

STARSS: A Spatio-Temporal Action Rating System for Soccer

Tom Decroos¹, Jan Van Haaren², Vladimir Dzyuba¹, and Jesse Davis¹

¹ KU Leuven, Department of Computer Science,
Celestijnenlaan 200A, 3001 Leuven, Belgium

² SciSports, Hengelosestraat 500, 7251 AN Enschede, The Netherlands

Abstract. An important task in sports analytics is to devise player-performance metrics that allow managers to take better-informed decisions. While several such metrics have been proposed for baseball, basketball, and ice hockey, this task has virtually remained unexplored to date for soccer. This paper presents an approach for automatically rating the actions performed by soccer players based on historical match data. The approach considers all player actions that contribute to a team’s offensive output and accounts for the context of the actions.

1 Introduction

One of the key objectives in sports analytics is to quantify player and team performances as objectively as possible. Ideally, each player’s contributions to his or her team could be summarized in a single number that allows managers to take better-informed decisions. There are two approaches to compute such a number. The first style of approach to this task focuses on aggregating a variety of statistics into a single number. Typically, these metrics consider statistics that can be derived from a boxscore. Two well-known examples of this approach are Wins Above Replacement (WAR)³ in baseball and the Player Efficiency Rating (PER)⁴ in basketball. The second type of approach tries to assign values to the actions performed during a match. These approaches are based on the analysis of eventstream data that describes the actions performed in a match, possibly in conjunction with optical tracking data. Examples of this are the Expected Possession Value (EPV) model [4] for basketball and a conceptually similar model for ice hockey [10]. In soccer, however, the task of devising objective player-performance metrics has been almost unexplored.

Building an objective player-performance metric for soccer players is an extremely challenging task. Due to the low-scoring nature of the game, distinguishing between “successful” and “unsuccessful” actions performed by players is not straightforward. We believe that a credible and reliable player-performance metric for soccer needs to satisfy at least the following two criteria. First, the metric should not be biased towards a particular style of play. For example, players in

³ http://www.baseball-reference.com/about/war_explained_position.shtml

⁴ <http://www.basketball-reference.com/about/per.html>

possession-based teams should not be favored over players in counter-attacking teams. Second, the metric should account for the spatial context at the time of each action. For example, a key pass in the final third of the pitch should be valued higher than a pass in midfield.

In this paper, we present STARSS (**S**patio-**T**emporal **A**ction **R**ating **S**ystem for **S**occer), which leverages historical match data to assign a rating to the actions (e.g., a pass or a shot) performed by the players in a match. For a given match, the presented approach proceeds in three steps. First, the approach splits the match into phases, which are uninterrupted sequences of actions where one team is in possession of the ball. Second, it assigns a phase rating to each phase based on historical match data. The higher the assigned rating, the more likely that the phase will end in a goal. Hence, our approach focuses on the actions that contribute to the offensive output of the team. Third, the approach distributes the phase rating across the individual actions that constitute the phase.

We use STARSS to rate players and teams in individual matches as well as throughout the course of a season. We present the top-15 players for the 2015/2016 season in the English Premier League, the German Bundesliga, and the Spanish La Liga. We find that wingers and attacking midfielders tend to contribute more to a team’s STARSS rating than strikers, that five of the top-15 players in the German Bundesliga play for FC Bayern Munich, and that Lionel Messi is the best player in the world.

2 Dataset

Our dataset consists of play-by-play data for the English Premier League, German Bundesliga, and Spanish La Liga for the 2012/2013 through 2015/2016 seasons. Our dataset comprises 4253 matches, 7,569,802 game events, 110,290 shots, and 11,842 goals.

The data for each match consist of a stream of events. For each event, the following information is available: the type (e.g., a foul or cross), a timestamp, the player involved, the team involved, and the location on the pitch (i.e., the x- and y-coordinate of the ball). Depending on the type of the event, additional information is available. For example, the end location for a pass or the outcome (e.g., off target, on target, or goal) for a shot.

The presented approach rates players by assessing the “actions” they perform on the pitch. We define an “action” as an event performed by a player who either is in possession of the ball or attempts to gain possession of the ball. The set of considered actions includes, among others, passes, dribbles, crosses, shots, interceptions, and tackles.

3 Approach

This section introduces STARSS, which is an approach for automatically rating the actions performed by soccer players. Unlike traditional approaches,⁵ which

⁵ For instance, the approaches used by WhoScored.com and Squawka.com.

simply compute a weighted sum of the frequencies of a hand-picked set of actions, our approach accounts for the spatio-temporal context in which the actions were performed. More specifically, our approach leverages the outcomes of similar actions that were performed in similar circumstances in the past to assess the value of a particular action to the team.

To rate the player actions in a given match, our approach performs the following three steps. First, it splits the match into phases of related actions. Second, it assigns *phase ratings* to the resulting phases based on historical match data. Third, it assigns *action ratings* to the individual actions that constitute the phases.

3.1 Splitting matches into phases

We view a match as a sequence $M = (a_1, \dots, a_n) = (a_i)_{i=1}^n$, where each a_i is an action performed by a player and n is the total number of actions. We start a new phase each time possession switches from one team to the other or too much time (i.e., 10 seconds) has passed between consecutive actions (Figure 1). For example, when the ball goes out of play for a throw-in or corner kick, when a goal is scored, or when a free kick is awarded.




Fig. 1. A sequence of actions being split in coherent phases.

This way, a match M is split into subsequences $(P_j)_{j=1}^m$, where each P_j is a phase and m is the total number of phases. Each phase P_j is a subsequence of consecutive actions $(a_k)_{k=j_1}^{j_2}$ appearing in the sequence M .

3.2 Rating phases

We assign a rating to a given phase P in two steps as follows. In the first step, we search the k most-similar phases in terms of their spatial location on the pitch in historical match data. This historical match data is split up in phases using the same procedure as in Section 3.1. To measure the similarity between

two phases P and P' , we employ dynamic time warping (DTW) [8] as it does not require that two sequences have the same length and is insensitive to minor mismatches. Specifically, we use a multivariate variant of dynamic time warping:

$$d(P, P') = \sqrt{\sum_{i=1}^2 DTW_i(P, P')^2} \quad (1)$$

where DTW_i is the DTW-based similarity of P and P' in the i^{th} dimension. Here, we only consider the x and y coordinates of the individual actions that make up the phases P and P' . The core idea behind this approach is to reward teams and players for phases that get the ball into dangerous places, even if the phases do not lead to a shot. Furthermore, we are less interested in which action is used to get the ball into the location, as we do not want to favor one style of play over another. Most existing approaches to rate phases in soccer (e.g., [5, 3, 6]) only take phases that lead to a goal or a shot into account.

In the second step, we compute the *phase rating* as the proportion of similar phases that end in a goal:

$$phase_rating(P) = \frac{\sum_{P' \in NN_k(P)} \mathbb{1}_{goal}(P')}{k} \quad (2)$$

where $\mathbb{1}_{goal}(P)$ is an indicator function that is 1 if P ends in a goal and 0 if it does not and $NN_k(P)$ is the set of the k most similar phases to P (i.e., the Nearest Neighbors) according to distance function $d(P, P')$ in Equation 1. For example, if we want to rate a phase P using $k = 100$ similar phases and 13 out of the 100 most similar phases end in a goal, then the rating of P amounts to 0.13.

3.3 Rating actions

Assume we are given a phase $P = (a_i)_{i=k}$ with a_k the first action of the phase and a_l the last action of the phase. We then assign a rating to each action a_i in P in two steps. We first compute an action weight w_i that indicates the relevance of the action a_i to the phase P , and then compute the *action rating* by multiplying the rating for the phase P with the normalized action weight w'_i .

We use an exponential-decay-based approach to compute the action weights. We consider actions that happen at the end of a phase more important than actions that happen at the start of a phase. Hence, we assign low weights to actions that happen at the start and high weights to actions that happen at the end. Intuitively, this makes sense as the last few actions in a phase have the highest influence on its outcome. We compute the action weights in reverse order, starting with the last action a_l and working our way down to the first action a_k , using the following formula:

$$w_i = \begin{cases} 1 & \text{if } i = l, \\ (1 - \lambda) \times w_{i+1} & \text{if } k \leq i < l \end{cases} \quad (3)$$

where λ is a user-specified parameter.

We normalize the weights such that they sum to one, which means that the phase rating is completely distributed across the individual actions.

$$w'_i = \frac{w_i}{\sum_{j=k}^l w_j} \quad (4)$$

Finally, the rating for an action $a_i \in P$ is computed by multiplying the phase rating with the normalized action weight.

$$action_rating(a_i) = phase_rating(P) \times w'_i \quad (5)$$

4 Methodology

This section explains how we leverage the proposed STARSS approach to address the case studies presented in the next section.

We use STARSS to rate players in individual matches as well as throughout the course of an entire season as follows. For a given player, we first sum the ratings for the actions performed by the player during a match or season, then divide this sum by the total number of minutes that the player played, and finally multiply by 90 to obtain a rating normalized per 90 minutes. Similarly, we also rate teams in individual matches and over the course of a season.

We respect the chronological order of the matches when rating player actions. To rate a player or team in a particular match, we only leverage matches from the same league that had already been played at the time of that match to discover similar phases. For example, to rate the Leicester City players in their 1-3 win over Manchester City on 6 February 2015, we only consider the Premier League matches in our dataset that were played before that date.

For each of the case studies, we set the parameters $k = 100$ and λ for computing the action weights to 0.25 based on domain knowledge and an empirical analysis of the available data.

5 Case studies

This section presents case studies that illustrate the utility of the presented STARSS approach. Concretely, we address the following two questions:

1. Can STARSS identify the top-performing players in a league?
2. Can STARSS identify the top-performing players in a match?

5.1 Can STARSS identify the top-performing players in a league?

To answer this question, we compute the player ratings for all players that played at least 10 hours in the 2015/2016 season of the English Premier League, German Bundesliga, and Spanish La Liga. We present the top-15 players for each league

in tables 1, 2, and 3, respectively. Note that these rankings differ from simply ranking players based on a combination of goals and assists per 90 minutes, and hence they provide insight beyond using these traditional metrics. Alexis Sánchez was a key player for championship contenders Arsenal and tops the Premier League ranking. Zlatko Junuzović, whose assists were instrumental for Werder Bremen in their battle against relegation, is the top-ranked player in the Bundesliga. Lionel Messi, who helped Barcelona claim the league title, tops the La Liga ranking. Unsurprisingly, the five-time FIFA Ballon d’Or winner is also the top-ranked player across the three leagues. Unlike the Premier League and Bundesliga, the La Liga ranking exhibits a clear gap between Lionel Messi and the rest, indicating that the Argentine forward is a class apart as was also suggested by earlier work [7].

The Premier League ranking suggests that Arsenal’s offensive compartment excelled in the 2015/2016 season. The Gunners, who eventually finished second, have four players in the top 15 with Santi Cazorla, Mesut Özil, and Olivier Giroud alongside top-ranked Alexis Sánchez. In contrast, surprise champions Leicester City have not a single player in the top 15, despite the fact that Riyad Mahrez won the Player’s Player of the Year award and Jamie Vardy won the Premier League Player of the Season Award. These players were ranked first and second in the league according to summing total goals and assists. Traditional metrics like expected-goals indicated that Leicester were hugely over-performing last season, that is, the results were much better than the underlying numbers. Additionally, Leicester were also awarded a remarkably high number of penalties (13 in total, while the league average was 4.5 league average).

The Bundesliga ranking clearly shows FC Bayern Munich’s superiority in the 2015/2016 season. Although Werder Bremen’s free-kick specialist Zlatko Junuzović tops the ranking, the eventual champions have five players in the top 15. The La Liga ranking sees most of the usual suspects near the top of the ranking with Lionel Messi (Barcelona), Neymar (Barcelona), Cristiano Ronaldo (Real Madrid), and Gareth Bale (Real Madrid) occupying the first four positions.

These rankings also indicate several highly-ranked players who made moves to larger clubs following the season. These include Ilkay Gundogan and Nolito, who both transferred to Manchester City, and Henrikh Mkhitaryan, who moved to Manchester United.

5.2 Can STARSS identify the top-performing players in matches?

To answer this question, we perform three steps. First, we compute the rating for each player in each match in the 2015/2016 season of the English Premier League, German Bundesliga, and Spanish La Liga. Second, we compute the team rating for each team in each match by summing the individual player ratings. Third, we compute each player’s share of the team rating.

Figure 2 shows the shares of the player ratings in the team ratings for El Clásico, a match between fierce rivals FC Barcelona and Real Madrid in La Liga, on 2 April 2016. Unsurprisingly, Lionel Messi, Luis Suarez, and Neymar occupy

Rank	Team	Player	Rating per 90	Minutes played	Goals per 90	Assists per 90
1	Arsenal	Alexis Sanchez	0.289	2446	0.478	0.147
2	West Ham United	Dimitri Payet	0.279	2571	0.315	0.420
3	West Ham United	Mauro Zarate	0.262	790	0.342	0.000
4	Chelsea	Willian	0.249	2749	0.164	0.196
5	Liverpool	Philippe Coutinho	0.244	2003	0.359	0.225
6	Arsenal	Santi Cazorla	0.240	1292	0.000	0.209
7	Arsenal	Mesut Ozil	0.240	3047	0.177	0.561
8	Sunderland	Wahbi Khazri	0.240	1077	0.167	0.084
9	Aston Villa	Rudy Gestede	0.237	1657	0.272	0.109
10	Manchester City	Kevin De Bruyne	0.233	2003	0.315	0.404
11	Tottenham Hotspur	Christian Eriksen	0.231	2938	0.184	0.398
12	Arsenal	Olivier Giroud	0.231	2424	0.594	0.223
13	Liverpool	Christian Benteke	0.229	1518	0.474	0.178
14	Tottenham Hotspur	Erik Lamela	0.228	2383	0.189	0.340
15	Manchester City	David Silva	0.222	1800	0.100	0.550

Table 1. Top 15 players in the 2015/2016 English Premier League season.

Rank	Team	Player	Rating per 90	Minutes played	Goals per 90	Assists per 90
1	Zlatko Junuzovic	SV Werder Bremen	0.271	2352	0.153	0.383
2	Ilkay Gundogan	Borussia Dortmund	0.264	1993	0.045	0.135
3	Alexandru Maxim	VfB Stuttgart	0.263	1047	0.086	0.430
4	Tobias Werner	FC Augsburg	0.258	698	0.000	0.129
5	Sandro Wagner	SV Darmstadt 98	0.257	2476	0.436	0.145
6	Henrikh Mkhitaryan	Borussia Dortmund	0.257	2578	0.384	0.524
7	Thiago Alcantara	FC Bayern Munchen	0.255	1636	0.110	0.165
8	Franck Ribery	FC Bayern Munchen	0.253	681	0.264	0.396
9	Robert Lewandowski	FC Bayern Munchen	0.252	2654	0.950	0.068
10	Thomas Muller	FC Bayern Munchen	0.250	2343	0.576	0.192
11	Arjen Robben	FC Bayern Munchen	0.248	1101	0.245	0.082
12	Dario Lezcano	FC Ingolstadt 04	0.247	1245	0.145	0.072
13	Hakan Calhanoglu	Bayer 04 Leverkusen	0.246	2263	0.080	0.199
14	Daniel Didavi	VfB Stuttgart	0.245	2432	0.407	0.111
15	Daniel Ginczek	VfB Stuttgart	0.244	630	0.286	0.143

Table 2. Top 15 players in the 2015/2016 German Bundesliga season.

Rank	Team	Player	Rating per 90 Minutes played	Goals per 90	Assists per 90
1	Lionel Messi	Barcelona	0.387	2728	0.759
2	Neymar	Barcelona	0.339	3057	0.559
3	Cristiano Ronaldo	Real Madrid	0.320	3183	0.820
4	Gareth Bale	Real Madrid	0.310	1738	0.984
5	Duda	Malaga	0.304	1061	0.085
6	Nolito	Celta de Vigo	0.297	2472	0.400
7	James Rodriguez	Real Madrid	0.286	1516	0.416
8	Yevhen Konoplyanka	Sevilla	0.273	1612	0.167
9	Ever Banega	Sevilla	0.266	1725	0.209
10	Isco	Real Madrid	0.251	1826	0.148
11	Luis Suarez	Barcelona	0.247	3150	1.057
12	Angel Correa	Atletico de Madrid	0.247	948	0.475
13	Jese	Real Madrid	0.243	827	0.544
14	Orellana	Celta de Vigo	0.238	2915	0.185
15	Saul Berjon	Eibar	0.238	1400	0.064

Table 3. Top 15 players in the 2015/2016 Spanish La Liga season.

the first three spots for FC Barcelona, while Gareth Bale, Cristiano Ronaldo, and Karim Benzema occupy the first three spots for Real Madrid.

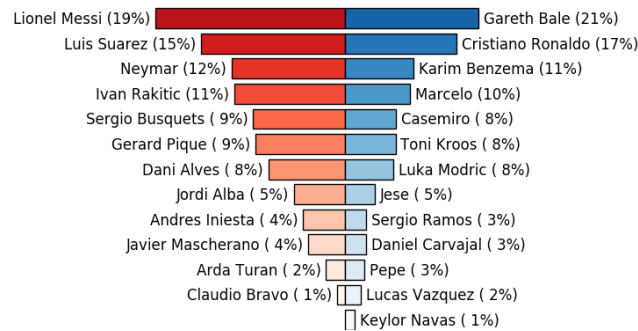


Fig. 2. The shares of the player ratings in the team ratings for El Clásico, a match between fierce rivals Real Madrid and FC Barcelona in La Liga, on 2 April 2016. FC Barcelona is shown at the left, while Real Madrid is shown at the right.

6 Related work

This section discusses related work in soccer as well as other sports.

Our approach is related to the work on expected-goals models, which have been a hot topic in the soccer-analytics community for the past few years.

Expected-goals models aim to objectively quantify the quality of goal attempts and several different models have been proposed in recent years [5, 3, 6, 1]. However, our work differs from these existing approaches in two crucial aspects. First, our approach is not restricted to shots and rates players based on all actions contributing to the team’s offensive output. Second, our approach explicitly takes the spatio-temporal context of the actions into account, as suggested by [2].

There are a number of websites such as WhoScored.com and Squawka.com that provide player ratings for soccer on a match-by-match basis. These websites have hand-crafted formulas that simply compute a weighted sum of frequencies for a hand-picked set of actions (e.g., shots, tackles, etc.). The weights associated with each action are set by hand according to domain knowledge. The importance of some of the defensive statistics such as the number of tackles is debatable. A high percentage of successful tackles is often considered a good thing but can also be the result of poor positioning. Our method differs from these approaches in that we avoid hand-crafting and use an automated data-driven approach to assign rankings. Furthermore, we consider the spatio-temporal context in which the actions were performed which the hand-crafted models ignore.

While virtually unexplored to date for soccer, the task of objectively quantifying player actions has been investigated for other sports. [4] propose the Expected Possession Value (EPV) model for basketball, which estimates the number of points a team is expected to score during a possession. uses a multiresolution semi-Markov stochastic model that defines a probability distribution over what the ballhandler is likely to do next, given the spatial configuration of the players on the court. Hence, this approach requires optical tracking data for all players. [10] introduces a conceptually similar model for ice hockey. They note that valuing actions can be posed within a reinforcement learning framework, which is challenging as in sports there is only access to a fixed data set, not a dynamic environment in which we can run new trials. Their approach considered a discrete state space and ignored locational information, which is highly important in soccer. [9] assesses the offensive productivity of hockey players based on the context in which they score goals. [11] introduces the Total Hockey Rating which goes beyond shots and goals to rate hockey players by taking all game events into account.

7 Conclusions

This paper introduces STARSS, which is an approach for automatically rating the actions performed by soccer players. Viewing a soccer match as a sequence of actions performed by players, the approach proceeds in three steps to rate these actions. First, it splits the match into phases of related actions. Second, it assigns a rating to each phase, indicating how likely it is that the phase will end in a goal. Third, it distributes the assigned rating over the individual actions that constitute the phase.

Unlike most existing approaches for rating soccer players, our approach goes beyond shots and goals. It considers all the actions that contribute to a team’s

offensive output and accounts for the spatio-temporal context in which these actions were performed. Several case studies show that our approach is able to identify top-performing players in individual matches as well as throughout the course of an entire season.

Acknowledgements

Tom Decroos is supported by the KU Leuven Research Fund (C22/15/015) and FWO-Vlaanderen (G.0356.12). Jan Van Haaren was supported by the Agency for Innovation by Science and Technology in Flanders (IWT). Vladimir Dzyuba is supported by FWO-Vlaanderen. Jesse Davis is partially supported by the KU Leuven Research Fund (C22/15/015) and FWO-Vlaanderen (G.0356.12, SBO-150033).

References

- [1] Aalbers, B.: Expected Goals Model: An Alternative Approach (2016), <http://www.scisports.com/news/2016/expected-goals-model-an-alternative-approach>
- [2] Altman, D.: Beyond Shots: A New Approach to Quantifying Scoring Opportunities (2015), <http://www.northyardanalytics.com/Dan-Altman-NYA-OptaPro-Forum-2015.pdf>
- [3] Caley, M.: Premier League Projections and New Expected Goals (2015), <http://cartilagefreecaptain.sbnation.com/2015/10/19/9295905/premier-league-projections-and-new-expected-goals>
- [4] Cervone, D., D'Amour, A., Bornn, L., Goldsberry, K.: POINTWISE: Predicting Points and Valuing Decisions in Real Time with NBA Optical Tracking Data. In: Proceedings of the 8th MIT Sloan Sports Analytics Conference, Boston, MA, USA. vol. 28 (2014)
- [5] IJtsma, S.: A Close Look at My New Expected Goals Model (2015), <http://11tegen11.net/2015/08/14/a-close-look-at-my-new-expected-goals-model/>
- [6] 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: Proc. of MIT Sloan Sports Analytics Conference (2014)
- [7] Morris, B.: Lionel Messi is Impossible (2014), <http://fivethirtyeight.com/features/lionel-messi-is-impossible/>
- [8] Müller, M.: Dynamic Time Warping, chap. 4, pp. 69–84. Information Retrieval for Music and Motion, Springer (2007)
- [9] Pettigrew, S.: Assessing the Offensive Productivity of NHL Players Using In-game Win Probabilities (2015), <http://www.sloansportsconference.com/wp-content/uploads/2015/02/SSAC15-RP-Finalist-Assessing-the-offensive-productivity-of-NHL-players2.pdf>
- [10] Routley, K., Schulte, O.: A Markov Game Model for Valuing Player Actions in Ice Hockey. In: Proc. of UAI. pp. 782–791 (2015)
- [11] Schuckers, M., Curro, J.: Total Hockey Rating (THoR): A Comprehensive Statistical Rating of National Hockey League Forwards and Defensemen based upon all On-Ice Events. In: 7th Annual MIT Sloan Sports Analytics Conference (2013)