Player Valuation in European Football

Edward Nsolo, Patrick Lambrix and Niklas Carlsson

Conference article

Cite this conference article as:

Nsolo, E., Lambrix, P., Carlsson, N. Player Valuation in European Football, In Brefeld, U., Davis, J., Van Haaren, J., Zimmermann, A. (eds), *Proceedings of the 5th Workshop on Machine Learning and Data Mining for Sports Analytics: co-located with 2018 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD 2018)*, Springer; 2018, pp. 42-54. ISBN: 9783030172732

DOI: https://doi.org/10.1007/978-3-030-17274-9_4

Lecture Notes in Computer Science, ISSN: 0302-9743, No. 11330

Copyright: Springer

The self-archived postprint version of this conference article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-153594>

Player valuation in European football

Edward Nsolo Patrick Lambrix Niklas Carlsson

Linköping University, Sweden

Abstract. As the success of a team depends on the performance of individual players, the valuation of player performance has become an important research topic. In this paper, we compare and contrast which attributes and skills best predict the success of individual players in their positions in five European top football leagues. Further, we evaluate different machine learning algorithms regarding prediction performance. Our results highlight features distinguishing top-tier players and show that prediction performance is higher for forwards than for other positions, suggesting that equally good prediction of defensive players may require more advanced metrics.

Keywords: Sports analytics · Data mining · Player valuation.

1 Introduction

The success of a football team depends a lot on the individual players making up that team. However, not all positions on a team are the same and there are significant differences in the style of game being played in different leagues. It is therefore important to take into account both the position and the league when evaluating individual players.

In this paper, we compare and contrast which attributes and skills best predict the success of individual players, within and across leagues and positions. First, we investigate which performance features are important in five European top leagues (Italy, Spain, England, Germany, and France) for defenders, midfielders, forwards, and goal keepers, respectively. Second, as part of our analysis, we evaluate different techniques for generating prediction models (based on different machine learning algorithms) for players belonging to the top segment of players. To capture further differentiation among the best players of a league, in our experiments, we have investigated top sets corresponding to 10%, 25% and 50% of the most highly ranked players in each considered category.

Our results provide interesting insights into what features may distinguish a top-tier player (in the top 10%, for example), from a good player (in the top 25%), or from an average player, within and across the leagues. The results also suggest that predicting performance may be easier for forwards than for other positions. Our work distinguishes itself from other work on player valuation or player performance, by working with tiers of players (in contrast to individual ratings) as well as by dealing with many skills (in contrast to some work with focus on a particular skill).

The remainder of the paper is organized as follows. Section 2 presents related work. Section 3 discusses the data sets and the data preparation. Sections 4 and 5 present the feature selection and prediction methods, respectively, and show and discuss the corresponding results. Finally, conclusions are presented in Section 6.

2 Related work

In many sports, work has started on valuation of player performance. For the sake of brevity, we address the related work in football.

Much work focuses on game play, including rating game actions [9, 28], pass prediction [18, 5, 7, 31, 13], shot prediction [25], expectation for goals given a game state [8], or more general game strategies [24, 3, 15, 11, 12, 4, 14, 1]. Other works try to predict the outcome of a game by estimating the probability of scoring for the individual goal scoring opportunities [10], by relating games to the players in the teams [26], or by using team rating systems [21].

Regarding player performance, in [29] a weighted +/- measure for player performance is introduced. The skill of field vision is investigated in [20] while in [32] the authors develop a method to, based on the current skill set, determine the skill set of a player in the future. The performance versus market value for forwards is discussed in [17]. Further, the influence of the heterogeneity of a team on team performance is discussed in [19].

3 Data collection and preparation

3.1 Data collection

Data was collected from five top European leagues, i.e., the highest English (English Premier League, EPL), Spanish (La Liga), German (BundesLiga), Italian (Serie A) and French (Ligue 1) leagues for the 2015-16 season. WhoScored (<https://www.whoscored.com>) was used as the main data provider. Most of its data is acquired from Opta (<http://www.optasports.com>) which is used by many secondary data providers. WhoScored and other secondary data providers use internal schemes developed by a group of soccer experts to rate player and team performances. The differences in rating players is relatively small. Therefore, we used the WhoScored rating as the gold standard in our experiments. The data contained information about 2,606 players. The list of attributes is given in Table 1. These are the attributes of WhoScored except for some where we aggregated the attributes, as for instance, goals represents goals by right foot, goals by left foot, goals by head and goals by other body parts.

We generated six groups of data sets, one group for each league as well as one for all leagues together. This allows us to find differences and commonalities between the different leagues. For each group, we divided the players with respect to their position on the field into 222 goalkeepers (GKs), 970 defenders (DFs), 1,109 midfielders (MDs), and 305 forwards (FWs). For the field players, players whose primary role and position on the pitch were defending and back were

Table 1. List of attributes. The defensive, offensive and passing categories are as they are defined in WhoScored. Note that some attributes are in several of these categories. We added the two other categories for the remaining attributes used in WhoScored.

Type	attributes
Identifying	player id, nationality, player name, age, position, height, weight, team name, league
Defensive	tackles, interceptions, offsides won, clearance, blocks, own goals, dribbles, fouls committed
Offensive	goals, assists, shots per game, key passes, fouled, offsides committed, dispossessed, bad control, dribbles
Passing	assists, key passes, passes per game, pass success rate, crosses, long balls, through balls
All	man of the match, fulltime, halftime, minutes played, aerals won, red cards, yellow cards
Class	rating

categorized as defenders. Players whose primary role and position on the pitch were playmaking and central were categorized as midfielders. Finally, players whose primary role and position on the pitch were attacking and front, were categorized as forwards.

3.2 Data preparation

For each data set representing a position (GK, DF, MD, FW) within a group (specific league or all leagues), we generated three final data sets representing the top 10%, 25% and 50% players for that position and group. We used Weka [16] with TigerJython in this phase.

We filled in some missing values in the data, such as nationality for some players, and used a uniform way to represent missing values over the data sets. We detected duplicates of players which transferred to other teams during the season and merged the data for these players. This affected 3 GKs, 20 DFs, 20 MDs, and 71 FWs. Each data set was normalized using the min-max method which transformed given numerical attribute values to a range of 0 to 1. The normalized value of value x is $(x - \min)/(\max - \min)$, where $\min$ and $\max$ are the minimum and maximum value for the attribute, respectively. The rating of WhoScored was used for the binary class attribute (deciding whether the player is in the top X% or not). The data was discretized to reflect this. We used SMOTE [6] to overcome the class imbalance for the 10% and 25% data sets. It is an oversampling technique that synthetically determines copies of the instances of the minority class to be added to the data set to match the quantity of instances of the majority class. We refer to these final data sets (after all preparation steps) as LEAGUE-POSITION-Xpc where LEAGUE has the values EPL, Bundesliga, La-Liga, Ligue-1, Serie-A and All; POSITION has the values GK, DF, MD, and FW; and X has the values 10, 25, 50 (e.g., EPL-GK-25pc).

4 Feature selection

For each data set we used two ways implemented in Weka [16], a filter method and a wrapper method, to select possible features that are important for player performance.

4.1 Filter method

With the filter method we computed absolute values of the Pearson correlation coefficients between each attribute and the class attribute. In our application we retained all attributes that had a Pearson correlation coefficient of at least 0.3.

The results for the combined leagues are shown in Table 2. For the results regarding the individual leagues, we refer to [27].

This method is fast and runs in milliseconds for the different cases.

4.2 Wrapper method

Wrapper methods use machine learning (ML) algorithms and evaluate the performance of the algorithms on the data set using different subsets of attributes. We used the Weka setting where we started from the empty set and used best-first search with backtracking after five consecutive non-improving nodes in the search tree. We used seven ML algorithms that can handle numeric, categorical, and binary attributes and that were chosen from different types; i.e., BayesNet (Bayesian nets), NaiveBayes (Naive Bayes classifier), Logistic (builds linear logistic regression models), IBk (k-nearest-neighbor), J48 (a C4.5 decision tree learner), Part (rules from partial decision trees), and RandomForest (random forests). The merit of a subset was evaluated using WrapperSubsetEval that uses the classifier with, in our case, five-fold cross-validation. Then we computed a support count for each attribute reflecting how often it occurred in a selected attribute subset computed by the different ML algorithms. Attributes that were selected at least twice were retained as important attributes.

The results for the data sets for the combined leagues are shown in Table 3. For the results for the individual leagues we refer to [27].

The run times for this method depend on the ML algorithms and the data sets ranging from below a minute to several hours (see [27]). Logistic and RandomForest were the slowest. For the Bayesian approaches there was not much difference between the different data sets.

4.3 Discussion

Filter method: As expected, the selected attributes did not include any of the identifying attributes such as team name, nationality, player name and league, which all received low correlation coefficients.

As we used the absolute value of the Pearson correlation coefficients, the impact of the selected attributes could be positive as well as negative on the

Table 2. Attributes filter method for combined leagues. Attributes in *italics* are in common with the other data sets for the same position (All-X-10pc, All-X-25pc, All-X-50pc). Attributes in **bold** are in common with the same data set for the wrapper method.

Selected attributes	Data set
minutes played, fulltime, clearance, dribble, yellow cards	All-GK-10pc
< none >	All-GK-25pc
man of the match	All-GK-50pc
<i>interceptions, man of the match, aerial won, crosses, tackles</i> , halftime, assists, fulltime, minutes played	All-DF-10pc
<i>man of the match, crosses, aerial won, interceptions</i> , minutes played, fulltime, tackles, goals, halftime, shots per game	All-DF-25pc
<i>crosses, fulltime</i> , minutes played, <i>interceptions, man of the match, aerial won, tackles, goals, through balls</i>	All-DF-50pc
<i>man of the match, key passes, fulltime, minutes played, shots per game, assists, crosses, goals, fouled, dispossessed</i> , halftime, long balls, through balls, bad control, tackles	All-MD-10pc
fulltime, minutes played, <i>man of the match, key passes, crosses, assists</i> , fouled, shots per game, goals, dispossessed , tackles, <i>halftime</i> , through balls, long balls, bad control, dribble	All-MD-25pc
fulltime, minutes played, crosses, key passes, assists, man of the match, fouled, tackles, shots per game , interceptions, yellow cards, <i>goals, through balls, dribble, dispossessed, bad control, long balls</i> , fouls committed, clearance	All-MD-50pc
<i>crosses, shots per game, fulltime, goals, man of the match, minutes played, assists, passes per game, key passes, fouled, aerial won, through balls, dispossessed, bad control, offsides committed, clearance, yellow cards, halftime, age, tackles, fouls committed</i>	All-FW-10pc
fulltime, minutes played, shots per game , <i>crosses, goals, passes per game, key passes, man of the match, assists, dispossessed, bad control, fouled, offsides committed, aerial won, through balls, tackles, clearance, yellow cards, fouls committed, age</i>	All-All-FW-25pc
<i>crosses, fulltime, minutes played, shots per game, key passes, passes per game, bad control, goals, dispossessed, fouled, aerial won, assists, offsides committed, through balls, fouls committed, man of the match, clearance, tackles, yellow cards, age</i>	All-FW-50pc

class variable. In our results most of these are considered positive (e.g., goals for forwards). We note that for the offensive and passing attributes the values are the highest for the FWs, while for the defensive attributes they are usually highest for the DFs.

The top 10 of selected attributes over all the different leagues [27] (i.e., they are selected for most combinations of league, position and top X%) contain attributes related to all player responsibilities, such as tackles for defense, shots per game and goals for offense, crosses for passing, and key passes and assists for both offense and passing. Further, there are other attributes such as man of the match, minutes played, full time and aeriels won. Identifying attributes as well as red cards, own goals, and offsides won are rarely selected.

For all leagues, fewer attributes are selected for GKs and DFs than for MDs and FWs. In Tables 2 and 3 we have marked the attributes that are in common

Table 3. Attributes wrapper method for combined leagues. Attributes in *italics* are in common with the other data sets for the same position (All-X-10pc, All-X-25pc, All-X-50pc). Attributes in **bold** are in common with the same data set for the filter method. The numbers in parentheses show the support counts. For attributes without a number the support count is 2.

Selected attributes	Data set
<i>player name</i> (5), player id (3), tackles, own goals, halftime, long balls	All-GK-10pc
<i>player name</i> (4), team name (3), player id (3) own goals, through balls	All-GK-25pc
man of the match (7), <i>player name</i> (3), aerial won, player id	All-GK-50pc
<i>interceptions</i> (4), player name (4), aerial won (3), man of the match (3), <i>shots per game</i> (3), tackles (3), <i>team name</i> (3), age, passes per game, blocks, red cards, yellow cards	All-DF-10pc
player name (5), man of the match (5), <i>shots per game</i> (4), tackles (4), <i>team name</i> (4), crosses (3), halftime (3), goals , own goals, disposessions, fulltime , offsides committed, <i>interceptions</i> , minutes played	All-DF-25pc
man of the match (7), <i>interceptions</i> (7), crosses (6), <i>shots per game</i> (5), tackles (4), assists (4), blocks (3), goals , <i>team name</i> , through balls , offsides committed, league	All-DF-50pc
player name (5), team name, fouled , own goals, man of the match	All-MD-10pc
man of the match (6), player name (4), halftime (4), assists (3), shots per game , crosses , pass success rate, league, red cards, goals , blocks, key passes	All-MD-25pc
fulltime (7), crosses (6), man of the match (4), key passes (3), tackles (3), shots per game (3), pass success rate, assists	All-MD-50pc
man of the match (5), player name (5), own goals (4), blocks (3), player id	All-FW-10pc
man of the match (7), player id (4), name (4), aerial won (3), shots per game (3), halftime (3), goals , offsides committed , age	All-FW-25pc
man of the match (6), crosses (5), shots per game (3), assists (3), <i>interceptions</i> (3), red cards, offsides won, fouls committed , weight, bad control	All-FW-50pc

for the same position in *italics*. In particular, for FWs and MDs many attributes are selected for each of the top 10%, top 25% and top 50% data sets, for DFs there are fewer common attributes while very few for GKs.

Interception is selected more often for Serie A than for other leagues and is important for all positions in Serie A. Through balls is selected less in La Liga than in other leagues. Height is only selected for top 10% GKs, DFs and FWs in the Bundesliga and top 10% GKs in the EPL. Tackles are selected for all DFs in the Bundesliga, Ligue 1 and Serie A, while only for some of the DF data sets for EPL and La Liga. Offsides committed is selected for all FWs in all leagues, except for the EPL where it is only selected for the top 50% data set.

Wrapper method: The subsets produced by WrapperSubsetEval had merit scores of over 68% for the data sets of the combined leagues and 77% for the data sets of the individual leagues. IBk, NaiveBayes and Logistic had on average the highest merit scores (of over 90%) across all data sets, except for some GK and top 50% data sets. The selected subsets for the top-10% data sets for all four categories of players for all individual and for the combined leagues had the highest merit scores while the selected subsets for the top-50% data sets

had the lowest merit scores. This might be caused by the amount of instances added when handling class imbalances. Further, the merit scores for the selected subsets for FWs were always higher than those for MDs, which in their turn were higher than for DFs and GKs.

In general, few selected attributes are common for the data sets related to the same position. For instance, for the Bundesliga there were no selected attributes in common for the different data sets at all. For the EPL, clearance was in common for GKs, aerials won and man of the match for DFs and man of the match for MDs. In contrast to the filter method, identifying attributes do appear in the selected lists.

Both: We note that the selected attributes for the combined leagues and the individual leagues are different. This suggests that players doing well in one league will not necessarily do well in another leagues and may reflect the fact that the playing style is different in different leagues.

A larger ratio of the selected attributes for the wrapper methods are in common with the selected attributes for the filter methods for the same data set. For the combined leagues the overlap is larger than for the individual leagues and the overlap is largest for DFs and MDs. For the individual leagues the Bundesliga has a small overlap which is mostly situated in the MD data sets.

5 Prediction

5.1 Methods

For each LEAGUE-POSITION-Xpc data set, based on the results of the feature selection procedure we created two data sets with the selected attributes from the wrapper and filter methods, respectively. Using Weka [16] we ran several ML algorithms: RandomForest, BayesNet, Logistic, DecisionTable (a decision table majority classifier), IBk, KStar (nearest neighbor with generalized distance), NaiveBayes, J48, Part, and ZeroR (predicts the majority class for nominals or the average value for numerics). ZeroR was used as a baseline. We split each data set into a training (66% of the instances) and a testing set (34 % of the instances). All experiments were run 10 times and we calculated averages for the performance values.

5.2 Results

As performance metrics we used the standard measures of accuracy, precision, recall, F1 score (or f-measure), and AUC-ROC. Given the true and false positives (TP and FP) and the true and false negatives (TN and FN), the accuracy is $(TP + TN) / (TP + TN + FP + FN)$, the precision is $TP / (TP + FP)$ and the recall is $TP / (TP + FN)$. The F1 score is a harmonic mean over precision and recall. AUC-ROC plots recall against precision.

Figures 1 and 2 show summary statistics for F1 scores and prediction performance, respectively, from full factor experiments, in which we evaluated every

Fig. 1. Distribution statistics when keeping the reported factor fixed and varying all other factors.

combination of (i) league or the combined leagues, (ii) position, (iii) top-set category, (iv) filter/wrapper selection, and (v) ML technique being applied. In total, this resulted in $6 \times 4 \times 3 \times 2 \times 10 = 1,440$ prediction evaluations (14,400 when taking into account 10 runs per configuration). Other and more detailed results are presented in an extended version [27].

5.3 Discussion

Figure 1 shows the distribution statistics for the individual leagues when keeping the reported factor fixed and varying all other factors. In particular, for each such factor and level, we present the maximum F1 score, the 90%-ile F1 score, the 75%-ile F1 score, the median F1 score (equal to the 50%-ile score), and the average F1 score. Here, the first three factors (i.e., the left hand side of Figure 1) allow us to compare and contrast the prediction scores obtained for different (i) leagues, (ii) player positions, and (iii) top-set categories. These results suggest a clear ordering in the prediction scores based on position and top set. For example, FWs have the highest F1 scores, MDs the second highest, DFs the second lowest, and GKs the lowest F1 scores. This suggests that ML may be more successfully applied to evaluate more offensive positions, but may be less successful to predict the skills of more defensive positions (with GKs the other extreme of the spectrum). The techniques also appear much better at predicting players in the top-10 set than predicting players in the top-50 set (again with a clear ordering of the three sets). Smaller differences are observed across the leagues, although the EPL and Bundesliga appear to be the two leagues for which it is easiest to use basic ML to predict top players.

The right hand side of Figure 1 compares the prediction success of the different methods. In general, the wrapper method is more successful than the filter method, and BayesNet and RandomForest provide the highest F1 scores across the distribution metrics. All methods significantly outperform the baseline.

Finally, we look closer at how well the above methods (the baseline excluded) can predict the top X players of each position in each of the leagues. Figure 2 presents these results. Here, we show results for the best predictor results (max)

Fig. 2. F1 scores of the top-X player sets (X=10, 25, 50%) of each position and league.

for each top-X set, where X=10%, 25%, and 50%, and the corresponding medians (when considering all filter/wrapper combinations with the different ML techniques). Again, the top-10% set is much easier to predict across the leagues (with medians above 0.9 across all leagues and positions, except GKs in France). As the top sets become larger, the F1 scores decrease, and again the more offensive positions obtain substantially higher prediction scores across the leagues.

6 Conclusion

In this paper we generated attributes and skills sets that best predict the success of individual players in their positions in five European top leagues. In contrast to other work we focused on the top tiers of players (top 10%, 25% and 50%). Further, we evaluated different ML algorithms on their performance of predicting the tiers using the generated attributes. For the sake of brevity we have shown some of the results in the paper while for more results we refer to [27]. Among other things, our prediction results show (i) a clear ordering in the prediction scores based on position (e.g., F1 of FW > F1 of MD > F1 of DF > F1 of GK) and top set (e.g., F1 of top 10% > F1 of top 25% > F1 of top 50%), (ii) that basic ML techniques are most successful to predict top players in EPL and Bundesliga (although good performance across the leagues), (iii) that the wrapper method is more successful than the filter method, and (iv) that BayesNet and RandomForest provide the highest F1 scores of the considered ML techniques.

One limitation of the approach is that the method is based on a ranking of experts, which is the current state of expertise. However, as in the case of baseball, the opinions of which properties constitute a 'good' player may change [22]. Future work could consider longitudinal drift in the rankings. It is also interesting to investigate the correlation of the rankings with the players' market value. Another limitation is that the approach does not take the quality of the team mates into account (although *team name* was sometimes a selected attribute). Although important, we are not aware of work that takes this explicitly into account. However, there is work on related questions such as team performance [5, 7, 21, 26, 19] and, in other sports, the performance of pairs or triples of players

(e.g., [30, 23] for ice hockey and [2] for basketball). Another direction for future work is to apply the methodology only on game-related attributes. Some of the identifying attributes may not be that interesting for teams for identifying possible future players for the team. Further, for some attributes, such as shots per game, we may investigate using normalized values based on actual play time. It could also be interesting to look at different tiers and to develop more advanced features and metrics for the more defensive positions, aiming to provide equally good prediction for DF and GK as for FW, for example.

References

1. Andrienko, G., Andrienko, N., Budziak, G., Dykes, J., Fuchs, G., von Landesberger, T., Weber, H.: Visual analysis of pressure in football. *Data Mining and Knowledge Discovery* **31**(6), 17931839 (2017), doi:10.1007/s10618-017-0513-2
2. Ayer, R.: Big 2's and big 3's: Analyzing how a team's best players complement each other. In: MIT Sloan Sports Analytics Conference (2012)
3. Bialkowski, A., Lucey, P., Carr, P., Yue, Y., Sridharan, S., Matthews, I.: Large-Scale Analysis of Soccer Matches Using Spatiotemporal Tracking Data. In: Kumar, R., Toivonen, H., Pei, J., Huang, J.Z., Wu, X. (eds.) *Proceedings of the 2014 IEEE International Conference on Data Mining*. pp. 725–730 (2014)
4. Bojinov, I., Bornn, L.: The Pressing Game: Optimal Defensive Disruption in Soccer. In: 10th MIT Sloan Sports Analytics Conference (2016)
5. Brandt, M., Brefeld, U.: Graph-based Approaches for Analyzing Team Interaction on the Example of Soccer. In: Davis, J., van Haaren, J., Zimmermann, A. (eds.) *Proceedings of the 2nd Workshop on Machine Learning and Data Mining for Sports Analytics*. CEUR Workshop Proceedings, vol. 1970, pp. 10–17 (2015)
6. Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research* **16**, 321–357 (2002), doi: 10.1613/jair.953
7. Cintia, P., Rinzivillo, S., Pappalardo, L.: Network-based Measures for Predicting the Outcomes of Football Games. In: Davis, J., van Haaren, J., Zimmermann, A. (eds.) *Proceedings of the 2nd Workshop on Machine Learning and Data Mining for Sports Analytics*. CEUR Workshop Proceedings, vol. 1970, pp. 46–54 (2015)
8. Decroos, T., Dzyuba, V., Van Haaren, J., Davis, J.: Predicting Soccer Highlights from Spatio-Temporal Match Event Streams. In: *Proceedings of the 31st AAAI Conference on Artificial Intelligence*. pp. 1302–1308 (2017)
9. Decroos, T., Van Haaren, J., Dzyuba, V., Davis, J.: STARSS: A Spatio-Temporal Action Rating System for Soccer. In: Davis, J., Kaytoue, M., Zimmermann, A. (eds.) *Proceedings of the 4th Workshop on Machine Learning and Data Mining for Sports Analytics*. CEUR Workshop Proceedings, vol. 1971, pp. 11–20 (2017)
10. Eggels, H., van Elk, R., Pechenizkiy, M.: Explaining Soccer Match Outcomes with Goal Scoring Opportunities Predictive Analytics. In: van Haaren, J., Kaytoue, M., Davis, J. (eds.) *Proceedings of the 3rd Workshop on Machine Learning and Data Mining for Sports Analytics*. CEUR Workshop Proceedings, vol. 1842 (2016)
11. Fernando, T., Wei, X., Fookes, C., Sridharan, S., Lucey, P.: Discovering Methods of Scoring in Soccer Using Tracking Data . In: Lucey, P., Yue, Y., Wiens, J., Morgan, S. (eds.) *Proceedings of the 2nd KDD Workshop on Large Scale Sports Analytics* (2015)

12. Gyarmati, L., Anguera, X.: Automatic Extraction of the Passing Strategies of Soccer Teams. In: Lucey, P., Yue, Y., Wiens, J., Morgan, S. (eds.) *Proceedings of the 2nd KDD Workshop on Large Scale Sports Analytics* (2015)
13. Gyarmati, L., Stanojevic, R.: QPass: a Merit-based Evaluation of Soccer Passes. In: Lucey, P., Yue, Y., Wiens, J., Morgan, S. (eds.) *Proceedings of the 3rd KDD Workshop on Large Scale Sports Analytics* (2016)
14. Haaren, J.V., Davis, J., Hannosset, S.: Strategy Discovery in Professional Soccer Match Data. In: Lucey, P., Yue, Y., Wiens, J., Morgan, S. (eds.) *Proceedings of the 3rd KDD Workshop on Large Scale Sports Analytics* (2016)
15. Haaren, J.V., Dzyuba, V., Hannosset, S., Davis, J.: Automatically Discovering Offensive Patterns in Soccer Match Data. In: Fromont, E., De Bie, T., van Leeuwen, M. (eds.) *Proceedings of International Symposium on Intelligent Data Analysis. LNCS*, vol. 9385, pp. 286–297 (2015)
16. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA Data Mining Software: An Update. *SIGKDD Explorations* **11**(1), 10–18 (2009), doi: 10.1145/1656274.1656278
17. He, M., Cachucho, R., Knobbe, A.: Football Player’s Performance and Market Value. In: Davis, J., van Haaren, J., Zimmermann, A. (eds.) *Proceedings of the 2nd Workshop on Machine Learning and Data Mining for Sports Analytics. CEUR Workshop Proceedings*, vol. 1970, pp. 87–95 (2015)
18. Horton, M., Gudmundsson, J., Chawla, S., Estephan, J.: Classification of Passes in Football Matches using Spatiotemporal Data. In: Lucey, P., Yue, Y., Wiens, J., Morgan, S. (eds.) *Proceedings of the 1st KDD Workshop on Large Scale Sports Analytics* (2014)
19. Ingersoll, K., Malesky, E., Saiegh, S.M.: Heterogeneity and team performance: Evaluating the effect of cultural diversity in the world’s top soccer league. *Journal of Sports Analytics* **3**(2), 67–92 (2017), doi: 10.3233/JSA-170052
20. Jordet, G., Bloomfield, J., Heijmerikx, J.: The hidden foundation of field vision in English Premier League (EPL) soccer players. In: *7th MIT Sloan Sports Analytics Conference* (2013)
21. Lasek, J.: EURO 2016 Predictions Using Team Rating Systems. In: van Haaren, J., Kaytoue, M., Davis, J. (eds.) *Proceedings of the 3rd Workshop on Machine Learning and Data Mining for Sports Analytics. CEUR Workshop Proceedings*, vol. 1842 (2016)
22. Lewis, M.: *Moneyball: The Art of Winning an Unfair Game*. W. W. Norton & Company (2003), isbn: 978-0-393-05765-2
23. Ljung, D., Carlsson, N., Lambrix, P.: Player pairs valuation in ice hockey. In: Brefeld, U., Davis, J., van Haaren, J., Zimmermann, A. (eds.) *Proceedings of the 5th Workshop on Machine Learning and Data Mining for Sports Analytics* (2018)
24. Lucey, P., Bialkowski, A., Carr, P., Foote, E., Matthews, I.: Characterizing Multi-Agent Team Behavior from Partial Team Tracings: Evidence from the English Premier League. In: *26th AAAI Conference on Artificial Intelligence*. pp. 1387–1393 (2012)
25. 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: *9th MIT Sloan Sports Analytics Conference* (2015)
26. Maystre, L., Kristof, V., Ferrer, A.J.G., Grossglauser, M.: The Player Kernel: Learning Team Strengths Based on Implicit Player Contributions. In: van Haaren, J., Kaytoue, M., Davis, J. (eds.) *Proceedings of the 3rd Workshop on Machine Learning and Data Mining for Sports Analytics. CEUR Workshop Proceedings*, vol. 1842 (2016)

27. Nsolo, E., Lambrix, P., Carlsson, N.: Player valuation in European football (extended version) (2018), <https://www.ida.liu.se/research/sportsanalytics/projects/conferences/MLSA18-soccer>
28. Sarkar, S., Chakraborty, S.: Pitch actions that distinguish high scoring teams: Findings from five European football leagues in 2015-16 . *Journal of Sports Analytics* **4**(1), 1–14 (2018), doi: 10.3233/JSA-16161
29. Schultze, S.R., Wellbrock, C.M.: A weighted plus/minus metric for individual soccer player performance. *Journal of Sports Analytics* **4**(2), 121–131 (2018), doi: 10.3233/JSA-170225
30. Thomas, A., Ventura, S.L., Jensen, S., Ma, S.: Competing process hazard function models for player ratings in ice hockey. *The Annals of Applied Statistics* **7**(3), 1497–1524 (2013)
31. Vercruyssen, V., Raedt, L.D., Davis, J.: Qualitative Spatial Reasoning for Soccer Pass Prediction. In: van Haaren, J., Kaytoue, M., Davis, J. (eds.) *Proceedings of the 3rd Workshop on Machine Learning and Data Mining for Sports Analytics*. CEUR Workshop Proceedings, vol. 1842 (2016)
32. Vroonen, R., Decroos, T., Haaren, J.V., Davis, J.: Predicting the Potential of Professional Soccer Players. In: Davis, J., Kaytoue, M., Zimmermann, A. (eds.) *Proceedings of the 4th Workshop on Machine Learning and Data Mining for Sports Analytics*. CEUR Workshop Proceedings, vol. 1971, pp. 1–10 (2017)