

# Measuring the Effectiveness of Pressing in Soccer

Simon Merckx<sup>1</sup>, Pieter Robberechts<sup>1</sup>, Yannick Euvrard<sup>2</sup>, and Jesse Davis<sup>1</sup>

<sup>1</sup> KU Leuven, Department of Computer Science; Leuven.AI, Belgium  
simon.merckx@student.kuleuven.be

{pieter.robberchts,jesse.davis}@kuleuven.be

<sup>2</sup> Royal Belgian Football Association, Belgium  
yannick.euvrard@rbfa.be

**Abstract.** Pressing is an important aspect of a soccer team’s defensive strategy. By exerting pressure on the player in possession of the ball, the goal is to win the ball back or at the very least deny the opponents the opportunity to develop an attack. Analyzing and evaluating the effectiveness of pressing strategies is a very important task for any professional match-analyst, but is currently being done exclusively manually by observing video footage. Automating the task saves analysts a tremendous amount of time, standardizes the otherwise subjective task, and allows identifying trends within larger data sets. Therefore, the purpose of this work is to automate the analysis of a soccer team’s defensive pressing strategy. Based on a combination of positional and event data, we first detect pressing situations using a set of expert-defined rules. These pressing situations are successively objectively evaluated by modelling pressing as a trade-off between the benefits of recovering the ball versus the cost of a team leaving its defensive structure, which makes passing through the lines easier for the opposition. We applied this analysis on all matches from a full regular season of the Belgian Jupiler Pro League and show how our metric can be used in practice through a number of use cases.

## 1 Introduction

Preparing for an upcoming soccer match is time-consuming since it requires watching hours of video to better understand an opponent’s tactics. Automated tools can help alleviate this burden by automatically analyzing the large amounts of data available about soccer matches. This spurred the proposal of new key performance indicators that quantify performance about certain aspects of the game, such as shooting [16, 10], passing [4, 19], the occupation of space [8, 5] and the creation of goal scoring opportunities [7, 9, 18].

This paper focuses on a key underexplored problem: valuing a team’s performance in pressing. Pressing is a defensive tactic whereby players exert pressure on the opponent in possession of the ball with the goal of winning the ball back in an advantageous position or at the very least deny the opponents the opportunity to develop an attack. This common definition implies that the defending team

attempts to close down the ball-carrying opponent. We do not consider other varieties of pressing that do not involve direct and aggressive forms of attacking the player on the ball (e.g., defensive pressing) or with other objectives in mind such as the creation of goal-scoring chances (e.g., counterpressing).

Existing work has focused on the spatial aspect [3] and the intensity [1] of pressing. Unfortunately, we currently lack metrics to quantify its effectiveness in different contexts. This arises for two reasons. First, pressing performance is one of the most challenging parts of the soccer game to quantify, because it is a collaborative effort between multiple players and with varying objectives. That is, it involves the coordination of players one or two lines removed from the immediate action and is reliant on the movement of a team in unison. Second, the data in the public sphere is manually annotated event data, which only provides information about the one player carrying the ball. Using this data, it is impossible to capture the coordinated movements that characterize pressing. At best, it allows proxy metrics for pressure that involve opponent pass completion percentage, as viewed relative to pass difficulty [11, 20], different parts of the pitch [24], or the number of defensive actions [22].

We propose a novel approach to evaluate pressing on the basis of positional tracking data. We make two contributions. First, we introduce a rule-based methodology to automatically identify pressing situations in tracking data. We use a set of expert-based rules that capture whether the press is effectively closing down the player in possession. Our second contribution is a metric to quantify pressing performance. Our metric is an extension of the VPEP risk-reward framework [15], which was originally defined to analyze pressing in the context of event data. This framework quantifies the effectiveness of pressing as a trade-off between the benefits of recovering the ball versus the cost of leaving the defensive structure, which makes passing through the lines easier for the opposition. We show how our metric can be used in practice through a number of use cases.

## 2 Data and Preprocessing

Our data set consists of positional tracking data and event data collected by Stats Perform over 241 matches during one season of the Belgian Jupiler Pro League. The positional data was captured by optical tracking systems and logs the positions of every player and the ball at a frequency of 10 frames per second with a resolution of 1 cm. The event data consists of manual annotations of the relevant on-the-ball actions (e.g., passes, shots, tackles and fouls). Synchronization of both data sources and quality control were performed by the data provider. The data was anonymized by removing player and team names.

Based on the tracking data, we compute the player velocities in the x and y direction at each timestamp. These velocities are smoothed using a moving average filter with a window size of 7 frames, skipping unrealistically high running speeds above 12 m/s. Additionally, velocities equal to zero are replaced by an epsilon value of  $1e-10$  to avoid errors while computing pitch control [21, 17].

### 3 Identifying Pressing Situations in Tracking Data

Measuring the effectiveness of pressing first requires a method to detect pressing situations in the data. Our approach includes three steps. First, for each frame in the positional tracking data, pressing-related features are calculated, which are then used to select the frames that involve pressing based on six expert rules. Finally, successive frames involving pressing are grouped into coherent situations.

#### 3.1 Rule-based identification of pressing frames

A rule-based approach to identify pressing was first proposed by StatsBomb. As a unique feature of their event data, they include a *pressure* event type that is triggered when a player is within a five-yard radius of an opponent in possession [12]. While this simple rule fulfills the premise that defending players have to be near the ball carrier in order to press him, it frequently includes events in which players enter the radius without effectively closing-down the ball carrier. We build further upon this rule by adapting the proximity requirement to tracking data and adding additional rules which ensure that the pressing team is well-positioned and actively reducing the space available to the ball carrier.

**Proximity.** Our definition of pressing (Section 1) implies above all that the opponent is put under pressure by closing down the ball-carrying opponent. Hence, in accordance to StatsBomb’s method [12], we require that a defender should be close enough to the ball. Additional rules ensure that there is indeed a player from the attacking team in possession of the ball. Without these rules, a moment when the ball is en route during a pass may be identified as pressure.

*Rules:*

- (1) The distance between the ball and the closest defending player should be less than 4 m.
- (2) The distance between the ball and the closest attacking player should be less than 1.5 m.
- (3) An attacking player should be closer to the ball than any defending player.

**Positioning of defenders.** Closing down the ball carrier requires defenders to be positioned such that they block the passing lanes and the direction towards the goal (i.e., the threat directions). If this is not the case, a distance of 4 m is often too large. To control for this, we use Andrienko et al.’s [1] pressing intensity metric, whose value increases if a defender moves closer to the ball, improves his position or if multiple defenders are near the ball. Using the default parameter settings, we require that the intensity of the press equals at least 35. This corresponds to a single defender being 4 m from the ball and being perfectly in a threat direction. Hence, this metric allows us to eliminate frames in which defenders are close, but poorly positioned and unable to close down the attacking player.

*Rules:*

- (4) The intensity of the press [1] should have a value above 35.

**Movement towards the ball.** The speed and direction in which players move are not taken into account in Andrienko’s model. Yet, this is an important element to identify pressing which requires that defenders are actively attempting to reduce the space available to the attacking team. To summarize the intensity with which the defending team moves towards the ball, we compute the defending team’s average pitch control [8] within a radius of 4 m around the ball and at the exact location of the ball. The pitch control is number in the range  $[0, 1]$  that is computed for every location on the pitch and that should be interpreted as the probability that a team (or player) can control the ball if it would be played to that location. The probabilities are calculated based on the speed and direction in which players move. To be labelled as pressing, we require that the defending team “owns” the space (which corresponds to a pitch control value larger than 0.5) in this radius or has a pitch control value larger than 0.48 at the ball’s location. The threshold values were chosen based on a manual review of the detected frames. A final rule was added to ensure that the closest defender does not stand still. Among other things, this rule filters out data points from static phases such as corner kicks and free kicks.

*Rules:*

- (5) The mean pitch control value in a 4 m radius around the ball should be at least 0.5 or the pitch control value on the ball’s location should be at least 0.48.
- (6) The speed of the defender closest to the ball should be larger than 1 m/s.

In addition to these six base rules, we identify two special cases. First, an extra rule excludes frames in which the ball is in the defending team’s box. Failure to press in this zone will most likely result in a dangerous opportunity for the attacking team. Therefore, in practice, there is no alternative option which we can evaluate.

The second special case occurs on the other side of the pitch. Analogous to StatsBomb’s approach [12], less strict rules are used when the attacking players are put under pressure in their own penalty area, since the potential danger following a mistake under pressure would be much greater compared to any other zone on the pitch. The closest distance from a defender to the ball is 6 m in this zone instead of 4 m. A corresponding threshold value of 15 was chosen for the pressing intensity and a threshold of 0.40 on the average pitch control in a 4 m radius around the ball. The pitch control value at the exact location of the ball is not taken into account. To put more emphasis on the active component, the speed rule is made stricter, requiring a speed of at least 2 m/s.

### 3.2 Grouping pressing frames

The set of rules established in the previous section can be used to select individual frames in the tracking data in which a team performs pressing. However, it is also interesting to see what triggered the pressing and how it evolved. Therefore, each period of possession is split into different pressing situations. A situation aims to capture more of the context surrounding the pressing. An additional advantage is that frames involving pressing that are close to each other are aggregated.

The division between pressing situations is based on time. Starting from the first pressing frame in a period of ball possession, all frames (including frames in which no pressing was detected) that are a maximum of 5 s before or after the selected frame are aggregated into the same pressing situation. A situation can therefore last up to 10 s. Next, we check whether the period of ball possession contains additional pressing frames more than 5 s after the first one. The additional pressing frames are handled analogously to the first one, but its preceding context is limited to the last non-pressing frame in the preceding situation (if the gap with the preceding pressing frame is smaller than 5 s). This ensures that no pressing frames are included in multiple situations.

Finally, we examine whether the average pitch control of the defending team around the ball and the speed of the nearest defender to the ball increase, and whether the distance from the nearest defender to the ball decreases in each of the constructed pressing situations. This final check emphasize the importance of the active component of pressing. The situations and associated frames where this is not the case are discarded.

## 4 Measuring the Effectiveness of Pressing

We measure the effectiveness of pressing as a trade-off between its risk and reward. In this section, we first motivate this risk-reward framework. Then we discuss our decomposed modelling approach to value the risk and reward.

### 4.1 The risk-reward trade-off of pressing

Pressing, like many other actions in football [7, 14], involves a trade-off between risk and reward. If pressing is executed effectively, the defending team can disrupt the opponent’s build-up and deprive them of the opportunity to launch an attack. Ideally, it results in the defending team recovering the ball. However, leaving the defensive structure to press also carries the risk of conceding counter-attacks when the opponent would manage to play through the first pressing line.

This trade-off is the principle on which the VPEP metric [15] is based and on which we build in this work. Let  $\{S_t(r), p_t(r)\}$  be a snapshot of the spatiotemporal tracking data and annotated events after  $t$  frames from the start of a pressing situation  $r$ , with  $p_t(r)$  capturing the pressing context and  $S_t(r)$  capturing the pressing-independent context. And let  $G(r)$  be the outcome of the pressing situation  $r$ , where  $G(r) \in \{1, -1\}$ , with  $-1$  being a shot on goal by the team in control of the ball (corresponding to the risk; denoted  $G^-$ ), and  $1$  being a ball recovery by the pressing team (corresponding to the reward; denoted  $G^+$ ). Analogous to VPEP, the pressing effectiveness  $V(S_t, p_t)$  in a game state  $S_t$  where the defending team puts pressure in a way  $p_t$  is defined as:

$$V(S_t, p_t) = \underbrace{\Delta \mathbb{E}[G^+ | S_t, p_t]}_{\text{Reward}} - C * \underbrace{\Delta \mathbb{E}[G^- | S_t, p_t]}_{\text{Risk}} \quad (1)$$

$$= \mathbb{E}[G^+ | S_t, p_t] - \mathbb{E}[G^+ | S_t, \emptyset] \quad = \mathbb{E}[G^- | S_t, p_t] - \mathbb{E}[G^- | S_t, \emptyset]$$

Hence, our metric assesses how pressing ( $p_t$ ) or a lack of pressing ( $\emptyset$ ) would change the long-term probability for both recovering the ball ( $\Delta \mathbb{E}[G^+ | S_t, p_t]$ )

and yielding a goalscoring opportunity ( $\Delta \mathbb{E}[G^-|S_t, p_t]$ ) in a game state  $S_t$ . Intuitively,  $\Delta \mathbb{E}[G^+|S_t, p_t]$  will be larger if pressing is more likely to succeed, while  $\Delta \mathbb{E}[G^-|S_t, p_t]$  will be larger if pressing is more likely to fail. Since the cost of failure is much higher than the reward of success in this framework, we add a factor  $C = 5$  to the risk term.

## 4.2 A decomposed modelling approach

To infer  $\mathbb{E}[G]$ , we propose a structured modeling approach where the risk and reward are decomposed into a series of subcomponents that model the influence of the press on the ball carrier, and each of his passing options. Each of these components can be estimated separately, simplifying the complexity of the learning task, providing the model with greater adaptability and facilitating the final estimate’s interpretation.

While the player that is pressed can take an infinite set of possible actions in practice, there is a small discrete set of actions that soccer practitioners frequently refer to when discussing a player’s options under pressure (i.e., passing, shooting, dribbling, or clearing the ball). For now, we make the simplifying assumption that the ball carrier will not try to dribble his way out of the press or simply clear the ball, but will try to pass the ball as quickly as possible to a teammate  $l \in L$  when he is pressed. Let  $D_t$  be (the location of) the selected pass receiver at time  $t$ , then based on this assumption, we can decompose  $\mathbb{E}[G]$  as:

$$\mathbb{E}[G^+|S_t, p_t] = \underbrace{P(G_t^+|S_t, p_t)}_{\text{Ball recovery probability}} + (1 - P(G_t^+|S_t, p_t)) \sum_{l \in L} \underbrace{P(D_t = l|S_t, p_t)}_{\text{Pass selection probability}} \underbrace{\mathbb{E}[G_t^+|D_t, S_t, p_t]}_{\text{Expected value of ball recovery}}$$

$$\mathbb{E}[G^-|S_t, p_t] = \underbrace{P(G_t^-|S_t, p_t)}_{\text{Shot selection probability}} + (1 - P(G_t^-|S_t, p_t)) \sum_{l \in L} \underbrace{P(D_t = l|S_t, p_t)}_{\text{Pass selection probability}} \underbrace{\mathbb{E}[G_t^-|D_t, S_t, p_t]}_{\text{Expected threat value}}$$

The structure of  $\mathbb{E}[G^+|S_t, \emptyset]$  and  $\mathbb{E}[G^-|S_t, \emptyset]$  is analogous, but these are calculated without explicitly taking into account the current pressure  $p_t$ . Thus, they capture the chance to recover the ball or encounter danger in a more average match situation described only by  $S_t$ . Implicitly, pressure is still taken into account. For example, a model will still capture that the receiver of an offensive pass close to the goal is likely to be marked. Yet, these models have no explicit information about the actual pressing behavior of the defenders and how it differs from the league average.

Modelling each of the components presents challenging tasks and requires a sufficiently comprehensive representation of the game state to produce accurate estimates. We build these representations from a wide set of fine-grained features describing the passing options, the match context and the defending team’s pressing. The full set of features and their usage within the different models presented in this work are detailed in Appendix A. Next, we discuss each subcomponent model.

**Pass selection model.** We consider pass selection as a ranking problem [23]. Therefore, we obtain the tracking frames for all 169,139 successful passes from the data set right before the pass was played, and construct groups of (usually) ten examples consisting out of the observed pass and the nine other passing options given by the location of each teammate. For each passing option, we then compute the set of features detailed in Appendix A. The observed pass receives a positive label, the other options receive a negative one. Next, we train a pairwise ranking model, learning to rank all passing options within the same group according to their likelihood to be chosen. Finally, a softmax activation function is applied to obtain the probability of each passing option.

**Ball recovery and shot selection.** We train two separate binary probabilistic gradient boosting classifiers to estimate the probabilities that the defending team will recover the ball or that the ball carrier will shoot on goal before a pass is given. The training set consists of 239,881 *Reception* and *Running with ball* actions. For the ball recovery model, an example is assigned a positive label if the defending team manages to regain possession within 5 s in a 5 m radius around the action’s start location. For the shot selection model, an example is assigned a positive label if the ball carrier shoots on goal within 5 s (without first passing).

**Expected risk and reward after a pass.** The models used to estimate the probability of recovering the ball ( $\mathbb{E}[G^+]$ ) or allowing a shot ( $\mathbb{E}[G^-]$ ) after a pass by the pressured player are similar, only differing in their labels. The features are constructed from the tracking data snapshots of successful passes right before the pass was played. The examples receive a positive label if the defending team manages to recover the ball within 5 s ( $G^+$ ) and if the attacking team manages to create a shot opportunity within 10 s after the pass ( $G^-$ ). Again, a binary probabilistic gradient boosting classifier is used to train the models.

### 4.3 Model training and evaluation

We chronologically split the available data into training (60%), validation (20%), and test sets (20%). We use the validation set for model selection and leave the test set as a hold-out data set for testing purposes. We train the models using XGBoost [6] with pairwise loss for the pass selection model and logloss for the other models. For all the models, we perform a grid search on the learning rate ( $\{0.1, 0.05, 0.01\}$ ) and max tree depth ( $\{3, 4, 8, 12\}$ ). To avoid overfitting, we use early stopping on the validation set based on the MAP score for the pass selection model, and the AUROC score for the other models.

Table 1 presents the results obtained in the test set for each of the proposed models. We can observe that the classification accuracy reported for each subcomponent outperforms its naive baseline. The models taking pressing information into account have a better score, which indicates that the pressing-related features provide additional information. Regarding the models’ calibration, we can observe low Adjusted Calibration Error (ACE) values [13] for all the models. Appendix B presents a fine-grained representation of the probability calibration of each of the models.

Table 1: The classification accuracy (AUROC) and calibration value (ACE) for each of the components of our model. Additionally, the table presents the optimal value of each model’s hyper-parameters.

Model	Evaluation metrics					Hyper-parameters	
	AUROC			ACE		Max tree depth	Learning rate
	Baseline	$(S_t, p_t)$	$(S_t, \emptyset)$	$(S_t, p_t)$	$(S_t, \emptyset)$		
Pass selection	0.79 <sup>1</sup>	0.94	0.91	0.0012	0.0019	12	0.05
Ball recovery	0.50 <sup>2</sup>	0.83	0.68	0.0025	0.0027	4	0.1
Shot selection	0.50 <sup>2</sup>	0.87	0.87	0.0009	0.0003	3	0.1
$\mathbb{E}[G^+]$	0.50 <sup>2</sup>	0.72	0.70	0.0095	0.0104	8	0.01
$\mathbb{E}[G^-]$	0.50 <sup>2</sup>	0.83	0.82	0.0030	0.0026	3	0.1

<sup>1</sup>  $P(D_t = l) = e^{-d/25} * \text{PitchControl}(x)$  [18]

<sup>2</sup> Predicting the class prior

Assessing the quality of the joint model is challenging, as there is no ground truth. Therefore, we compute its correlation with statistics that directly measure success like shots and goals at most 10s after the pressure; as well as standard metrics for pressing effectiveness, such as PPDA [22] (Table 2). Additionally, we specifically look into high pressures (on the opponent’s half). As expected, our pressing effectiveness metric is positively correlated with the number of shots created per pressure and negatively correlated with shots allowed. We also observe a positive correlation with league points, indicating that better teams also excel at pressing. Furthermore, it seems that teams that press more often are not necessarily better at it, while pressing high up the pitch and pressing aggressively (i.e., low PPDA) pays off in general.

Table 2: Pearson correlation between each team’s average pressing effectiveness, reward and risk, and several pressing-related metrics.

	Attempted shots / press	Number of high pressing situations	League points	Ball possession (%)	Goals scored / high press	Number of pressing situations	PPDA	Shots allowed / high press	Shots allowed / press
Effectiveness	0.49	0.35	0.30	0.13	0.12	-0.11	-0.34	-0.62	-0.72
Reward	0.59	0.63	0.55	0.62	-0.03	-0.04	-0.85	-0.23	-0.65
Risk	0.05	0.27	0.26	0.55	-0.18	0.05	-0.51	0.45	0.13

To evaluate the stability of our metric, we compare the Pearson correlation between the average values of teams in the first half of the data set with the average values in the second half of the data set. The mean effectiveness of teams in both halves is strongly correlated ( $\rho = 0.72$ ), as are the reward ( $\rho = 0.76$ ) and risk components ( $\rho = 0.57$ ). This implies that teams were consistently assigned comparable mean scores across the two halves of the data set and indicates that our metric detects consistent team behavior.

## 5 Use Cases

Our framework enables a series of novel practical applications. First, using our rule-based identification of pressing situations, we can provide match analysts with a video playlist or a quantitative report giving an aggregated overview of all pressing situations (e.g., how often a team (un)successfully presses in different locations on the pitch). This helps analysts in their daily processes by saving them time, but also by providing objective and comparable benchmarks [2].

More interestingly, using our novel metric, we can obtain direct insight in specific pressing situations at any frame during a match. The presented decomposed modeling approach provides the advantage of obtaining numerical estimates for valuing the positioning and actions of individual players, allowing us to understand the impact that each player has on the development of a pressing situation. In Figure 1, we illustrate this for a specific game situation.

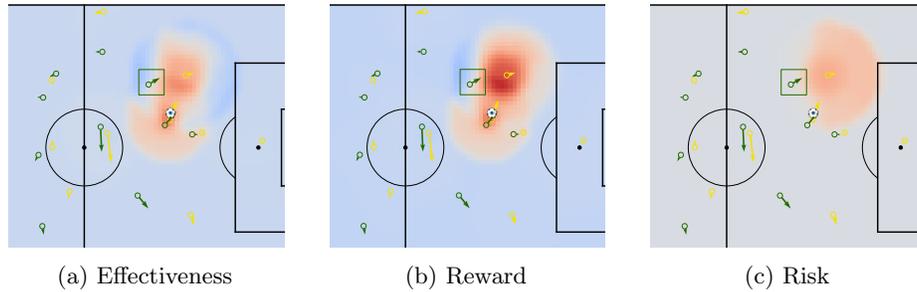


Fig. 1: Our metric can be used to identify the locations on the pitch that maximize the pressing effectiveness in a specific match situation. In this example, the highlighted player can improve his pressing effectiveness by pressing more aggressively. Higher (lower) values for effectiveness, risk and reward are plotted in red (blue). The green team attacks to the right.

Another application of the metric lies in more aggregated statistics to identify trends within larger data sets. This makes it possible to, for example, determine how a team’s pressing evolved over multiple matches, inspect average pressing behavior in specific match situations, and to evaluate the effect of tactical changes or player selections. To illustrate this, Figure 2 shows the effectiveness of a particular team’s pressure on each zone of the pitch over the course of an entire season. We observe an imbalance between the left and right side for defensive pressing, with the team pressing more often and more effectively on the right side; while pressing less and with more risk on the left side. Additionally, the team’s high press is ineffective at recovering the ball. This information could be valuable to the coaching staff.

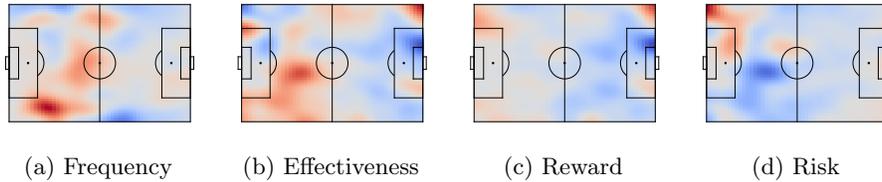


Fig. 2: Heatmaps illustrating the pressing frequency, effectiveness, reward and risk of a particular team relative to the league average. Blue corresponds to less than average; red to more than average. The team attacks to the right.

## 6 Conclusions

This work first introduced an expert-based set of rules to automatically identify all pressing situations in a soccer match based on positional tracking and event data. Using a decomposed modelling approach to evaluate the risk-reward trade-off of pressing, we were able to obtain relevant insights into pressing behavior and effectiveness. Future work will compare the performance of our gradient boosted trees models with deep learning models [9], extend the set of actions that can be performed by the player under pressure, and incorporate characteristics of the player under pressure and possible pass receivers.

**Acknowledgements** The authors would like to thank the Pro League and Stats Perform for providing the positional and event data. We would cordially like to thank dr. Matteo Balliauw (RBFA Knowledge Centre of the Royal Belgian Football Association) for his feedback and help. This research received funding from the Flemish Government (AI Research Program). Pieter Robberechts is supported by the KU Leuven Research Fund (C14/17/070).

## References

1. Andrienko, G., Andrienko, N., Budziak, G., Dykes, J., Fuchs, G., Landesberger, T., Weber, H.: Visual analysis of pressure in football. *Data Mining and Knowledge Discovery* **31**(6), 1793–1839 (Nov 2017). <https://doi.org/10.1007/s10618-017-0513-2>
2. Bauer, P., Anzer, G.: Data-driven detection of counterpressing in professional football. *Data Mining and Knowledge Discovery* pp. 1–41 (2021). <https://doi.org/10.1007/s10618-021-00763-7>
3. Bojinov, I., Bornn, L.: The pressing game: Optimal defensive disruption in soccer. In: *Proceedings of the 10th MIT Sloan Sports Analytics Conference*. pp. 1–8. SSAC '16, Boston, USA (2016)
4. Bransen, L., Van Haaren, J., van de Velden, M.: Measuring soccer players' contributions to chance creation by valuing their passes. *Journal of Quantitative Analysis in Sports* **15**(2), 97–116 (2019). <https://doi.org/10.1515/jqas-2018-0020>
5. Brefeld, U., Lasek, J., Mair, S.: Probabilistic movement models and zones of control. *Machine Learning* **108**(1), 127–147 (2019). <https://doi.org/10.1007/s10994-018-5725-1>

6. Chen, T., Guestrin, C.: XGBoost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 785–794. KDD '16, ACM, New York, NY, USA (2016). <https://doi.org/10.1145/2939672.2939785>
7. Decroos, T., Bransen, L., Van Haaren, J., Davis, J.: Actions speak louder than goals: Valuing player actions in soccer. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1851–1861. KDD '19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3292500.3330758>
8. Fernandez, J., Bornn, L.: Wide open spaces : A statistical technique for measuring space creation in professional soccer. In: Proceedings of the 12th MIT Sloan Sports Analytics Conference. pp. 1–19. SSAC '18, Boston, USA (2018)
9. Fernández, J., Bornn, L., Cervone, D.: A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. *Machine Learning* **110**(1), 1–39 (2021). <https://doi.org/10.1007/s10994-021-05989-6>
10. Lucey, P., Bialkowski, A., Monfort, M., Carr, P., Matthews, I.: Quality vs quantity: Improved shot prediction in soccer using strategic features from spatiotemporal data. In: Proceedings of the 8th MIT Sloan Sports Analytics Conference. pp. 1–9. SSAC '14, Boston, USA (2014)
11. Morgan, W.: Under pressure. <https://statsbomb.com/2017/09/under-pressure/> (September 2017), accessed: 2021-05-03
12. Morgan, W.: How statsbomb data helps measure counter-pressing. <https://statsbomb.com/2018/05/how-statsbomb-data-helps-measure-counter-pressing/> (May 2018), accessed: 2021-05-01
13. Nixon, J., Dusenberry, M.W., Zhang, L., Jerfel, G., Tran, D.: Measuring calibration in deep learning. In: CVPR Workshops. vol. 2 (2019)
14. Power, P., Ruiz, H., Wei, X., Lucey, P.: Not all passes are created equal: Objectively measuring the risk and reward of passes in soccer from tracking data. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1605–1613. KDD '17, ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3097983.3098051>
15. Robberechts, P.: Valuing the art of pressing. In: Statsbomb Innovation in Football Conference. London, United Kingdom (Nov 2019), <http://statsbomb.com/wp-content/uploads/2019/10/Pieter-Robberechts-Valuing-the-Art-of-Pressing.pdf>
16. Robberechts, P., Davis, J.: How data availability affects the ability to learn good xG models. In: Proceedings of the International Workshop on Machine Learning and Data Mining for Sports Analytics. pp. 17–27. Springer, Gent, Belgium (virtual) (2020). [https://doi.org/10.1007/978-3-030-64912-8\\_2](https://doi.org/10.1007/978-3-030-64912-8_2)
17. Shaw, L.: Laurie on tracking. <https://github.com/Friends-of-Tracking-Data-FoTD/LaurieOnTracking/> (2020)
18. Spearman, W.: Beyond expected goals. In: Proceedings of the 12th MIT Sloan Sports Analytics Conference. pp. 1–17. Boston, USA (2018)
19. Steiner, S., Rauh, S., Rumo, M., Sonderegger, K., Seiler, R.: Outplaying opponents—a differential perspective on passes using position data. *German journal of exercise and sport research* **49**(2), 140–149 (2019). <https://doi.org/10.1007/s12662-019-00579-0>
20. Szczepański, L., McHale, I.: Beyond completion rate: evaluating the passing ability of footballers. *Journal of the Royal Statistical Society. Series A* pp. 513–533 (2016). <https://doi.org/10.1111/rssa.12115>
21. Thomson, W.: Metrica pitch control. <https://github.com/anenglishgoat/Metrica-pitch-control> (2020)

22. Trainor, C.: Defensive metrics: Measuring the intensity of a high press. <https://statsbomb.com/2014/07/defensive-metrics-measuring-the-intensity-of-a-high-press/> (July 2014), accessed: 2021-03-02
23. Vercruyssen, V., De Raedt, L., Davis, J.: Qualitative spatial reasoning for soccer pass prediction. In: CEUR Workshop Proceedings. vol. 1842 (2016)
24. Ward, D.: Gladbach: Leaving protectionism behind. <https://statsbomb.com/2016/08/gladbach-leaving-protectionism-behind/> (2016)

## A Features

The table below presents the set of features used as input for each model. The features can be divided into three categories. The first category describes the possible passing options in a match situation. Each of these passes is described separately according to some characteristics of the player in possession, the receiver of the pass and the pass itself. The second category consists of match context features. These are identical for all possible passes, as they are independent of the sender and receiver of the pass. They describe the broader context of the match and are constructed based on event data. The passing and match context features together capture the match situation  $S_t$ . The last category are the pressing-related features  $p_t$ , which describe the pressure of the defending team on the player in possession and each possible receiver of a pass. The “category” column indicates to which group the feature belongs. For the remaining columns, a checkmark indicates in which models the feature is used, including: pass selection probability (PS), ball recovery probability (BR), shot selection probability (SS) and expected risk and reward after a pass (G). The ranking model that estimates the pass selection score for a pass and the classifiers that estimate the expected risk and reward after the pass only use the passing and pressing features that belong to that pass. However, the pass selection model implicitly takes into account the features of all possible passing options since the probabilities are determined by applying the softmax function over the predicted scores in a group.

Category	Feature description	PS	BR	SS	G
Passing	Distance to pass receiver	✓			✓
Passing	Distance between the location of the pass receiver and the center of the goal	✓			✓
Passing	Distance between the location of the ball and the center of the goal	✓	✓	✓	✓
Passing	Distance between the location of the pass receiver and the closest sideline	✓			✓
Passing	Distance between the location of the pass receiver and the goalline	✓			✓
Passing	Distance between the location of the ball and the closest sideline	✓	✓	✓	✓
Passing	Distance between the location of the ball and the goalline	✓	✓	✓	✓
Passing	Difference in distance to the center of the goal between the location of the ball and the location of the receiver	✓			✓

Passing	Difference in distance to the closest sideline between the location of the ball and the location of the receiver	✓			✓
Passing	Difference in distance to the goalline between the location of the ball and the location of the receiver	✓			✓
Passing	The speed of the player giving the pass	✓	✓	✓	✓
Passing	The speed of the pass receiver	✓			✓
Passing	Angle between the location of the ball and the center of the goal	✓	✓	✓	✓
Passing	Angle between the location of the receiver and the center of the goal	✓			✓
Passing	Angle between the movement direction of the player giving the pass and the receiver	✓			✓
Passing	Angle between the passer, the receiver and the center of the goal.	✓			✓
Passing	Angle between the receiver, passer and the center of the goal.	✓			✓
Match context	Time elapsed since the start of the game	✓	✓	✓	✓
Match context	Time elapsed since the last action of the defending team	✓	✓	✓	✓
Match context	Number of goals scored by the defending team	✓	✓	✓	✓
Match context	Number of goals scored by the attacking team	✓	✓	✓	✓
Match context	Goal difference	✓	✓	✓	✓
Pressing	Distance between the pass-receiver and his closest defender	✓			✓
Pressing	Speed of the defender closest to the pass-receiver	✓			✓
Pressing	Distance between the player giving the pass and his closest defender	✓	✓	✓	✓
Pressing	Speed of the defender closest to the player giving the pass	✓	✓	✓	✓
Pressing	Difference in speed between the pass-receiver and his closest defender	✓			✓
Pressing	Difference in speed between the player giving the pass and his closest defender	✓	✓	✓	✓
Pressing	Number of defenders in a 3 m radius around the ball		✓	✓	
Pressing	Number of defenders in a 5 m radius around the ball		✓	✓	
Pressing	Distance between the location of the defender closest to the ball and the nearest sideline		✓	✓	
Pressing	Distance between the location of the defender closest to the ball and the goalline		✓	✓	
Pressing	Difference in distance to the goalline between the location of the ball and the closest defender		✓	✓	

Pressing	Difference in distance to the closest sideline between the location of the ball and the closest defender	✓	✓		
Pressing	Number of defenders in a 10m radius around the ball	✓	✓	✓	✓
Pressing	Number of defenders in a 10m radius around the pass-receiver	✓			✓
Pressing	Average pitch control in a 4m radius around the location of the ball	✓	✓	✓	✓
Pressing	Average pitch control in a 4m radius around the pass-receiver’s location	✓			✓
Pressing	Number of defenders on or nearby the pass line <sup>3</sup>	✓			
Pressing	Difference in distance to the goalline between the location of the pass-receiver and his closest defender	✓			✓
Pressing	Difference in distance to the closest sideline between the location of the pass-receiver and his closest defender	✓			✓
Pressing	Angle between the location of the pass-receiver, his closest defender and the center of the goal	✓			✓

## B Evaluation

Figure 3 presents a detailed representation of the probability calibration of each model. The x-axis represents the predicted probabilities, while the y-axis represents the empirical probability among the examples in our test set. Additionally, the histogram represents the distribution of the predicted probabilities. In these plots, we can observe that the different models provide calibrated probability estimates along their full range of predictions, which is crucial to be able to use them to analyze the impact that specific decisions have on the effectiveness of pressing. Furthermore, we can observe that the (expected) ball recovery probabilities cover a wider range than the (expected) shot selection probabilities, indicating that it is harder for the attacking team to create a shot opportunity than it is for the defending team to recover the ball. Hence, our  $C = 5$  constant in formula 1 to balance the risk and reward.

<sup>3</sup> A rectangle is constructed around the pass line. Within this rectangle the number of defenders is counted. The width of the rectangle is based on the distance  $d$  of the pass. For passes shorter than 10m the width is  $d/4$ . For passes with  $d > 10$ , the width is equal to 5m.

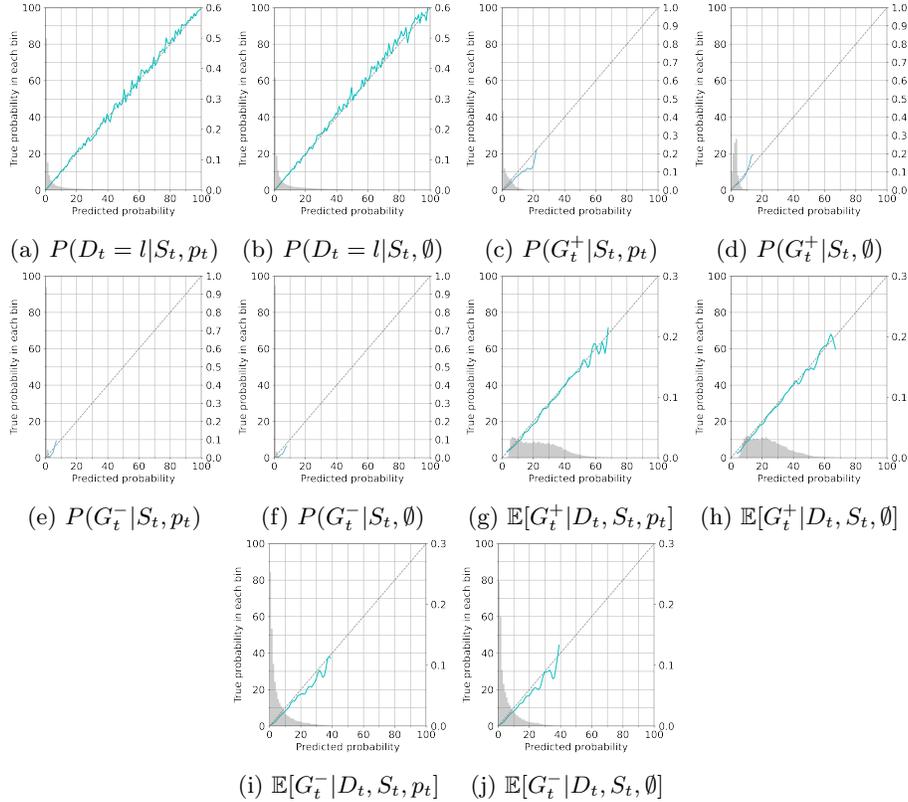


Fig. 3: Probability calibration plots for the pass selection (a and b), ball recovery probability (c and d), shot selection probability (e and f), expected reward (g and h), and expected risk (i and j) models, respectively with and without pressing-related features. Values on the x-axis represent the mean predicted probability by bin, among 100 equally-sized bins. The y-axis represents the mean observed outcome by bin. The histogram represents the percentage of examples in each bin relative to the total number of examples for each model. Kernel density estimation smoothing is applied among the bins in figure c – j and the probability calibration curve is not displayed for bins containing less than 20 examples.