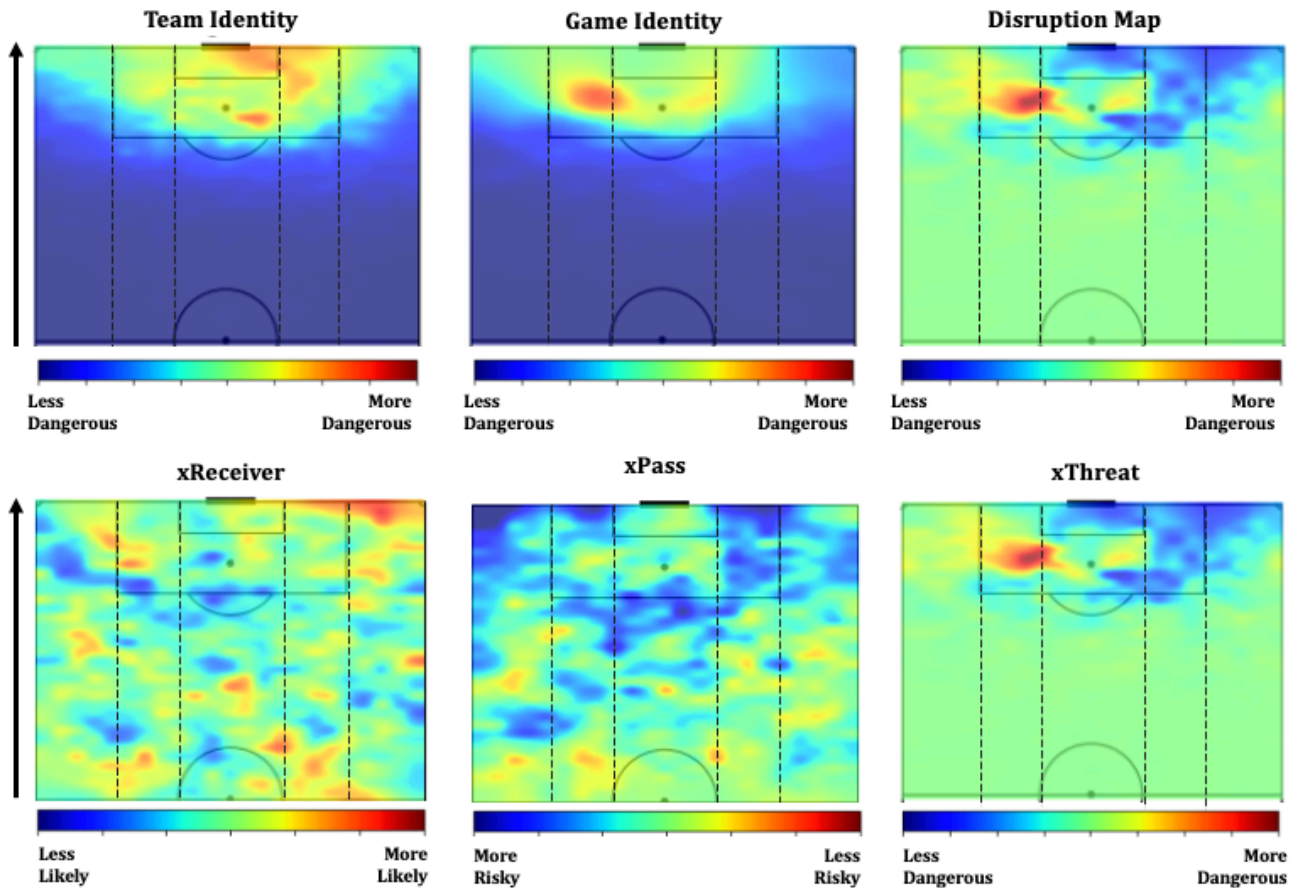


Figure 4. Top row shows how Disruption Maps are composed; they are the difference between a Team Identity and a Game Identity; Bottom row shows Juventus’ Disruption Maps for xReceiver, xPass, and xThreat

To re-emphasise how poor Juventus’ performance was viewed, the German football magazine Kicker stated in the first half, “Cristiano Ronaldo was out of the game, and Dybala also failed to give momentum to the Bianconeri’s offense.” [14]. Our player Disruption Maps back this up (Figure 5). Ronaldo’s probability of receiving a pass was below average across the pitch in first half especially in right channels. This marginally improved in the second half but only in the deep right channel. His threat didn’t exist in the first half (Figure 5 lower left) however, he took up dangerous positions on the left channel. But, as shown in his second half xReceiver Disruption Map (Figure 5 upper right) he was never viewed as a good option to pass to.

Dybala took up significantly more dangerous positions on the left of the penalty area and around the penalty spot. The risk of completing the pass was significantly lower as well indicating he took both threatening and low risk positions in attacking areas. However, he was rarely viewed as a realistic passing option (low xReceiver) with him only being above average on the left flank and the middle of the pitch. This is the critical insight we are now able to surface compared to other methods. We are not just able to measure if a player is in a good position (high threat/low risk), we can measure if they are a realistic option to receive the ball. To understand why Dybala was not a viable option despite his good positioning we need to go a deeper level of analysis which is out of scope for this paper.

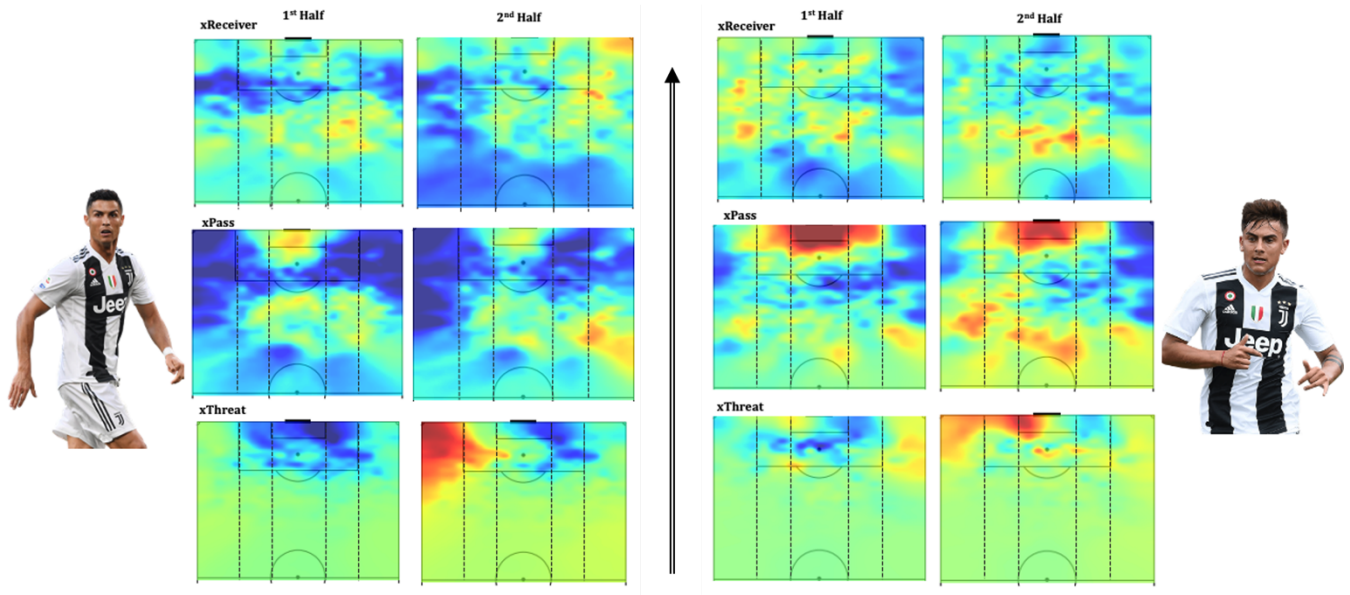


Figure 5: Player xReceiver and xThreat Disruption Maps by half: Left Ronaldo, Right Dybala. Blue areas are below average values whereas red indicates above average regions.

4. Going Deeper – Measuring Decision Making

A primary defensive tactic in soccer is to target key players in possession. These may be players who are important to the team's build up or conversely players who are more likely to give the ball away and cause transition moments. Targeting can happen in three ways, *ball orientated* where defenders will specifically target the player in possession through pressure; *man orientated*, where the ball carrier is not targeted but their passing options are; or a combination of both with the aim of either causing a transition or for the ball to be played to a less dangerous area [16] (Figure 3).

These high-level complex coaching concepts are incredibly difficult to capture and rely on domain experts watching hours of video to find and analyze. However, by using the output from the xReceiver model, we can create a detector to find these moments by simply finding when the primary target changes for the ball carrier. We define the primary target as the attacker with the highest xReceiver value at each frame. If there is a change in the primary target, we make an assumption that a proactive action has occurred. This could be a run from a supporting attacker, the player in possession moving with the ball or a defender(s) actions.

An obvious next step could be to train a model to predict these situations. However, these labels do not currently exist and to ask a set of domain experts to annotate thousands of examples would be highly time consuming and expensive. Instead, we used a programmatic labelling functions approach [17] and asked a domain expert to identify simple functions that capture the expected behaviour of players for different defensive contexts (see Appendix A3 for a sample). Based on observing our initial detected moments with a domain expert, we defined three defensive and two attacking situations (Table 3).

Table 3: Defensive off ball contexts and their definitions as defined by the domain expert.

Situations	Definition
Ball orientated defending	Defender moves closer to the ball carrier at high speed to increase the chance of the ball being given away or pass to less dangerous area.
Man orientated defending	Defender moves closer to the primary target to reduce the probability of them being a receiver.
Ball and man orientated defending	Combination of the previous two.
Active off Ball Runs	An attacker moves at high speed to increase their probability of being a receiver.

4.1 Ronaldo: Null and Void

We can see from the Disruption Maps that in the first half Lazio were able to nullify Ronaldo as a viable passing option. However, how did they do this? Using our Defensive Impact toolbox, we can now assess if Ronaldo was targeted when in possession (ball orientated defending) and out of possession (man orientated defending). In addition, we can assess if Ronaldo made any active runs to become more available for a team mate.

4.1.1 Targeting the Supply Line

To understand why Ronaldo went missing in the first half, we identified the primary players who pass the ball to Ronaldo from the last 5 games (Table 5). We chose the last 5 games as this is how opposition scouts would do it. We measured how often Ronaldo was identified as a primary target for these players (Table 6). We clearly see that the relationship between Matuidi and Ronaldo is heavily targeted, with Ronaldo only being considered an option for Matuidi once during the Lazio match. We also see the difference in Ronaldo's performance between the first and second half, with Ronaldo only being identified 5 times in the first and 8 times in the second. This further supports the observations in Kicker.

4.1.2 Targeting Ronaldo in Possession

Ronaldo was targeted 8 times in the game when in possession with Lazio applying both ball and man orientated defending. Figure 6 (right) shows a primary example of how Lazio closed down Ronaldo's options by applying pressure to Ronaldo as he carried the ball and also applying pressure to his primary targets; Matuidi (red 14) and Sandro (red 12). We can see how Sandro made an active run to become the primary target and how Correa (blue 11) attempted to apply man orientated pressure. However, the defensive run of Luis Alberto (blue 10) to press Ronaldo was responsible for the decrease in Sandro's xReceiver value. At the same time Matuidi (red 14) became the most likely receiver (xR 0.98), however, due to the combined pressure of Milinković-Savić (blue 21) and Parolo (blue 16) his xPass was only 0.66, meaning there was a high chance of a transition.

Table 5. Top 3 passers who started the game from the previous 5 matches

Passer	Pass Total	Games Played Together	Pass Per 90
Alex Sandro	31	4	7.23
Matuidi	22	4	5.31
Dybala	22	4	5.29

Table 6. Ronaldo's availability for his three primary suppliers

Supplier	Ronaldo Target First Half	Ronaldo Target Second Half
Paulo Dybala	1	5
Blaise Matuidi	1	0
Alex Sandro	3	3

Table 7. The availability of Ronaldo's three primary suppliers when he is in possession

Option	Ronaldo In Possession First Half	Ronaldo In Possession Second Half
Paulo Dybala	4	4
Blaise Matuidi	2	1
Alex Sandro	3	6

We have demonstrated how effective Lazio's defensive play was in nullifying Ronaldo and how they did it. The other aspect we can examine is if Ronaldo actively tried to make himself available. Ronaldo made 7 active runs (Figure 6 Right [Link to Video](#)) however only 3 were attacking forward runs. If we compare this to Lazio's forwards, Immobile and Luis Alberto (Figure 7 Left), we see they were more dangerous with their runs. Figure 6 (Left) shows an example of Ronaldo making an active run into the penalty area to become a primary target for Douglas Costa. When Costa received the ball Ronaldo only had an xReceiver of 0.22. However, as we follow his run, we can see him becoming the primary target at the edge of the penalty area. Interestingly, the plot reveals he was always in a highly threatening position for the entirety of the move by playing on the shoulder of the last defender.




Figure 6. Left: Example of Ronaldo (7) making an active off ball run to become the most likely receiver. His trail shows both how dangerous his run was and how his xReceiver increased. **Right:** Ronaldo being targeted with combined man orientated and ball orientated defending. [LINK TO VIDEO](#)




Figure 7. Active Run Maps for Lazio's (Ciro Immobile and Luis Alberto) and Juventus' (Paulo Dybala and Cristiano Ronaldo) main attackers.



4. Summary

We proposed a novel GNN architecture that allows to deal with unstructured data. The GNN allows us to directly learn from a lightweight feature representation the collective behaviour of the soccer players on the pitch. This circumvents the need to order players and for heavy feature crafting which allows us to apply an end-to-end inference pipeline in real-time applications.

Based on the model outputs, we introduced Defensive Impact, a toolbox to measure the influence of defensive strategy on the opposition at a team and player level. Disruption Maps create a compact visual representation of a team's effect on the opposition's xThreat, xPass and xReceiver values compared to their global overage. It allows users to determine where a defence has been successful or struggled based on reducing/increasing an opposition's threat, pass risk and player availability.

We were able to identify different defensive styles (man and ball orientated defending) and off ball runs. Due to the lack of labelled data we utilised a task programming approach – creating labelling functions based on a domain expert's insights. Moving forward an active learning approach where labels are generated, trained against and then assessed would be a recommended next step.

This research enables the evaluation of defensive behaviour and provides tools and insights for coaches and fans to engage with.

5. References

- [1] Xinyu W., Long S., Lucey, P., Morgan, S., & Sridharan, S.. Large-scale analysis of formations in soccer. In Digital Image Computing: Techniques and Applications (DICTA), 2013 International Conference on, pages 1–8. IEEE, 2013.
- [2] 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. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [3] Fernández, J., & Bornn, L. (2020). SoccerMap: A Deep Learning Architecture for Visually-Interpretable Analysis in Soccer. *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*.
- [4] Brefeld, U., Lasek, J., & Mair, S. (2019). Probabilistic movement models and zones of control. *Machine Learning*, 108(1), 127–147. <https://doi.org/10.1007/s10994-018-5725-1>
- [5] Kim, K., Grundmann, M., Shamir, A., Matthews, I., Hodgins, J., & Essa, I. (2010). Motion fields to predict play evolution in dynamic sport scenes. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 840–847
- [6] Cervone, D., D’amour, A., Bornn, L., & Goldsberry, K. (2014). POINTWISE: Predicting Points and Valuing Decisions in Real Time with NBA Optical Tracking Data. *MIT Sloan Sports Analytics Conference*.
- [7] Wei, X., Lucey, P., Morgan, S., Reid, M., & Sridharan, S. (2016). “The Thin Edge of the Wedge”: Accurately Predicting Shot Outcomes in Tennis using Style and Context Priors. *MIT Sloan Sports Analytics Conference*.
- [8] Wei, X., Lucey, P., Vidas, S., Morgan, S., & Sridharan, S. (2015). Forecasting events using an augmented hidden conditional random field. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9006, 569–582. https://doi.org/10.1007/978-3-319-16817-3_37
- [9] Spearman, W. (2019). Beyond Expected Goals. *MIT Sloan Sports Analytics Conference*.
- [10] Franks, A., Miller, A., Bornn, L., & Goldsberry, K. (2015). Counterpoints: Advanced Defensive Metrics for NBA Basketball. *MIT Sloan Sports Analytics Conference*.
- [11] Seidl, T., Cherukumudi, A., Hartnett, A., Carr, P., & Lucey, P. (2018). Bhostgusters: Realtime Interactive Play Sketching with Synthesized NBA Defenses. *MIT Sloan Sports Analytics Conference*.
- [12] Le, H. M., Carr, P., Yue, Y., & Lucey, P. (2017). Data Driven Ghosting using Deep Imitation Learning. *MIT Sloan Sports Analytics Conference*.
- [13] Battaglia, P, et al. (2018). Relational inductive biases, deep learning, and graph networks.
- [14] <https://www.espn.com/soccer/report?gameId=522613>
- [15] <https://www.kicker.de/lazio-gegen-juventus-2019-serie-a-4509058/spielbericht> (in German)
- [16] <https://spielverlagerung.com/2014/07/07/counterpressing-variations/>
- [17] Sun, J. J., Kennedy, A., Eric, Z., Yue, Y., & Perona, P. (2020). Task Programming: Learning Data Efficient Behavior Representations. *arXiv Preprint*, (2011.13917).
- [18] Horton, M. (2020). Learning Feature Representations from Football Tracking. *MIT Sloan Sports Analytics Conference*.
- [19] Mehrasa, N., Zhong, Y., Tung, F., Bornn, L., Mori, G. (2018) Deep Learning of Player Trajectory Representations for Team Activity Analysis. *MIT Sloan Sports Analytics Conference*

A. Appendix

A1. Speed of Inference

Our three GNN models use the same features and the number of features is rather small as described above. This raised our interest whether these models are applicable in real-time or at least close to real-time. We ran a test how long the feature engineering (custom python code) and a GNN model inference of a single frame take on an off-the-shelf notebook³, respectively. Table 2 shows that on average the whole process from crafting the features until the inference of the model is made took less than 0.05s for one tracking timestamp on our test notebook. This enables us to use our GNN models nearly in real-time.

Table A.1: Speed tests how long feature crafting and the model inference takes in seconds [s]; descriptive results of 1000 attempts

	Feature crafting (s)	Inference (s)	Total (s)
mean (s)	0.044	0.001	0.045
std	0.004	0.0002	0.006
min	0.038	0.0007	0.04
max	0.068	0.0029	0.08

³ MacBookPro with a 2.3 GHz Intel Core i5 processor and 16GB RAM

A2. Spatial identity maps for all teams from Serie A season 2018/19 ordered by league table

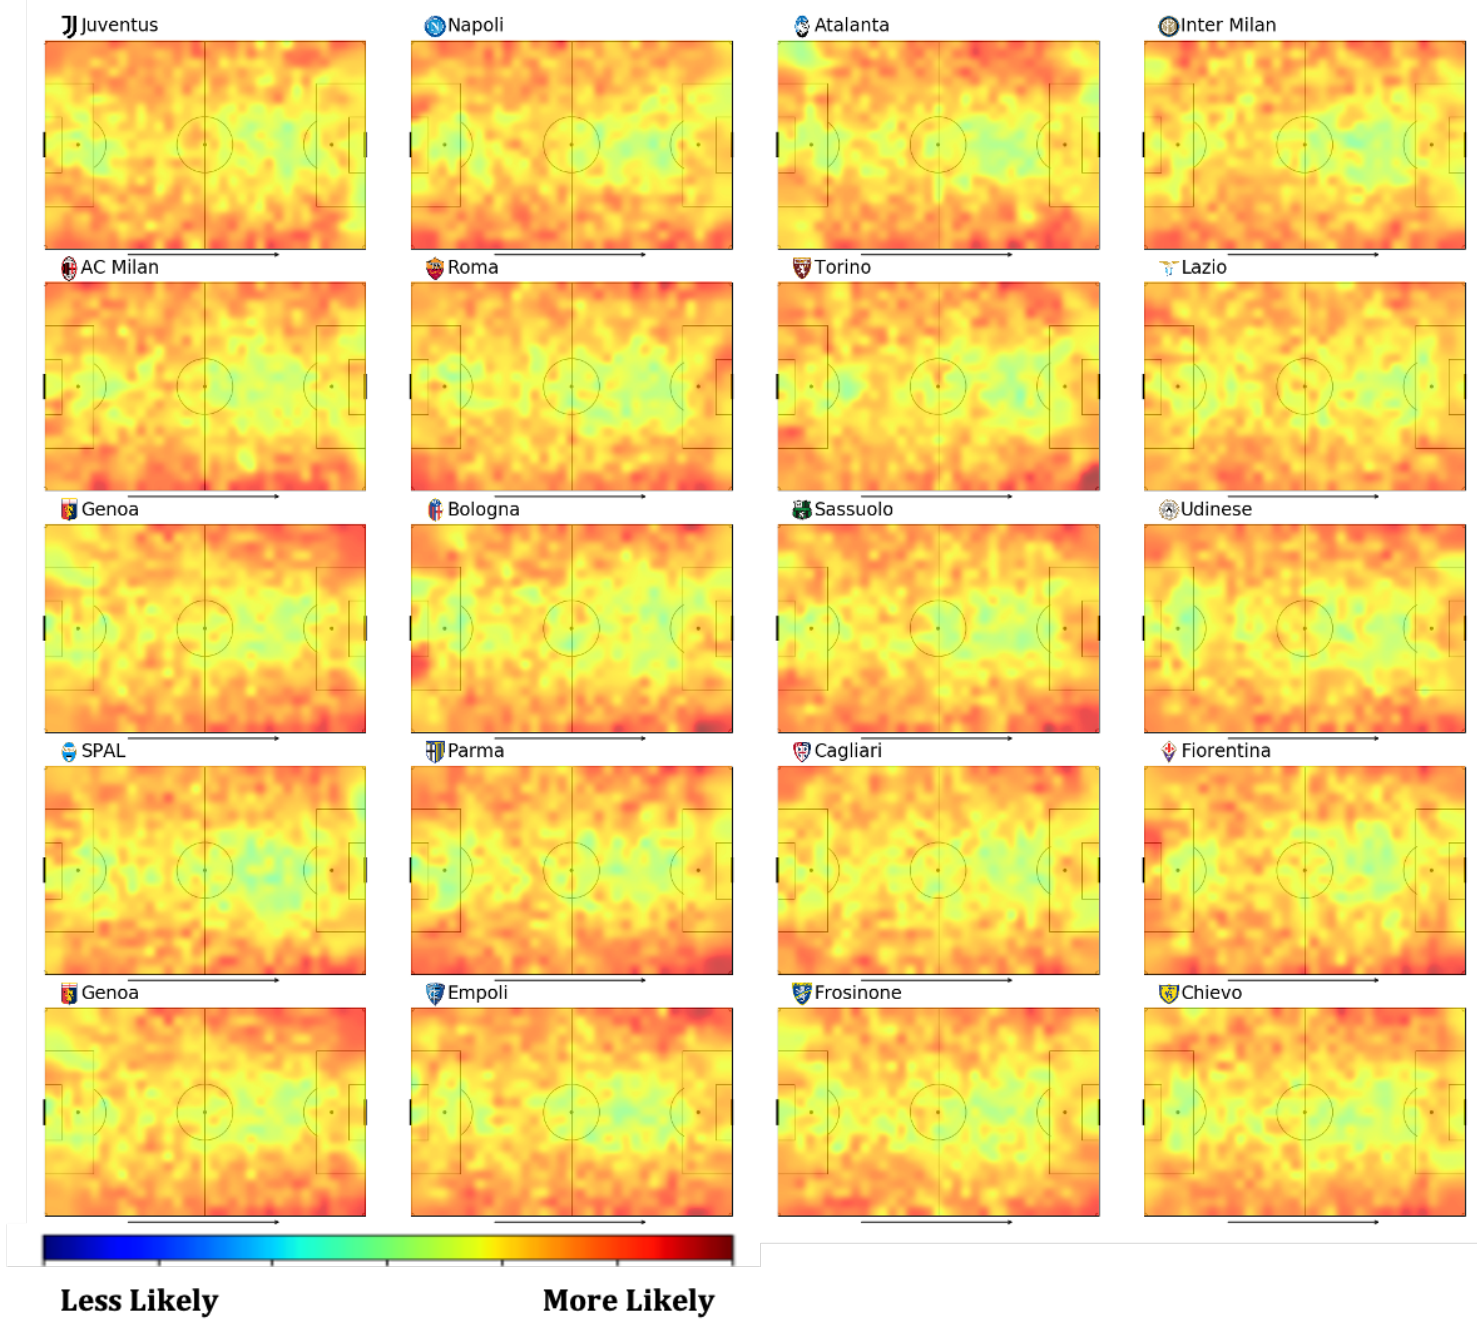


Figure A.1: “Spatial Identity for xReceiver” for all teams from Serie A 2018/19. Teams playing from left to right.

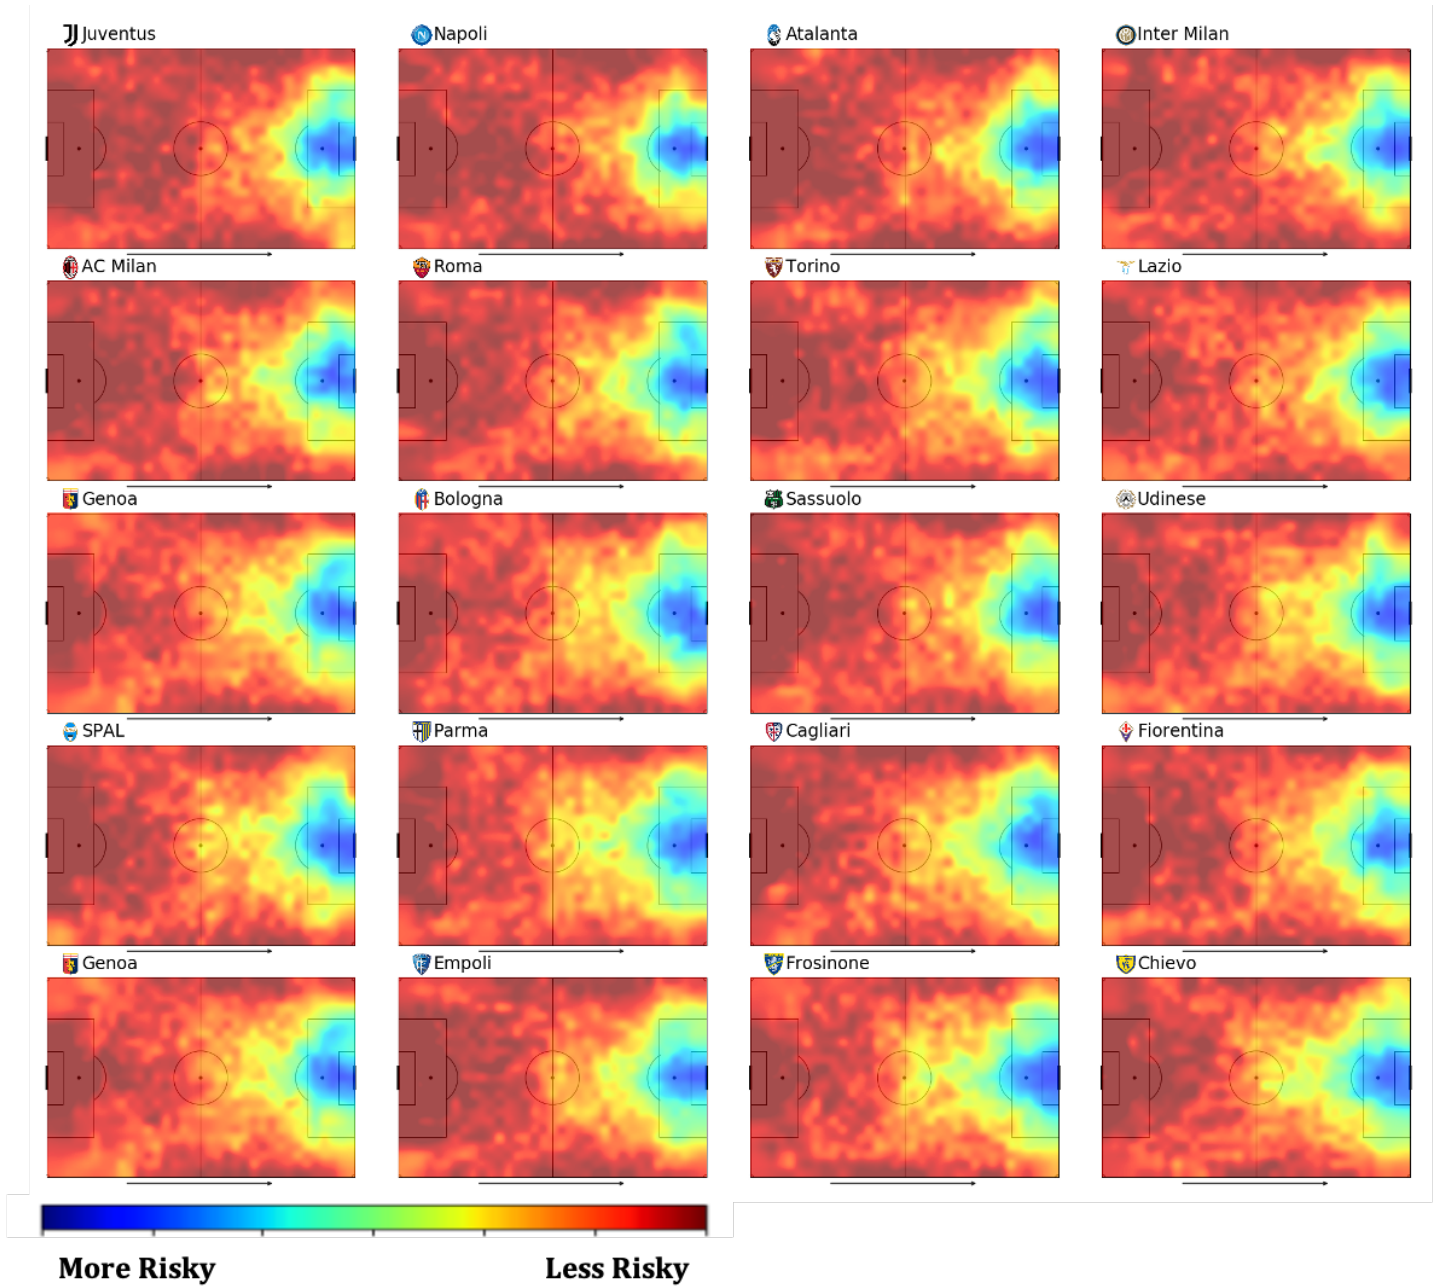


Figure A.2: “Spatial Identity for xPass” for all teams from Serie A 2018/19. Teams playing from left to right.

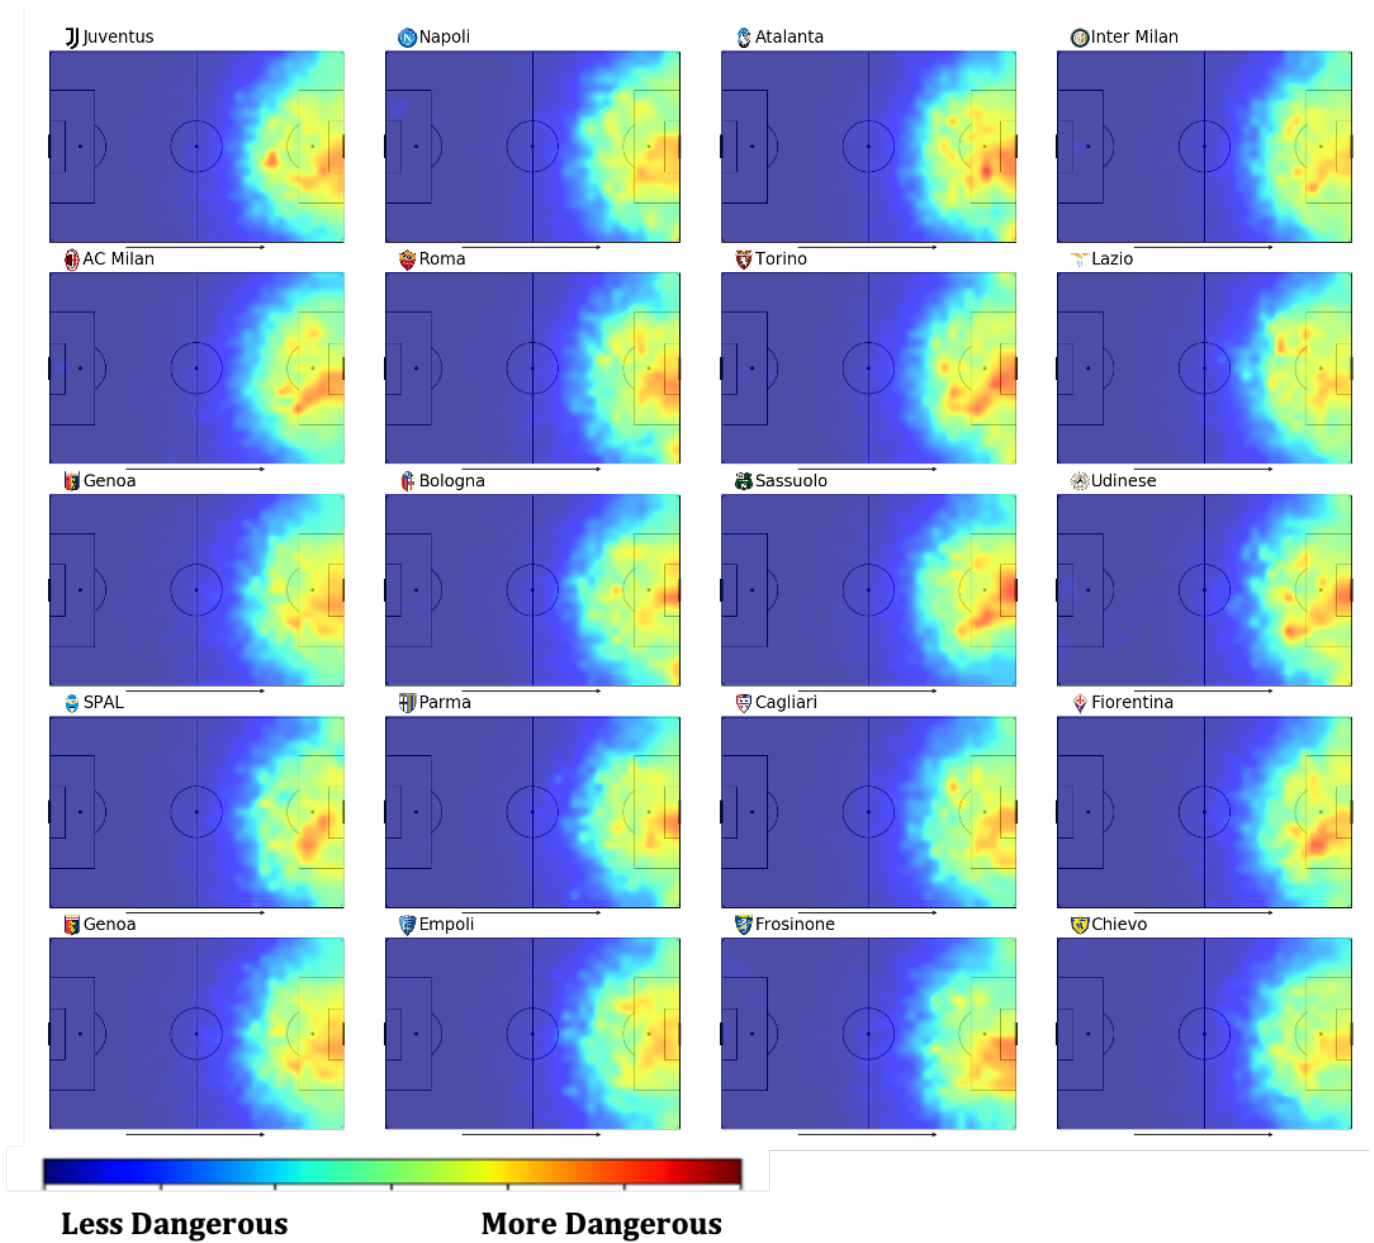


Figure A.3: “Spatial Identity for xThreat” for all teams from Serie A 2018/19. Teams playing from left to right.

A3. Sample pseudo code for a programmatic labelling function

Based on discussions with domain experts the task of identifying man orientated defending was translated into the following pseudo code:

“Defender moves closer to the primary target to reduce the probability of them being a receiver.”

```
Set xr_threshold = 0.8           # player a needs to be a likely pass receiving option
Set speed_threshold = 5 (m/s)    # player b need to actively move
Set distance_threshold = 5 (m)   # player b needs to be close to player a

For team a in possession:
  For each player pa in team a:
    For each player pb in team b:
      If (xReceiver(pa) > xr_threshold) & (speed(pb) > speed_threshold) & (d(pa, pb) < distance_threshold):
        potential_man_orientation(pb) = True
        potential_target = pa
```

Figure A.4: Pseudo code for programmatic labelling of *man orientated defending*.