

# A Markov Framework for Learning and Reasoning About Strategies in Professional Soccer

**Maaike Van Roy**

**Pieter Robberechts**

**Wen-Chi Yang**

*Department of Computer Science, KU Leuven; Leuven.AI  
Leuven, Belgium*

MAAIKE.VANROY@KULEUVEN.BE

PIETER.ROBBERECHTS@KULEUVEN.BE

WENCHI.YANG@KULEUVEN.BE

**Luc De Raedt**

*Department of Computer Science, KU Leuven; Leuven.AI  
Leuven, Belgium*

LUC.DERAEDT@KULEUVEN.BE

*Centre for Applied Autonomous Sensor Systems, Örebro University  
Örebro, Sweden*

**Jesse Davis**

*Department of Computer Science, KU Leuven; Leuven.AI  
Leuven, Belgium*

JESSE.DAVIS@KULEUVEN.BE

## Abstract

Strategy-optimization is a fundamental element of dynamic and complex team sports such as soccer, American football, and basketball. As the amount of data that is collected from matches in these sports has increased, so has the demand for data-driven decision-making support. If alternative strategies need to be balanced, a data-driven approach can uncover insights that are not available from qualitative analysis. This could tremendously aid teams in their match preparations. In this work, we propose a novel Markov model-based framework for soccer that allows reasoning about the specific strategies teams use in order to gain insights into the efficiency of each strategy. The framework consists of two components: (1) a learning component, which entails modeling a team's offensive behavior by learning a Markov decision process (MDP) from event data that is collected from the team's matches, and (2) a reasoning component, which involves a novel application of probabilistic model checking to reason about the efficacy of the learned strategies of each team. In this paper, we provide an overview of this framework and illustrate it on several use cases using real-world event data from three leagues. Our results show that the framework can be used to reason about the shot decision-making of teams and to optimise the defensive strategies used when playing against a particular team. The general ideas presented in this framework can easily be extended to other sports.

## 1. Introduction

There is currently a vast amount of data collected from professional sports matches. Clubs, news media, researchers, and fans are increasingly employing techniques from artificial intelligence to analyze this data. Markov models are a key analysis tool because they enable modeling the dynamic environment that characterizes many sports such as soccer, American football, ice hockey, basketball, volleyball, and (table) tennis. Specifically, they are used to model the observed team and player behavior and they have provided insights into areas such as valuing players' contributions for recruitment purposes (Cervone, D'Amour,

Bornn, & Goldsberry, 2016; Routley & Schulte, 2015; Rudd, 2011; Singh, 2019; Yam, 2019), performing match analysis (Fernández, Bornn, & Cervone, 2021; Liu & Hohmann, 2013; Pfeiffer, Zhang, & Hohmann, 2010; Wenninger & Lames, 2016), and assessing win probability (Dong, Shi, Chuong, Jiang, & Sun, 2015).

This paper presents a framework that can support a professional soccer coach’s tactical planning. To motivate the need for this, consider the following questions that a coach is confronted with:

1. Which areas of the pitch are most effective to help a team reach the area around the opponent’s goal?
2. Given that a player possesses the ball in a specific location, what is the chance of generating a shot with a higher probability of resulting in a goal later on in the possession?
3. How would shooting more often from outside the penalty box affect the number of goals the team would be expected to score over the course of a season?

Unfortunately, these are difficult questions to answer. On the one hand, one must be able to learn a model that accurately captures a team’s observed behavior in matches. On the other hand, one must reason, often in a counterfactual way, to understand the implications of what would happen if the team were to behave differently.

We develop a novel framework for addressing tactical questions such as those listed above based on the combination of learning and reasoning. The learning component entails learning the model underlying a Markov decision process (MDP), specifically the transition probabilities, where the state space consists of locations on the pitch and the actions involve moving between these zones or shooting on goal. The MDP only focuses on modeling offensive behavior by estimating the observed policy and the transition function when a team possesses the ball. Moreover, a separate MDP is learned for each team by estimating their observed policy from data about their matches. By averaging over, e.g., a season of data,<sup>1</sup> the MDP captures a team’s general behavior and is agnostic to the opponent. This setting roughly corresponds to a passive model-based reinforcement learning setting, albeit one where the reward is known (a reward of 1 is received for scoring a goal).

The MDPs will be learned from event stream data, which annotates various information (e.g., location, time, players involved) about all on-the-ball events (e.g., passes, tackles, and shots) that occur during a match. The data poses two challenges from a learning perspective. First, the observational nature of the data means that when an action is unsuccessful, its intended end location is not recorded and hence unknown. For example, if a player attempts a cross that the opposition clears, we are unsure of where the player was aiming for. We address this challenge using a combination of domain knowledge and predictive modeling. Second, the data is sparse due to the size of the pitch, dynamic nature of the game, and the relatively short seasons. Therefore, we estimate the probability model of each team’s MDP using a hierarchical Bayesian approach that uses a prior based on a “typical team”. The probability model is then specialized to an individual team on the basis of their data.

---

1. Data can go out of date due to switches in team personnel and general changes in playing style.

The reasoning component operates on two different levels. First, we demonstrate a novel application of probabilistic model checking (Kwiatkowska, Norman, & Parker, 2011) to reason about a team’s fixed, learned policy (i.e., one that is observed in the data). This enables reasoning about the probabilities of certain patterns of movement that a team may use to generate scoring opportunities. Moreover, we show how to reason about an opponent’s policy to evaluate the effect that certain defensive strategies would have on reducing the chance of conceding a goal. Second, we discuss several different ways to (slightly) modify a team’s observed policy. This enables reasoning about the effects of, e.g., shooting more or less often from certain areas of the pitch.

Using event stream data from three leagues, we provide a number of illustrative use cases. The first use case focuses on investigating the analytical-driven trend that long-distance shots have steadily decreased over the past decade in all major leagues. We reason about various aspects of long-distance shooting behavior, including the effects of shooting more or less often from outside the penalty box, which was a question posed by the director of analytics at a professional soccer club.<sup>2</sup> Intriguingly, our analysis goes against the new conventional wisdom and indicates that teams have overcorrected and may now be taking too few long-distance shots. Namely, teams would roughly score an extra goal per season if they shot more frequently from distance. The second use case focuses on optimizing a defending team’s game plan. We use our techniques to reason about (1) how an opponent may generate scoring opportunities and (2) the effect of certain defensive strategies on reducing the chance of conceding a goal. For each team, we found various regions on the pitch that are crucial for their chance-creation patterns. By forcing these teams to avoid these regions, a defending team could decrease their chances of conceding a goal, even if the opponent were to adapt to it.

## 2. Preliminaries

In this section, we provide background on soccer, the data, the notion of expected goals in soccer, Markov decision processes, and probabilistic model checking.

### 2.1 Soccer

Soccer is a ball sport that is played between two teams on a grass pitch with a dimension of 105 meters by 68 meters. A pitch is illustrated in Figure 1. Each match consists of two 45-minute halves where the clock runs continuously. The objective is to score more goals than your opponent.

Each team may have (at most) eleven players on the pitch consisting of one goalkeeper and ten outfield players. The goalkeeper guards the goal and is the only player that is allowed to touch the ball with their arms or hands, and only within a designated area called the penalty box (see Figure 1). The outfield players can move the ball over the pitch using their head or feet. Typical ways in which outfield players move the ball are by (1) a shot, (2) a pass, (3) a cross, and (4) a dribble. A shot is any type of action that deliberately tries to move the ball into the opponent’s goal. With a pass, a player tries to move the ball to another player on his team by kicking the ball with his foot or head (i.e., a headed

---

2. <https://twitter.com/devinpleuler/status/1226919308762193920>




Figure 1: Illustration of a soccer pitch. Each team has their own goal (drawn outside the pitch) at opposite sides of the pitch. The two rectangles near each team’s goal are mainly important for goalkeepers: from within the six-yard box (small rectangle) the goalkeepers can perform goal-kicks, within the penalty box (large rectangle) the goalkeepers are allowed to touch the ball with their hands. The dots and arcs near each goal are used for penalties. The dot and circle in the middle of the field are used for kick-offs.

pass). A cross is a special type of pass that originates from near the sides of the pitch close to the opponent’s goal and ends in or near the penalty box. Crosses typically follow a parabolic path through the air with the aim of bypassing several opposing players to reach a teammate who is in a shooting position. With a dribble, a player aims to move the ball by repeatedly kicking it while running. A dribble involves only the current on-the-ball player and possibly an opposing player that is bypassed or that interrupts the dribble. Typically, data providers do not agree on terminology, but they all distinguish between actions that just progress the ball and actions that aim to pass a defender. The definition of dribbles used in this work includes both actions that aim to carry the ball forward and actions that try to move the ball past an opposing player. Dribbles can fail when the player loses control of the ball and the other team recovers possession of the ball. On defense, players can also perform various actions, such as tackling an opposing player that possesses the ball (i.e., to force them to lose control of the ball) or trying to intercept a pass made by an opposing player.

## 2.2 Event Stream Data

While there are a variety of sources of data collected about soccer matches (e.g., box scores, video), one of the most widely available is event stream or play-by-play data with vendors collecting this data for 100s of leagues worldwide. Event stream data describes all on-the-ball actions that occur during each game, which is typically collected by human annotators while watching videos. A match typically has around 1500 to 3000 on-the-ball actions.<sup>3</sup> For each on-the-ball action, event stream data records a number of features such as the type of

3. The number of on-the-ball actions depends on what is annotated in the data (e.g., pressure events, separate actions for pass and receipt).




Figure 2: Illustration of a sequence of four actions from the Manchester City versus Liverpool game on January 14, 2018, as recorded in the event stream data format.

the action (e.g., pass, dribble, shot, interception), the start and end locations of the action, the body part used to execute the action, the result of the action (i.e., successful or not), the time at which the action was performed, the player who performed the action, and the team the acting player belongs to. Figure 2 illustrates a sequence of four actions from a game between Manchester City and Liverpool as they were recorded in the event stream data format.

### 2.3 Expected Goals

While goals are the main currency in soccer matches, they are exceedingly rare as only about three goals are scored per match on average. While shots are more common, on average a shot results in a goal only 10% of the time. Hence, the number of goals scored fluctuates and is subject to some randomness (e.g., lucky bounces or deflections). To try to give a more accurate view of chance creation, the analytics community developed the expected goals (xG) metric which aims to quantify the quality of each scoring opportunity (Green, 2012). This has now become a well-accepted metric that appears in mainstream media outlets and is mentioned by managers.

From a technical perspective, xG assigns a probability to each shot that represents its chance of directly resulting in a goal. These models are typically trained using logistic regression or gradient boosted trees on large historical datasets of shots. Each shot is described by the game context from when it was taken, and how this is represented is the key difference among existing models (Caley, 2015; Decroos, Bransen, Van Haaren, & Davis, 2019; Decroos & Davis, 2020; Ijtsma, 2015; Knutson, 2020; Lucey, Bialkowski, Mofort, Carr, & Matthews, 2015; Robberechts & Davis, 2020). The context can range from relatively straightforward such as just a location (Spearman, 2018) to very complex including many features (e.g., the shot’s distance and angle to the goal, time in the match, score differential, location of the goalie, the previous action type).

### 2.4 Markov Decision Process

An MDP is a tuple  $M = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$  where  $\mathcal{S}$  is a finite set of states,  $\mathcal{A}$  is a finite set of actions,  $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the transition function,  $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  and  $\gamma$  is the discount factor. The transition function  $P$  must denote a proper probability distribution for all state-action pairs  $(s, a)$ , that is,  $\sum_{s' \in \mathcal{S}} P(s, a, s') = 1$ . A state is *absorbing* if one cannot escape from it once entering. When an absorbing state is entered, the episodic task ends (either success or failure). Given an MDP, a path of length  $n$  is denoted by

$\rho_n = s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} \dots \xrightarrow{a_{n-1}} s_n$  where  $s_i \in \mathcal{S}$ ,  $a_i \in \mathcal{A}$  and  $P(s_i, a_i, s_{i+1}) > 0$ . Similarly, an infinite path is denoted by  $\rho = s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} \dots$ .

A policy (or strategy)  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  specifies, for each state, a probability distribution over all actions. Given an MDP and a policy  $\pi$ , the value function  $V_\pi : \mathcal{S} \rightarrow \mathbb{R}$  represents the value of being in a state. The value function is denoted as  $V$  (Bellman, 1966).

$$V_\pi(s) = \sum_{a \in \mathcal{A}} \pi(a | s) \sum_{s' \in \mathcal{S}} P(s, a, s') (R(s, a, s') + \gamma V_\pi(s'))$$

A Markov reward process (MRP) is a specialization of an MDP which allows for exactly one fixed policy and a reward function (Baier & Katoen, 2008). When learning the value function of a team, we learn it for a fixed policy and use a state-based reward function  $R : \mathcal{S} \rightarrow \mathbb{N}$  that assigns to each state a non-negative integer reward. Thus, the MDP becomes an MRP.

## 2.5 Probabilistic Model Checking

Probabilistic model checking is a formal verification technique that determines whether a stochastic system satisfies a given desired stochastic property (Kwiatkowska et al., 2011; Hensel, Junges, Katoen, Quatmann, & Volk, 2022). It provides rigorous guarantees by formulating both the system and the property in mathematical forms. The standard formalism includes various discrete and continuous Markov models and probabilistic temporal logics (Forejt, Kwiatkowska, Norman, & Parker, 2011). We consider MRPs, MDPs, and PCTL\* (which subsumes Probabilistic Computational Tree Logic (PCTL) and Probabilistic Linear Temporal Logic (LTL)) (Forejt et al., 2011).

PCTL\* is a temporal logic that expresses model properties over time and allows for probabilistic quantification. The syntax of PCTL\* is as follows. We fix a finite set of atomic propositions  $\mathcal{AP}$ . A state formula (denoted by  $\phi$ ) and a path formula (denoted by  $\psi$ ) over  $\mathcal{AP}$  can be constructed by the grammar below. A PCTL\* property is always a state formula  $\phi$ .

$$\begin{aligned} \text{state formula } \phi &::= \text{true} \mid c \mid \phi \wedge \phi \mid \neg \phi \mid \mathbb{P}_{\bowtie \mathbf{p}}[\psi] \\ \text{path formula } \psi &::= \phi \mid \mathbf{X} \psi \mid \psi \mathbf{U}^{\leq \mathbf{k}} \psi \mid \psi \mathbf{U} \psi \mid \neg \psi \end{aligned}$$

where  $c \in \mathcal{AP}$ ,  $\mathbf{p}$  is a probability such that  $0 \leq \mathbf{p} \leq 1$ ,  $\mathbf{k} \in \mathbb{N}$  is a positive integer and  $\bowtie \in \{\leq, <, \geq, >\}$ . A PCTL\* property evaluates to either *true* or *false* in a state.

Intuitively,  $\mathbf{X} \alpha$  means  $\alpha$  holds in the next step,  $\alpha \mathbf{U} \beta$  means  $\alpha$  holds until  $\beta$  holds,  $\alpha \mathbf{U}^{\leq \mathbf{k}} \beta$  means  $\alpha$  holds until  $\beta$  holds and  $\beta$  will hold within  $\mathbf{k}$  steps, and  $\mathbb{P}_{\bowtie \mathbf{p}}[\alpha]$  means  $\alpha$  holds with a probability  $\bowtie \mathbf{p}$ . Additionally, the commonly used temporal component  $\mathbf{F} \alpha$  ( $\alpha$  will eventually hold) can be defined by the until operator:  $\mathbf{F} \alpha \equiv \text{true} \mathbf{U} \alpha$ . General examples of PCTL\* properties are  $\mathbb{P}[\mathbf{F} A]$  (“probability of finally reaching a state A”), or  $\mathbb{P}[\neg A \mathbf{U} B]$  (“probability of never entering A before reaching B”). A specific example for soccer could be  $\mathbb{P}[\text{penalty\_box}]$  (“probability of finally reaching the penalty box”) where *penalty\_box* denotes all states in the Markov model that correspond with being in the penalty box. We refer the readers to Baier (1998) for more details of PCTL\*.

### 3. Modeling Offensive Behavior as a Markov Model

To reason about a team’s observed behavior, we model each team’s offensive behavior using an MDP. This formalism allows us to model and analyze the dynamic environment of soccer and the in-game decisions made by players. On a high level, the proposed MDP models the probability of a team moving the ball from one location to another location on the pitch. This can be done by either (1) a movement action such as a pass or a cross to a teammate, or a dribble, and (2) a shot. Specifically, we model this at the level of a possession sequence, which is a sequence of consecutive on-the-ball actions made by the same team. A possession sequence corresponds to a path in the MDP (see Section 2.4).

Formally, the MDP consists of the state space  $\mathcal{S}$ , the set of actions  $\mathcal{A}$ , the transition function  $P$ , the policy  $\pi$ , the reward function  $R$ , and a discount factor  $\gamma$ . We discuss each of these in the following sections.

#### 3.1 State Space

The state space is defined as  $\mathcal{S} = \mathcal{E} \cup \mathcal{L}$ . We use  $\mathcal{E} = \{lost\_possession, no\_goal, goal\}$  to denote the set of absorbing states where *lost\_possession* signifies loss of possession (i.e., a failed movement action), *no\_goal* a failed shot, and *goal* a successful shot.  $\mathcal{L}$  is the set of transient states that can be entered and exited during a possession sequence. As event data only contains information about on-the-ball actions (i.e., the locations of players not possessing the ball are unknown), we will define the state space based entirely on the location of the ball. Specifically, the transient states are defined by partitioning the field into various grid cells. The precise set of transient states varies by use case.

#### 3.2 Action Space

For each state  $s \in \mathcal{L}$ , we consider the actions  $\mathcal{A} = \{move\_to(s) \mid s \in \mathcal{L}\} \cup \{shoot\}$ . The action *move\_to(s)* denotes that a player intends to move the ball to state  $s$ . Modeling the *intended* end location where the player wants to move the ball to allows us to explicitly model and analyze the goal-directed policies of players and teams. In contrast, existing approaches for soccer have used the *observed* end location of movement actions, which makes analyzing a team’s decisions not straightforward (Rudd, 2011; Singh, 2019; Yam, 2019).

Using an expressive action space means that the number of actions is larger than the number of states (i.e.,  $len(\mathcal{A}) > len(\mathcal{S})$ ), introducing an enormous policy space of the MDP. This indicates that the team behavior analysis has a high worst-case computational complexity and could potentially be a performance bottleneck. However, this is not an issue in practice because a number of actions “do not make sense” (e.g., passing from in front of the opponent’s goal back to your own goal) and hence are not observed in the real-world data.

#### 3.3 Transition Function

Each of the previously defined actions can either succeed or fail. This is modeled by the transition function  $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ . For the absorbing states  $\mathcal{E}$ ,  $\mathcal{A} = \emptyset$ , and the

only possible transition is a self-loop with a probability of one. For each state  $s \in \mathcal{L}$ , the transition function is defined as follows:

- $P(s, \text{move\_to}(s'), s')$  is the probability of successfully moving to state  $s' \in \mathcal{L}$  from state  $s$ ;
- $P(s, \text{move\_to}(s'), \text{lost\_possession})$  is the probability of unsuccessfully moving to state  $s' \in \mathcal{L}$  from state  $s$ , this is equal to  $1 - P(s, \text{move\_to}(s'), s')$ ;
- $P(s, \text{shoot}, \text{goal})$  is the probability of scoring a goal from state  $s$ , and this can be viewed as a location-based xG value;
- $P(s, \text{shoot}, \text{no\_goal})$  is the probability of failing to score a goal when shooting from state  $s$ , this is equal to  $1 - P(s, \text{shoot}, \text{goal})$ ;
- $P(s, a, s') = 0$  in all other cases.

### 3.4 Policy

The policy  $\pi$  defines the probability distribution over actions for each state:  $\pi(a | s) = Pr[A = a | S = s]$ . In this case, it denotes how likely a team is to choose each specific movement or shot action in each state. The policy and the transition function completely define the in-game behavior of the team. The policy is also the only part that the team or player can immediately control. We do not compute the optimal policy for each team, as removing the randomness from a team’s policy makes its behavior very predictable for an opponent. Often, teams will also not be acting optimally as the individual players are not perfect optimizers (Sandholtz & Bornn, 2020). Rather, we look into analyzing the current policy that teams use. Thus, we assume that the current policy of the MDP is fixed and is equal to the one that is observed in the data.

Analyzing the currently employed policy can already provide numerous insights regarding a team’s strategies. Furthermore, by adapting the policy to certain counterfactual scenarios and estimating the effects of these changes, further insights into what would happen if different strategies were to be employed can be provided to practitioners.

### 3.5 Reward Function and Discount Factor

The reward function  $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  of our model mimics the rewards obtained during a real-life soccer game. It assigns a reward of one when a goal is scored from a state  $s \in \mathcal{L}$ , i.e.,  $R(s, \text{shoot}, \text{goal}) = 1$ . All other actions do not receive a reward. The discount factor  $\gamma$  is set to 1. As soccer is a low-scoring sport and extreme scores (i.e., more than five goals) are uncommon, it is reasonable to assume that the rewards should not be discounted.

## 4. Learning the Markov Model from Event Stream Data

Because teams have unique styles, we learn a separate transition function and observed policy for each team from historical event stream data. The resulting MDP for each team is agnostic over its opponent and captures the *general* behavior employed by a team during their games. However, to learn the MDP, two important challenges must be tackled.



First, explicitly modeling the intended end location of actions complicates calculating the transition probabilities. Estimating these probabilities requires knowing the total number of times a team *chooses* to move from any state  $s \in \mathcal{L}$  to any other state  $s' \in \mathcal{L}$ , regardless of whether the action succeeded. However, because the data contains only the end locations of successful actions and not the intended end locations of failed actions, this total number is unknown and also unknowable without asking the player. We address this by combining domain knowledge with predictive modeling.

Second, the amount of data per team per season is limited, with teams typically performing between 400 and 700 shots and 10,000 and 25,000 passes per season. Using data that is more than a few seasons old is not of great use due to changes in players, managing personnel, and game trends. Most teams will also not often perform the same actions at the exact same locations. Moreover, the actions are not evenly distributed over the pitch. For example, there is much more action in midfield than there is near the opponent’s goal. Thus, the amount of relevant data available per team per state can be quite limited, and accurately estimating the probabilities for these states can become problematic. We mitigate these sparsity issues by estimating these probabilities using a hierarchical Bayesian approach with a prior that is based on a “typical team”.

Next, we outline our approaches for (1) predicting the intended end locations of failed actions, and (2) learning the transition function and policy with a hierarchical Bayesian approach.

#### 4.1 Predicting the Intended End Locations of Actions

For each failed movement action, we aim to predict a probability distribution over all possible intended end locations. As a movement action can both be a pass, a cross, or a dribble, which each have characterizing dynamics, we will treat each of these separately. Additionally, the predictions will be made on a per-team basis to capture each team’s distinctive behavior.

**Passes:** A gradient boosted trees ensemble is first trained on the end states of a team’s successful passes, as it is reasonable to assume that these passes reached their intended destination. Each pass is described by a set of 6 features, including the pass’ start state, the direction of the pass, the body part used to execute the pass, and the start states of the three preceding actions.<sup>4</sup> These features are commonly used in sports analytics to describe the context of actions (Caley, 2015; Decroos et al., 2019; Power, Ruiz, Wei, & Lucey, 2017; Robberechts & Davis, 2020). Appendix C provides and discusses an ablation study on these features. The trained model is subsequently used to obtain an initial probability distribution over all possible intended end states for a team’s failed passes. We improve the predicted distributions based on two reasonable assumptions. First, we assume that passes generally travel in straight lines. Thus, the intended end location must lie along the line through the pass’ observed start location, its (failed) end location, and the location where the ball would have gone out of bounds. Hence, all end states that are not on this

---

4. Sequences containing less than three actions are not considered during training. To infer the end location of failed passes with less than three preceding actions, we simply use the distribution over observed successful end locations in the training set.

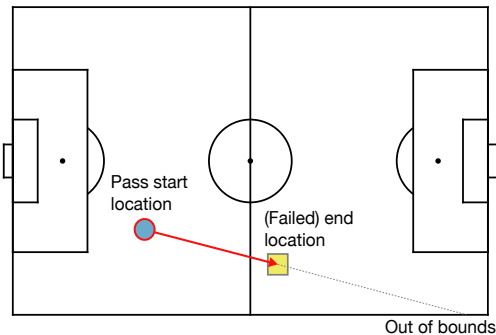


Figure 3: Illustration of an unsuccessful pass, starting at the circle and unsuccessfully ending at the square. The intended end location most likely lies along the line that starts from the pass’ start location, passes through its failed end location, and ends when it goes out of bounds. If the pass was intercepted, the intended end location most likely lies on the dashed part of the line.

straight line can be pruned from the distribution. Second, when a pass was intercepted by the opponent before it could reach the intended teammate, the intended end location should lie further on this line than the location where it was intercepted. Hence, when the pass was intercepted, we can additionally prune all end states on the straight line between its start location and its (failed) end location. Figure 3 illustrates these ideas. However, pruning these states has as a consequence that the remaining probabilities no longer form a probability distribution. Therefore, we normalize the distribution after pruning.

**Crosses:** Similarly to our approach for passes, we train a gradient boosted trees ensemble using the same set of features on all successful crosses of a team. Next, the ensemble is used to predict a probability distribution over all possible intended end states and is improved with domain knowledge. Here, we can make use of the data provider’s definition that crosses originate from the sides of the field near the opponent’s goal and end in the area in front of the opponent’s goal.<sup>5</sup> Additionally, we make a simplifying assumption that crosses follow a parabolic path through the air, thereby bypassing all states that lie in between these two locations.<sup>6</sup> Thus, with this definition in mind, the intended end location will most likely be a location that is close to the failed one. Specifically, we post-process the predicted distribution by only retaining those states that lie within a radius  $r$  from the failed one and normalize the probabilities afterwards.

**Dribbles:** Given that dribbles are extremely local actions, we assume that when dribbles fail the player had intended to move the ball to a location that is not too far from its failed end location. Thus, we assume that the most likely intended end state of a failed dribble is simply the same state where the dribble failed.

5. See <https://www.statsperform.com/opta-event-definitions/>

6. In reality, 86.5% of all crosses in the data set are “chipped into the air”.

## 4.2 A Bayesian Approach to Learning the Transition Model and Policy

Using a hierarchical Bayesian approach to estimate the transition function and policy allows sharing information between teams by using a prior that is based on a “typical team”. Conceptually, this can be seen as starting with a generic model for how a typical team behaves and then adapting it based on the observed actions that the specific team performed during games. If there is strong evidence that a team deviates from what is typical in a given location, this will be picked up by the model. However, the parameters will shrink towards the prior for locations where there is little data for a team.

We define a separate hierarchical Bayesian model for the transition function and policy of a team:

**Transition function model:** For a specific team  $t$ , we model the probability that a chosen action  $a \in \mathcal{A}$  in a certain state  $s \in \mathcal{L}$  succeeds as a Bernoulli random variable  $O_{t,s,a}$ :

$$\begin{aligned} O_{t,s,a} &\sim \text{Bernoulli}(p_{t,s,a}) \\ p_{t,s,a} &= \text{invlogit}(\gamma_{t,s,a}) \\ \gamma_{t,s,a} &\sim \mathcal{N}(\mu_{s,a}, \sigma_{t,s,a}^2) \\ \sigma_{t,s,a}^2 &\sim \text{Half-Normal}(5.0) \end{aligned}$$

Here,  $\text{invlogit}(x)$  stands for the inverse logit function  $1/(1 + e^{-x})$ . The log-odds of team  $t$  successfully completing action  $a$  in state  $s$  is represented by  $\gamma_{t,s,a}$  and is normally distributed with an overall prior mean  $\mu_{s,a}$  and a team-dependent variance  $\sigma_{t,s,a}^2$ . In turn, the team-dependent variance is half-normally distributed with a scale factor of 5.0. This corresponds to a weakly informative prior for the variance. The overall prior mean is computed based on the data of all other teams. The probabilities of this prior mean are calculated using simple counts and afterwards Gaussian smoothing is applied to ensure spatial coherence.

**Policy model:** For a specific team  $t$ , we model the probability of choosing a specific action in state  $s$  as a categorical random variable  $A_{t,s}$  with action type probabilities  $\vec{p}_{t,s}$  (i.e., a specific *move\_to* or *shoot* action):

$$\begin{aligned} A_{t,s} &\sim \text{Categorical}(\vec{p}_{t,s}) \\ \vec{p}_{t,s} &= \text{softmax}(\vec{\lambda}_{t,s}) \\ \lambda_{t,s,a} &\sim \mathcal{N}(\alpha_{s,a}, \nu_{t,s,a}^2) \\ \nu_{t,s,a}^2 &\sim \text{Half-Normal}(5.0) \end{aligned}$$

Here,  $\text{softmax}(\vec{x})$  is the softmax function where each entry  $i$  in  $\vec{x}$  is set equal to  $e^{x_i} / \sum_k e^{x_k}$ . The log-odds of team  $t$  choosing each different action  $a \in \mathcal{A}$  in state  $s$  is represented by  $\vec{\lambda}_{t,s}$ . For each chosen action  $a$ , each  $\lambda_{t,s,a}$  is normally distributed with an overall prior mean  $\alpha_{s,a}$  and a team-dependent variance  $\nu_{t,s,a}^2$ . In turn, the team-dependent variance is half-normally distributed with a scale factor of 5.0. This again corresponds to a weakly informative prior for the variance. Similarly to the transition function model, the overall prior mean is computed based on the data of all other teams using simple counting and Gaussian smoothing.

## 5. Reasoning About Learned Policies

The tactical planning to prepare for a match is an arduous task that involves identifying the movement patterns from which a team creates its best chances. Based on this information, a coach can then devise a tactical plan that utilizes his team’s strong points and disrupts the opponent’s most effective patterns. For example, from an offensive perspective, the coach may want to provide advice to players about the merits of long throw-ins to create immediate scoring chances vs. focusing on trying to maintain possession in order to create a better chance after a couple of additional actions. From a defensive perspective, the coach might want to identify which action sequences and areas of the pitch the opposing team tends to use to generate the most shooting opportunities.

In a highly dynamic environment such as soccer where the exact same patterns never repeat and the outcome is affected by variability due to luck, fatigue, and the skill of players, these are difficult questions to answer both for a coach and a data analyst. Analyzing the efficiency of strategic behaviors requires reasoning about the inherent uncertainty in terms of which subsequent actions will occur and whether these actions will succeed or fail. Unfortunately, assessing this is (nearly) impossible to do solely based on the raw data.

Our key insight is that probabilistic model checking techniques can be applied on a team’s MRP (i.e., their MDP with the team’s policy fixed) to reason about the efficacy of various offensive and defensive strategies in soccer. That is, probabilistic model checking techniques can provide formal guarantees about the probabilities of teams scoring or reaching other desired situations from certain possible behaviors (i.e., sequences of actions) that can arise in the system. Thus, these techniques allow us to (1) formulate various possible behaviors teams could exhibit, and (2) evaluate and compare their relative merits. We will now describe our proposal for how to use these techniques to reason about the efficacy of various offensive and defensive strategies in soccer.

### 5.1 Reasoning About the Effect of Offensive Strategies

Ultimately, a team’s offensive strategy revolves around generating shots and scoring goals. Consequently, a coach may want to provide general advice about how to optimize the chance of scoring such as whether his players should take a shot or pass in the hopes of generating a better shot later on. While a player may have a good sense of the chance of scoring from a particular shot or the chance of successfully completing a pass, it is harder to assess the probability of generating a shot within the next couple of actions. We propose to evaluate the relative merits of such options as follows. Given a finite sequence of actions denoted by their end states  $(s_1, \dots, s_{n-1})$  where  $s_1, \dots, s_{n-1} \in \mathcal{L}$ , the probability of yielding a goal immediately after this sequence can be formulated as:

$$P_{s=?}[\mathbf{X}(s_1 \wedge (\dots \wedge (\mathbf{X}(s_{n-1} \wedge (\mathbf{X} \text{ goal})))))] \quad (1)$$

where  $s \in \mathcal{L}$  is the state in which the sequence starts. By comparing the returned probabilities for multiple such sequences, the relative merits of different courses of actions can be compared.

More generally, when possessing the ball at a certain location, one might want to know *how likely* it is to ever generate a better shot within the same possession. For a particular

state  $s \in \mathcal{L}$ , this probability can be computed as:

$$P_{s=?}[\mathbf{X}(L_1 \wedge (\mathbf{F}(L_2 \wedge (\mathbf{X} \text{goal})))))] \quad (2)$$

where  $L_1, L_2 \subseteq \mathcal{L}$  with  $L_1$  the set of states to which the ball can be moved to from  $s$  (i.e., performing anything but a shot in the current state) and  $L_2$  the set of states with a higher success probability for executing a shot than the one associated with performing it in state  $s$ .

## 5.2 Reasoning About the Effect of Defensive Strategies

From a defensive perspective, broadly speaking, a soccer team’s objective is to minimize its chance of conceding a goal. This requires understanding how the opposing team tends to generate dangerous situations. Based on this understanding, a coach would want to evaluate the efficacy of different tactics for disrupting the opponent’s attack. For example, how would an opponent’s chance of generating a shot be affected by forcing them *to avoid reaching certain areas of pitch*? It is possible to reason about such scenarios using probabilistic model checking. Specifically, we reason about where a team should focus defensively to (1) limit or suppress the number of shots the opponent will generate and (2) disrupt movement in the midfield to reduce the chance of conceding.

### 5.2.1 SHOT SUPPRESSION

Johan Cruyff, one of the greatest soccer players of all time, famously said “You can’t score if you don’t shoot”. Hence, a defending team can ultimately reduce its chances of conceding a goal by attempting to reduce the number of shots the opponent takes. As over 90% of the shots are taken from the yellow-shaded region (denoted *shot\_locations*) in Figure 4, a first approach to reduce your chances of conceding is to *directly* suppress the shots the opponent takes from this region. However, as a dangerous situation was already created by reaching this region, this is not an ideal situation to be in as a defending team. Therefore, the second approach *indirectly* suppresses the number of shots taken by limiting the number of times the opponent reaches this favored region. Next, we explain the specific properties that can be evaluated by a model checker to reason about the effect of both approaches.

**Direct shot suppression.** To reason about how a team generates shots, consider the query:

$$P_{s=?}[\mathbf{F}(\text{shot\_locations} \cup (\text{goal} \vee \text{no\_goal}))] \quad (3)$$

which gives the probability of a sequence starting in  $s$  eventually reaching a state in *shot\_locations* after which it eventually reaches either *goal* or *no\_goal*. That is, it computes the unrestricted probability of a sequence starting in  $s$  resulting in a shot from *shot\_locations*. Suppose we could prevent an opponent from ever entering  $s' \in \text{non\_shot}$ , where  $\text{non\_shot} = \mathcal{L} \setminus \text{shot\_locations}$ . We can reason about the effect of forcing the opposing team to avoid state  $s'$  (i.e., the opposing team can never enter state  $s'$ ) on the probability of shooting using the following query:

$$P_{s=?}[\neg s' \cup (\text{shot\_locations} \cup (\text{goal} \vee \text{no\_goal}))] \quad (4)$$




Figure 4: Illustration of three interesting areas used to optimize the defensive game plan. The region where most shots are taken from is shaded in yellow, the final third entry region is shaded in gray, and the middle third of the pitch is shaded in blue.

Both queries can be combined to reason about the effect of avoiding  $s'$  on the probability of eventually shooting when starting the sequence in  $s$ . However, a sequence can begin in any state. Therefore, to compute the average reduction in your opponent's probability of shooting when forced to avoid  $s'$ , we must sum over all locations in  $non\_shot$  in numerator and denominator:

$$1 - \frac{\sum_{s \in non\_shot} P_{s=?}[\neg s' \text{ U } (shot\_locations \text{ U } (goal \vee no\_goal))]}{\sum_{s \in non\_shot} P_{s=?}[\text{F } (shot\_locations \text{ U } (goal \vee no\_goal))]} \quad (5)$$

By computing this for all  $s'$ , we can measure each state's importance for directly suppressing shots.

**Indirect shot suppression.** To reason about how likely your opponent is to reach  $shot\_locations$  from a location  $s$ , consider the following query:

$$P_{s=?}[\text{F } shot\_locations] \quad (6)$$

We can reason about the effect of a counterfactual policy that forces the opposing team to avoid state  $s'$  on this probability using the following query:

$$P_{s=?}[\neg s' \text{ U } shot\_locations] \quad (7)$$

By combining these two queries, we can reason about the effect on your opponent's probability of ever reaching  $shot\_locations$  if you can force them to avoid entering location  $s'$ :

$$1 - \frac{\sum_{s \in non\_shot} P_{s=?}[\neg s' \text{ U } shot\_locations]}{\sum_{s \in non\_shot} P_{s=?}[\text{F } shot\_locations]} \quad (8)$$

By computing this for all  $s'$ , we can measure each state's percent decrease in the probability of reaching  $shot\_locations$  when prevented from entering  $s'$ . This gives an indication of the importance of  $s'$  for indirectly suppressing shots.

### 5.2.2 MOVEMENT SUPPRESSION

Disrupting your opponent’s movement further away from the penalty area can be equally effective because an opponent will often start their possession in their own defensive half or near midfield. By employing such defensive strategies, the defending team may be able to make it more challenging to get the ball into dangerous positions.

Figure 4 visualizes two critical regions (denoted *region*) for which we want to decrease the attacking team’s overall chances of scoring from: the final third entry and middle third regions. Once in the middle third of the pitch, the attacking team has already bypassed some defending players and created some threat. Once the team reaches the final third entry region, the threat increases as the team possesses the ball closer to goal.

The goal of **movement suppression** is to identify a set of states *area* that if your opponent was forced to avoid, would decrease their chance of scoring from each state  $s \in region$  by at least  $x$  percentage points. Formally, this can be framed as the following query:

$$\forall s \in region : (\mathbb{P}_{s=?}[F goal] - \mathbb{P}_{s=?}[\neg area \cup goal]) \geq x \quad (9)$$

Thus, this query computes both the unrestricted chance of scoring from state  $s$  and the restricted chance of scoring from state  $s$  when forced to avoid any state in *area*. Then it checks whether the difference is at least  $x$  percentage points for each state in *region*. The different possible sets of states that form *area* can be found using the following query:

$$area = \{s'' \in \mathcal{L} \setminus region \mid \mathbb{P}_{s''=?}[F s'] \geq b\} \quad (10)$$

This query finds all states  $s''$  around a central state  $s' \in \mathcal{L} \setminus region$  such that they reach  $s'$  with a probability greater than threshold  $b$ . The specific values for  $b$  and  $x$  can be chosen dependent on the team and use case. While we focus on the final third entry and middle third regions, this approach is generally applicable to any region that is interesting to practitioners.

## 6. Reasoning About Alternative Policies

The soccer MDPs are learned from observed team behavior. From a coaching and tactical point of view, it would be incredibly valuable to assess the merits of alternative strategies (i.e., slightly different policies). For example, a coach may want to know what the effect is of passing more aggressively (e.g., fewer backwards or lateral passes, more through balls) or shooting more frequently from outside the penalty box. The small number of games and inherent randomness of soccer make it difficult to evaluate slightly different policies based on match results. This is compounded by the high-stakes nature of professional soccer, and the incredibly short tenure of most managers discourages such experimentation. Therefore, a simulation-based approach to assessing the efficacy of slight modifications would be extremely useful.

This section describes techniques to reason about these tactical counterfactual questions in a simulated manner by estimating the effect of (slightly) altering a team’s observed policy. From a technical perspective, one way to modify a policy is by altering the probabilities of selecting various actions in a particular state. However, such a modification raises two questions:

1. If one action becomes more likely, then other actions must become less likely to maintain a probability distribution. Thus, which actions become less likely and by how much does each one’s probability of being selected decrease?
2. How can we estimate the effect of the altered policy? This is complicated by the fact that for some actions, namely shots, their probability of succeeding is unlikely to remain constant as their volume increases (or decreases) due to a quality-quantity trade-off. Thus, estimating the effect of the modified policy requires accounting for this trade-off.

We treat movement actions and shots separately and we now describe how we address these two points for each action type.

### 6.1 Counterfactual Movement Actions

First, we describe two approaches for altering the movement policy. Second, we provide a general approach to evaluate the effect on the created danger of a team.

#### 6.1.1 CHANGING THE MOVEMENT POLICY

It is possible to modify a policy to select an action both more or less often. For simplicity, the following discussion focuses on increasing the probability of selecting actions with the discussion of decreasing the probability deferred to Appendix B.1.

Suppose that, given a start state  $s$ , the objective is to move to a set of states  $area$  by an additional  $x$  percent. This requires increasing the probability of moving from  $s$  to each state  $s' \in area$  as follows:

$$\pi'(move\_to(s') | s) = (1 + x) * \pi(move\_to(s') | s) \tag{11}$$

However, to maintain a probability distribution, the probability of moving from  $s$  to all other states  $s'' \in \mathcal{L} \setminus area$  must decrease. We describe two approaches to address this: the **proportional approach**, which was originally proposed for basketball (Sandholtz & Bornn, 2018, 2020), and the novel **same aggressiveness approach** which is tailored towards soccer.

**Proportional approach.** This approach assumes that when the probability of performing an action increases in a state, the probability of all other actions decreases proportionally to their original share. Hence, in the adapted policy the team will try to move to all states not in  $area$  slightly less often. Concretely, when increasing the probability of trying to move to any state in  $area$ , the probability of moving to any other state  $s'' \notin area$  is decreased as follows:

$$\pi'(move\_to(s'') | s) = \frac{\pi(move\_to(s'') | s)}{total\_other} * (total\_other - x * total\_area) \tag{12}$$

with  $total\_other$  and  $total\_area$  defined as:

$$total\_other = \sum_{s'' \in \mathcal{L} \setminus area} \pi(move\_to(s'') | s) \tag{13}$$

$$total\_area = \sum_{s' \in area} \pi(move\_to(s') | s) \tag{14}$$



**Same aggressiveness approach.** In soccer, certain movements, namely backwards and lateral passes, generally tend to be less risky than playing the ball forward. Consequently, a team’s aggressiveness in a state  $s$  roughly corresponds to the total probability of selecting any action that moves the ball forward. The previous proportional modification could alter this balance. This is not always ideal as a team might want to assess what would happen if they slightly avoided the middle of the field but kept their aggressiveness the same (i.e., use the flanks more often). Hence, this approach aims to keep a team’s aggressiveness unchanged.

Concretely, when increasing the probability of trying to move to any state in  $area$ , the probability of moving to any other state  $s'' \notin area$  is decreased using the following approach. If  $s' \in area$  is a state that lies higher up the pitch than state  $s$  (i.e., the ball needs to be moved forward to reach state  $s'$ ), the additional probability of choosing to move from  $s$  to  $s'$  is removed from all other states  $s'' \in s^h \subseteq \mathcal{L} \setminus area$  that also lie higher up the pitch than state  $s$  as follows:

$$\pi'(move\_to(s'') | s) = \frac{\pi(move\_to(s'') | s)}{total} * (total - x * \pi(move\_to(s') | s)) \quad (15)$$

with  $total$  defined as:

$$total = \sum_{s'' \in s^h} \pi(move\_to(s'') | s) \quad (16)$$

For all other states  $s'' \notin s^h$ , the policy remains unchanged. Similarly, if  $s' \in area$  lies at an equal height or lower on the pitch than state  $s$  (i.e., the ball needs to be moved backward or laterally to reach state  $s'$ ), the additional probability of choosing to move from  $s$  to  $s'$  is removed from all other states  $s'' \in s^l \subseteq \mathcal{L} \setminus area$  that also lie at an equal height or lower on the pitch than state  $s$  using the same formulas. To keep the probability of moving forward equal, the policy of moving to the other states (i.e., higher up the pitch) then remains unchanged.

### 6.1.2 ESTIMATING THE EFFECT OF AN ALTERED MOVEMENT POLICY

Estimating the effect of employing the adapted policy can be done by measuring the effect on the team’s overall chances of scoring. Specifically, we measure a team’s chances of scoring from each state using the value function, both *before* and *after* they adapt their policy, and compare them. The following formula computes the percent change in the team’s overall chances of scoring when employing a new policy  $\pi'$  compared to their original policy  $\pi$ :

$$1 - \frac{\sum_{s \in \mathcal{L}} V_{\pi'}(s)}{\sum_{s \in \mathcal{L}} V_{\pi}(s)} \quad (17)$$

We assume that the number of possession sequences starting in each state is not affected by the changes in the policy. Thus, the probability of starting a sequence in each state is not taken into account in the above formula.

## 6.2 Counterfactual Shot Actions

Similar to movement actions, we consider both increasing and decreasing the shot volume of teams. Changing the policy itself is relatively straightforward. However, care must be

taken when evaluating the effect of policy changes as increasing (decreasing) the propensity to shoot in a state will not only affect the number of shots a team takes, but also the success probability of these shots as there is likely a quantity-quality trade-off. We propose a method that takes both the quantity and the quality of the shots into account when evaluating the effect of shooting more or less often.

### 6.2.1 CHANGING THE SHOT POLICY

When altering the shot volume in a particular state, we assume that when the probability of performing a shot increases (decreases) in a state, the probability of all other actions decreases (increases) proportionally to their original share. Mathematically, increasing the shot probability by  $x$  percent in a state  $s$  can be done using the following formulas:

$$\pi'(shoot | s) = (1 + x) * \pi(shoot | s) \quad (18)$$

$$\forall s' \in \mathcal{L} : \pi'(move\_to(s') | s) = \frac{\pi(move\_to(s') | s)}{total} * (total - x * \pi(shoot | s)) \quad (19)$$

Decreasing the shot probability by  $x$  percent in a state  $s$  can be realized as:

$$\pi'(shoot | s) = (1 - x) * \pi(shoot | s) \quad (20)$$

$$\forall s' \in \mathcal{L} : \pi'(move\_to(s') | s) = \frac{\pi(move\_to(s') | s)}{total} * (total + x * \pi(shoot | s)) \quad (21)$$

with in both cases

$$total = \sum_{s' \in \mathcal{L}} \pi(move\_to(s') | s) \quad (22)$$

### 6.2.2 ESTIMATING THE EFFECT OF AN ALTERED SHOT POLICY

Ultimately, the objective of shooting more frequently is to score more goals. However, goals are extremely rare events in soccer matches, and even most very high-quality shots from open play have less than 20% chance of resulting in a goal. Therefore we evaluate the effect of a new shot policy by estimating the number of goals that the team is expected to score over the course of a season. This can be done by knowing (1) the expected number of shots in each state over the course of a season, and (2) the probability of a shot resulting in a goal for each state. Formally, this can be computed as:

$$E[goals | \pi] = \sum_{s \in \mathcal{L}} E[shots\ in\ s | \pi] * P(s, shoot, goal) \quad (23)$$

**Estimating the expected number of shots.** Estimating the expected number of shots in each state is relatively straightforward, and can be done as follows:

$$E[shots\ in\ s | \pi] = \pi(shoot | s) * E[visits\ to\ s | \pi] \quad (24)$$

In turn, estimating the expected number of visits to a state over the course of a season can be computed by knowing (1) how many possessions start in each state  $s' \in \mathcal{L}$  and (2) given that a possession starts in  $s'$ , how many times state  $s$  will be visited prior to absorption. The number of possessions starting in each state  $s'$  can be computed by calculating it from

the event stream data. We assume that this number is not affected by the changes in the policy. The number of visits to each state can be derived directly from the MDP by computing its fundamental matrix  $N$ . Each entry  $N_{ij}$  is equal to the expected number of times state  $j$  will be visited before an absorbing state is reached, given that the possession started in state  $i$ . Formally,  $N$  can be computed as:

$$N = (I - Q)^{-1} \quad (25)$$

where  $I$  is the  $|\mathcal{L}|$ -by- $|\mathcal{L}|$  identity matrix and  $Q$  describes the probability of transitioning from one state  $s \in \mathcal{L}$  to another. Each entry  $Q_{ij} = P(i, \text{move\_to}(j), j) * \pi(\text{move\_to}(j) | i)$  and  $\pi$  can either be the observed (to estimate the current expected number of goals) or altered policy. Appendix B.2 provides an analysis on the accuracy of the outlined approach when using the observed policy.

**Estimating the success probability of shots.** Finally, estimating  $P(s, \text{shoot}, \text{goal})$  for the observed policy  $\pi$  is also straightforward as it is simply the probability learned from the data. However, when we modify the policy, this will likely affect the success probability of the shots taken because there is likely a quantity-quality trade-off. For example, increasing the shot volume from a particular location would likely entail a team taking more low-quality shots from that location. When decreasing their shot volume, it stands to reason that the team likely continues to take high-quality chances and tries to remove the lower-quality shots from their profile. This relation between the frequency of a given shot type and its probability of succeeding has previously been discussed in basketball (Goldman & Rao, 2014; Sandholtz & Bornn, 2018, 2020).

It seems reasonable to assume that such a relationship would also exist for soccer. Therefore, we adjust for this trade-off when computing  $P(s, \text{shoot}, \text{goal})$  when an altered policy  $\pi'$  is in place. Recall that  $P(s, \text{shoot}, \text{goal})$  is the xG value for any shot taken in state  $s$  because our model uses a purely location-based xG score. To obtain a fine-grained distribution over possible xG values for a given state  $s$  we compute the xG values of each shot using the XGBoost model with advanced feature set included in the SOCCER\_XG package.<sup>7</sup> For a detailed description of the model, we refer to Robberechts and Davis (2020). We then use the distribution of these xG values to derive an adjustment to the estimate of  $P'(s, \text{shoot}, \text{goal})$  for an altered policy  $\pi'$ . We apply a different modification for increasing and decreasing the number of shots.

*Case 1: Increasing the propensity to shoot.* In this case, each additional shot taken will likely be of lower quality. In the new policy,  $E[\text{shots in } s | \pi'] - E[\text{shots in } s | \pi]$  more shots will be taken than in the original policy. We only modify the xG value for the additional shots taken in  $\pi'$  by awarding them an xG of:

$$P(s, \text{shoot}, \text{goal}) - (\mu_s - \mu_s^{low}) \quad (26)$$

where  $\mu_s$  is the average computed xG of all shots occurring in state  $s$  and  $\mu_s^{low}$  is the average computed xG of the below-average shots occurring in state  $s$ . We illustrate this for one

7. [https://github.com/ML-KULeuven/soccer\\_xg](https://github.com/ML-KULeuven/soccer_xg)




Figure 5: Possible expected goals (xG) distribution in one state. Orange indicates the average xG value ( $\mu_s$ ), green indicates the average of the below-average xG values ( $\mu_s^{low}$ ), and purple indicates the average xG value of the shots with an xG value in the top 90% of all shots in the considered zone ( $\mu_s^{high}$ ).

state in Figure 5.

*Case 2: Decreasing the propensity to shoot.* In this case, fewer shots will be taken than in the original policy. Hence, the xG estimate for shooting in each state is too low. We counter this by modifying the xG values in each state where shooting has decreased and award them an xG of:

$$P(s, \text{shoot}, \text{goal}) + (\mu_s^{high} - \mu_s) \quad (27)$$

where  $\mu_s$  is again the average computed xG of all shots occurring in state  $s$  and  $\mu_s^{high}$  is the average computed xG of the  $1-x$  highest-quality shots occurring in state  $s$ , when decreasing the propensity to shoot by  $x$  percent. Thus, we drop the  $x$  percent lowest-quality shots and compute the average xG of the remaining shots. We illustrate this for  $x = 10\%$  in Figure 5.

## 7. Experimental Setup and Evaluation

In this section, we provide details on the experimental setup for learning the MDP. Next, we evaluate the proposed learning approaches to answer the following three questions: (1) whether the intended end locations can be accurately predicted by our proposed approach, (2) whether the proposed hierarchical Bayesian approaches better capture the data of each team than approaches without a global prior, and (3) whether the final learned MDP captures the same information as is available in the event data of the team. These evaluations allow us to validate whether the learned MDP faithfully mimics the observed behavior of the team.

Table 1: Average number of actions that are passes, dribbles, crosses, or shots per team in each season for each of the considered competitions.

Competition	2018/19	2019/20
Bundesliga	22,166	22,190
LaLiga	23,154	22,864
Premier League	24,584	24,946

## 7.1 Data Sets

Our analyses focus on the event stream data of the 2018/19 and 2019/20 seasons of the English Premier League, German Bundesliga, and Spanish LaLiga. Each of these leagues employs a format where each team plays every other team twice: once at home and once away. The English and Spanish competitions contain 20 teams, resulting in a total of 380 matches per season. The German competition contains 18 teams, resulting in a total of 306 matches per season. Table 1 shows, for each season and competition, the average number of considered actions (i.e., pass, dribble, cross, shot) per team to construct the models. We use a team’s actions during both seasons to learn their MDP. In our analyses, we focus on 20 teams, including 7 English Premier League teams (Burnley, Chelsea, Everton, Liverpool, Manchester City, Manchester United, and Newcastle), 7 German Bundesliga teams (Bayern Munich, Dortmund, Leipzig, Hoffenheim, Mönchengladbach, Schalke, and Wolfsburg), and 6 Spanish LaLiga teams (Atlético Madrid, Barcelona, Eibar, Real Madrid, Sevilla, and Valencia).

## 7.2 Transient State Spaces

We consider two different sets of transient states (Figure 6), which are tailored to the use cases.

**An offensive grid.** This grid divides the offensive half of the field into zones of  $3\text{m} \times 3\text{m}$  and captures the defensive half in one state (Figure 6a). This set of states is very fine-grained around the opponent’s goal area and will be used in the use cases to analyze a team’s shot policy.

**A custom grid.** This grid is more fine-grained where chances of scoring are higher and more coarse-grained where the chances of scoring are lower. These states also capture the differences between locations in the defensive half (Figure 6b). The size of the states ensures that there is sufficient data available in each state to analyze the different movement strategies and their effects. Additionally, the states also correspond to zones that are frequently used by practitioners when discussing movement strategies (e.g., the final third and their entry zones, the penalty box, half-spaces). Therefore, this set of states will be used in the use cases to analyze (defensive) movement policies.

## 7.3 Learning the MDP

Learning each team’s MDP consists of two steps. In the first step, the proposed team-specific models to predict the end location of passes and crosses (Section 4.1) are trained




Figure 6: Illustration of two sets of transient states that are used. Each transient state is determined by their location on the field. Teams play left to right. In (a), the defensive half denotes one state, the offensive half is split into zones of  $3\text{m} \times 3\text{m}$ . This state space is more suited for analyzing shot strategies as it is more fine-grained around the opponent’s goal area. In (b), the field is divided into custom zones, which are more suited for analyzing movement strategies.

using XGBoost (Chen & Guestrin, 2016).<sup>8</sup> The training data consists of the successful passes or crosses performed by the team. This data set is split into a train and test set using a 70-30 split and 5-fold cross-validation is performed on the train set to tune XGBoost’s hyperparameters. The test set is used for evaluation. For crosses, a radius  $r = 2.5$  meters is used, which is a reasonable error margin for professional soccer players. Appendix A.1 contains more information on the experimental setup. After training, the intended end locations of failed actions are predicted for each team that played during both seasons in the same league. The predictions of all teams in the same league over both seasons are used to create the prior for the Bayesian approach.

In the second step, the team-specific MDPs are learned using a hierarchical Bayesian approach. The probabilistic programming package PyMC3 (Salvatier, Wiecki, & Fonnesbeck, 2016) is used to model each of the Bayesian models (Section 4.2).<sup>9</sup> We use PyMC3’s Auto-Differentiation Variational Inference (ADVI) implementation to train the models. This allows us to scale to large data sets. All pass, cross, dribble, and shot actions in a team’s data set are used to train the models. For failed actions, there are multiple possible intended end locations in the data set, each with its own predicted probability. This means that each failed example is represented multiple times in the data set, once for each possible end location. To ensure that failed examples are not overly represented in the data set and to not count these examples more than once, we weigh each observation when passing it as input to the Bayesian models. The weight used for each failed observation is taken from the probability distribution over all possible intended end locations. For successful examples, the weight is set to one, as their intended end location is known. This ensures that failed and successful examples are weighted equally. To fit the transition function models (policy models) until convergence, 50,000 (40,000) iterations are used. Each model is then sampled

8. See <https://xgboost.ai/>. We used XGBoost v1.5.

9. See <https://docs.pymc.io/>. We used PyMC3 v3.11.

Table 2: AUROC and Brier scores for the end location prediction models averaged ( $\pm 1$  std) over all teams that played during both 2018/19 and 2019/20 seasons in the English Premier League, German Bundesliga, or Spanish LaLiga. Results are shown for each state space (OFF = offensive grid, CUSTOM = custom grid). The baseline predicts the distribution over observed end locations in the training set. DK refers to the post-processing step using domain knowledge. For AUROC, higher values are better. For the Brier score, lower values are better.

State Space	Action	Model	AUROC	Brier Score
OFF	Pass	Baseline	0.50 ( $\pm 0.00$ )	0.73 ( $\pm 0.04$ )
		XGBoost	0.73 ( $\pm 0.04$ )	0.57 ( $\pm 0.04$ )
		DK	<b>0.99 (<math>\pm 0.00</math>)</b>	<b>0.51 (<math>\pm 0.03</math>)</b>
	Cross	Baseline	0.50 ( $\pm 0.00$ )	0.94 ( $\pm 0.01$ )
		XGBoost	0.52 ( $\pm 0.03$ )	0.97 ( $\pm 0.04$ )
		DK	<b>0.92 (<math>\pm 0.02</math>)</b>	<b>0.78 (<math>\pm 0.07</math>)</b>
CUSTOM	Pass	Baseline	0.50 ( $\pm 0.00$ )	0.98 ( $\pm 0.00$ )
		XGBoost	0.93 ( $\pm 0.01$ )	0.80 ( $\pm 0.02$ )
		DK	<b>0.99 (<math>\pm 0.00</math>)</b>	<b>0.52 (<math>\pm 0.02</math>)</b>
	Cross	Baseline	0.50 ( $\pm 0.00$ )	0.87 ( $\pm 0.02$ )
		XGBoost	0.53 ( $\pm 0.03$ )	0.89 ( $\pm 0.03$ )
		DK	<b>0.96 (<math>\pm 0.01</math>)</b>	<b>0.55 (<math>\pm 0.05</math>)</b>

4,000 times and the average of these samples is used to compute the final probabilities. Appendix A.2 contains more information on the experimental setup.

#### 7.4 Evaluation of the Intended End Location Predictions

The held-out test set of each team’s successful actions is used to evaluate the performance of the end location prediction models. Table 2 reports the area under the ROC curve (AUROC) and the original Brier score (Brier, 1950) averaged across all teams that played during both the 2018/19 and 2019/20 seasons in the English Premier League, German Bundesliga, or Spanish LaLiga (i.e., 49 teams in total) and for both proposed state spaces. The model’s ability to distinguish between classes is measured by AUROC. Whether the model’s probability estimates are well-calibrated is measured by the Brier score. In this case, calibration is more important as the estimates are used to learn the transition model of the MDP. The baseline model corresponds to naively predicting a prior probability of an action ending in each state, which can be retrieved from the data. The results show that the learned models outperform the baseline and that using domain knowledge to post-process the predictions substantially improves performance.

#### 7.5 Evaluation of the Bayesian Approach

Evaluating the predictive accuracy of the Bayesian models is done by computing the expected log pointwise predictive density using Pareto-smoothed importance sampling leave-one-out cross-validation (PSIS-LOO-CV).<sup>10</sup> As training the models can be quite time-

10. See <https://arviz-devs.github.io/arviz/>. We used the PSIS-LOO-CV implementation from Arviz v2.0.

Table 3: Expected log pointwise predictive density (in thousands) using PSIS-LOO-CV for each proposed hierarchical model and state space (OFF = offensive grid, CUSTOM = custom grid). The results are averaged ( $\pm$  standard error) over all 20 teams that were considered (7 Premier League, 7 Bundesliga, and 6 LaLiga teams). We compare the results against non-hierarchical versions of the models. Higher values indicate models with better predictive accuracy.

State Space	Model	$\pi(\cdot)$	$P(\cdot)$
OFF	Unpooled	-180.1 ( $\pm$ 0.73)	-27.1 ( $\pm$ 0.13)
	Hierarchical	<b>-161.4 (<math>\pm</math> 0.68)</b>	<b>-26.2 (<math>\pm</math> 0.28)</b>
CUSTOM	Unpooled	-159.1 ( $\pm$ 0.29)	-19.4 ( $\pm$ 0.14)
	Hierarchical	<b>-157.5 (<math>\pm</math> 0.29)</b>	<b>-18.9 (<math>\pm</math> 0.18)</b>

consuming depending on the used state space, this approach can be used to efficiently estimate the expected log pointwise predictive density without the need for refitting and evaluating each model for each left out sample.

Table 3 shows the estimated log pointwise predictive density values for our proposed Bayesian models, averaged over all 20 teams for which an MDP was constructed and for both state spaces. We compare this to the results of a non-hierarchical approach in which the prior means  $\mu_{s,a}$  and  $\alpha_{s,a}$  are estimated for each team separately using a prior normal distribution with zero mean and standard deviation 5.0. The results show that the hierarchical models outperform the non-hierarchical approaches for both the policy and transition probability models and for both state spaces. The ability to use information about a global prior helps with resolving sparsity issues and in turn increases predictive accuracy.

## 7.6 Evaluation of the Complete Model

Evaluating the correctness of the complete MDP (i.e., that it faithfully captures the behavior of the team) is not straightforward as there is no ground-truth model against which can be compared. Instead, we propose to evaluate the model on a soccer-relevant metric. We compute the probability of eventually scoring from each state and compare (1) the results when computing these values using the model, and (2) the values obtained when only using the raw event stream data. This metric evaluates all learned parts of the MDP because all learned probabilities in the model (i.e., the transition function and policy) are needed to calculate the probability of scoring from each state. That is because teams can use both movement actions and shot actions during their possession sequences to score goals, and each action can either fail or succeed.

Computing these probabilities from the model can be done by using the value function. Each state’s value is equal to the probability of eventually scoring as only scoring results in a reward. Computing these probabilities empirically from the data can be done by identifying possession sequences that lead to a goal later on and computing the fraction of such possession sequences over all possession sequences per state.

Averaged over the 20 teams for which an MDP was constructed, the mean absolute error ( $\pm$  1 std) between the results from the model and the empirical results is 0.021 ( $\pm$  0.002) using the offensive grid (Figure 6a) and 0.008 ( $\pm$  0.002) using the custom grid (Figure 6b).



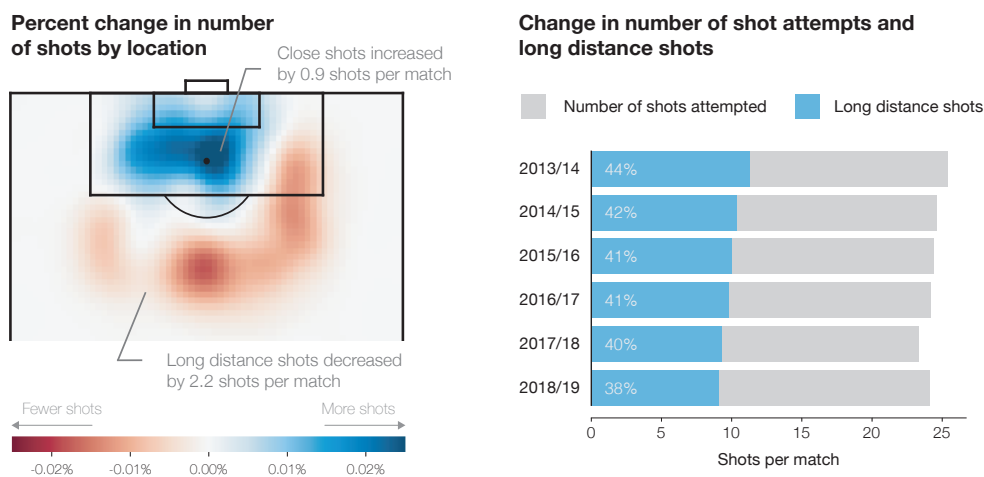


Figure 7: Evolution of shooting in the English Premier League between 2013/14 and 2018/19. A long-distance shot is defined as a shot taken from outside the penalty box. Over the course of this time period, the number of these shots has declined by around 20%.

This comparison illustrates that the models provide fairly accurate estimates of derived results. The small difference between the results from the model and the results obtained from the raw data is possibly due to the fact that the Markov model approach generalizes over all possible sequences to reach the goal, whereas not all these sequences are observed in the data set. There is also a difference visible between both state spaces. As could already be seen in the previous two evaluations, the results using the custom grid were slightly better than the results using the offensive grid, which can be attributed to the offensive grid being more fine-grained and containing fewer data in each state. Naturally, the ability to better predict the intended end locations and estimate the transition model for one state space over the other propagates to the quality of the final model.

## 8. Use Case: Reasoning About Shot Policies

A key trend in soccer that could be attributed to the rise of analytics is a change in teams' shooting policies: teams shoot much less frequently from outside the penalty box, which is illustrated in Figure 7 for the English Premier League. This has been driven by analysis of the xG metric that has shown that because long-distance shots have a much lower chance of resulting in a goal than closer ones, a (slightly) smaller number of high-quality attempts will likely yield more goals than a slew of low-quality ones. In this use case, we use the reasoning tools developed in this work to question this new conventional wisdom and analyze whether teams are now shooting *too infrequently from long distance*.

To gain better insights into long-distance shooting, we perform four analyses. First, we investigate in which situations players should forgo a shot outside the penalty box because the near-term chance of scoring is higher (e.g., by passing it to a teammate who then shoots). Second, we extend our first analysis by looking further into the future and aim to evaluate the odds of the team ever generating a better shooting chance than the one they




Figure 8: Illustration of the state space and interesting areas used to analyze shot policies. The long-distance zone from which shots can occur is shown in blue. The flanks to which the ball is often played to open up space are shaded in gray.

have now. Third, we reason about the possible effects of uniformly increasing or decreasing the number of long-distance shots on the expected number of goals a team would score in a season. Fourth, we perform a similar analysis as the third one but focus on a targeted increase instead of a uniform increase. For all analyses, we use the *offensive grid* (Figure 6) as transient states.

### 8.1 Shoot Immediately or Move Before Shooting?

Consider a player possessing the ball a few meters outside the penalty box with two choices:

1. Take the shot now from the present location;
2. Try to move the ball to another location (e.g., pass it to a teammate) and shoot from the new location.

What decision should the player make? Common sense dictates that the best option is the one that has the highest chance of generating a goal. However, knowing which choice is best according to this criterion involves a trade-off. Taking the shot now ensures the possibility of scoring the goal. Forgoing the shot may generate a better opportunity down the line, but does entail some risk. The pass may be errant or your teammate may miscontrol it, leading to the shooting opportunity evaporating. While the best decision will clearly be context-specific (e.g., how much pressure is the shooter under, where is the teammate located), we can use the techniques from Section 5.1 to get a general sense of when it is preferable to shoot and when it is better to pass.

First, we consider two specific movement conditions: (1) there is exactly one attempted move action after which a shot is attempted, and (2) there are exactly two move actions to any field location after which a shot is attempted. We use the approach of Section 5.1 to calculate the exact probability of scoring in both scenarios for each state inside the blue-colored box in Figure 8. Figure 9 visualizes for each movement scenario the difference in xG

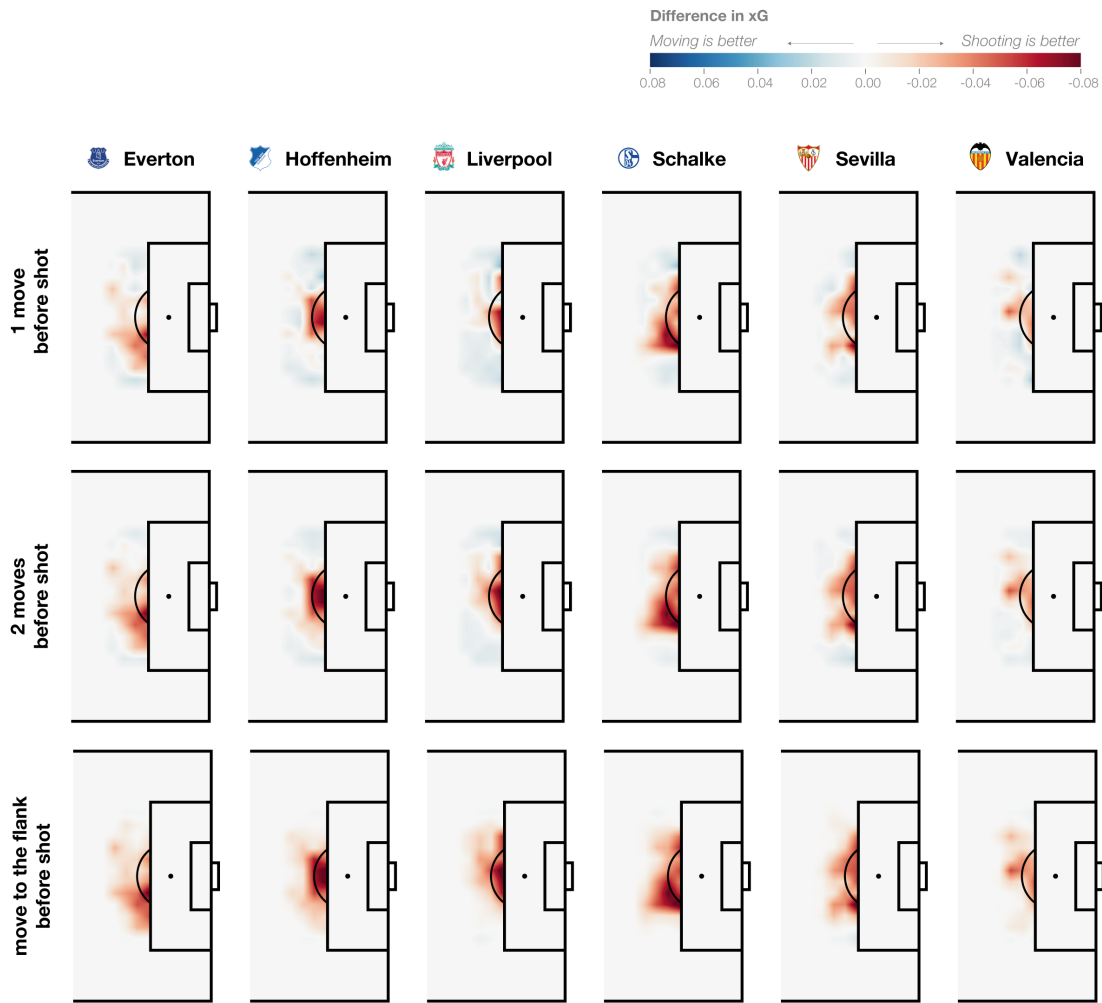


Figure 9: The difference in xG between moving once prior to shooting (top), moving exactly twice to any location prior to shooting (middle), or moving twice prior to shooting with the constraint that the first move action must move the ball to one of the flanks (bottom) versus directly shooting. Red indicates where immediately shooting is the better choice as moving would decrease your odds of scoring. Blue indicates when moving would optimize your chances of scoring.

between that scenario and directly shooting for the following teams: Everton, Hoffenheim, Liverpool, Schalke, Sevilla, and Valencia. These results show that for each team there are zones where moving is better and also zones where shooting is better. Interestingly, the differences in benefits of shooting and moving are not symmetric. When shooting is preferred, the payoff in terms of increase in the chance of scoring is much higher than when moving is preferred. The locations where shooting is preferred vary by team, but commonalities do arise. Zones that are in the front of the center of the goal are generally regarded as good locations to shoot from. Intuitively, this makes sense because when a

player attempts a forward central pass from this location, the receiving player almost always needs to turn before being able to shoot. Thus, immediately shooting is a better choice than a pass followed by a subsequent shot. Additionally, locations a bit further out and to the left/right of the penalty arc can also be spotted as good shooting locations for some teams. For example, for Everton, this could be due to their inverted winger Yannick Bolasie, who can cut infield to get the ball on his dominant foot, which facilitates shooting. When increasing the number of required move actions prior to shooting, the benefit of moving decreases (see middle row in Figure 9). Hence, in some zones shooting is now preferable whereas moving was better in the one-move scenario. This is most clearly visible at the border of the penalty box. In these states, performing more move actions increases the odds of losing the ball and consequently missing out on the chance of scoring. Concretely, it is preferable to take an available shot in these zones, unless you can move directly (i.e., in one step) to a player who will be able to shoot.

Next, we further restrict the second movement scenario by requiring that the first action in the sequence must move the ball to one of the flanks, which are defined as the gray-shaded regions in Figure 8. Often the flanks are used to open up space in the hopes of generating a better shot later in the possession sequence. This analysis also highlights the power and flexibility of probabilistic model checking to incorporate and reason about various constraints on an MDP. The bottom row in Figure 9 visualizes the difference in xG between this scenario and directly shooting for the same six teams. Imposing the additional restriction of first moving the ball to the flank decreases the chance of generating a goal later on and hence makes shooting more advantageous in many areas.

These types of analyses and the insights that they provide can aid teams in several ways. Offensively, it can help teams provide concrete advice to players about the desired decisions to make in a particular part of the pitch when specific situations arise. Defensively, it gives teams an indication of how an opponent should behave, which they can use to help craft a plan to disrupt them. For example, when deciding upon the line-up for a game, a team playing against Everton might want to use a holding defensive midfielder with the aim of clogging the area to the left of the penalty arc in order to force them to move the ball to the flank instead of allowing them to shoot.

## 8.2 Probability of Generating a Better Shot Later in the Possession

To complement the analysis in the prior section, we look further into the future and answer the question: “For a specific location, what is the probability of ever generating a better shot than the one they have when outside the box?”. To answer this question, we apply the techniques of Section 5.1 and calculate, for each long-distance state  $s$ , the probability of generating a shot with a higher xG value than the one associated with shooting from  $s$ .

Figure 10 shows these probabilities for the same six teams as in the previous analyses. Our method identifies various regions with higher and lower odds of generating a better shot, which mostly correspond to the same areas identified in the previous analyses. Interestingly, there are many locations where teams are very unlikely to ever generate a better shot. In fact, the probability of ever generating a better shot can be as low as 5% for some locations. An example of such a location for Chelsea and Schalke is the region to the right of the penalty arc. Similarly, when a team is on the edge of the penalty box, directly in front

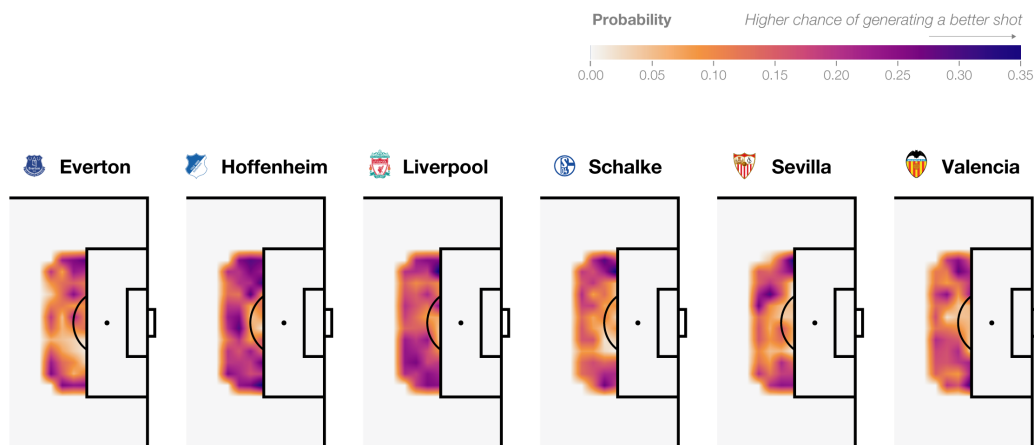


Figure 10: The probability of ever generating a better shot in the same possession sequence for Everton, Hoffenheim, Liverpool, Schalke, Sevilla, and Valencia. The regions with the lowest probabilities correspond to the zones where shooting was preferable to moving in the previous analyses. The magnitudes of these probabilities vary by team.

of goal, the chances of generating a better shot are only between 5% and 10%. In other locations, like Sevilla’s left side of attack, the probabilities can be quite high, at around 35%. The magnitude of the probabilities can vary substantially from team to team. For example, Hoffenheim is almost always more likely to generate a better shot, whereas Everton clearly has a large region from which immediately shooting would be preferred.

### 8.3 Uniformly Shooting More or Less From Distance

Practitioners often wonder what would happen if players decided to alter their shot volume during games. To answer such questions we use the reasoning techniques outlined in Section 6.2. Specifically, we explore the effect of increasing and decreasing the frequency of shooting from long-distance by 5%, 10%, and 20%.

Figure 11 shows the change in the expected number of goals each of the 20 considered teams would score over the course of a season as a result of altering the frequency of shooting from distance. For most teams, we see that shooting less from distance leads to a decrease in the number of goals a team would be expected to score. Uniformly increasing the number of long-distance shots would yield more goals for half of the teams. In each league, there are some exceptions that would see the opposite happening. For example, the previous analyses identified Hoffenheim as a team that almost always is more likely to generate a better shot later on in the possession. Thus, increasing the number of shots they take from long-distance is counterproductive. Other teams like Atlético, Bayern, and Chelsea can be identified as the teams that can expect bigger increases. All three teams have or had players with a good long-distance shot (e.g., Correa and Koke with Atlético, Lewandowski and Gnabry with Bayern, and Hazard with Chelsea).




Figure 11: Effect on the expected number of goals a team would score when uniformly increasing or decreasing the frequency of shooting from the long-distance region shown in Figure 8 by 5%, 10%, and 20% for all 20 considered teams in the English Premier League, Spanish LaLiga, and German Bundesliga. Most teams see an increase in the number of goals with an increase in shooting from distance. Some exceptions, like Manchester City and Newcastle, would see a decrease.

### 8.4 Targeted Increases of Shots From Distance

The previous analysis evaluated the effect of uniformly increasing a team’s propensity to shoot from all the considered long-distance locations. However, the first and second analyses clearly illustrate that there are a limited number of team-specific long-distance zones where it may be fruitful to consider shooting more often. Therefore, in this analysis we explore the effect of each team shooting 5%, 10%, or 20% more often but only from those locations where shooting was deemed to be the better choice than moving for the team.

Figure 12 shows the change in the expected number of goals each of the 20 considered teams would score over the course of a season. This more targeted approach now yields increases in the expected number of goals scored for almost all teams. Shooting 20% more often from long-distance tends to yield a gain of almost one extra goal for half of the teams. Figure 13 shows that every goal scored during a season equates to roughly one point in the table. Thus, scoring an extra goal is possibly very important as for the bottom teams it could mean the difference between being relegated or staying up. Relegation is a catastrophe

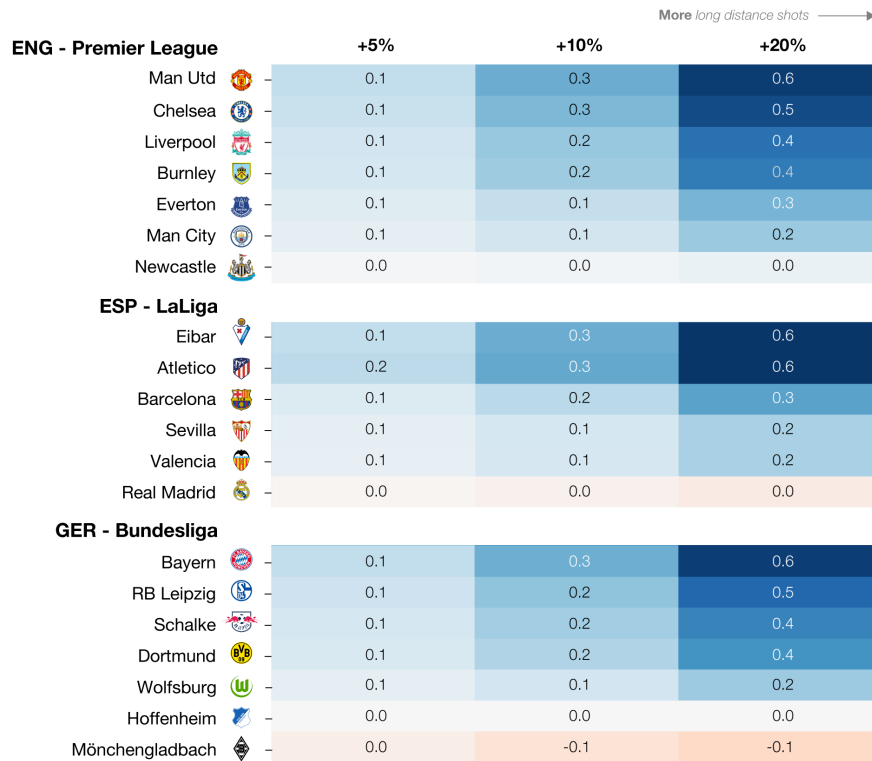


Figure 12: Effect on the expected number of goals a team would score when performing a targeted increase in the frequency of shooting from the long-distance region shown in Figure 8 by 5%, 10%, and 20% for all 20 considered teams in the English Premier League, Spanish LaLiga, and German Bundesliga. Half of the teams would score almost an extra goal. Interestingly, Mönchengladbach would still see a decrease.

with a cost that can be estimated at around \$250 million.<sup>11</sup> For the top teams, an extra goal could mean the difference between qualifying for the Champions League or not.

## 9. Use Case: Reasoning About Defensive Strategies

In this use case, we reason about the efficacy of various defensive strategies in terms of forcing teams to avoid states from which they create dangerous situations. Specifically, we address the following three questions:

1. Which states should a team be forced to avoid to decrease the number of shots they take?

11. <https://www.firstpost.com/sports/premier-league-250m-at-stake-as-aston-villa-bournemouth-and-watford-fight-to-avoid-relegation-on-final-day-8641391.html>




Figure 13: Points in the final table as a function of the goals scored. The image visualizes the data of ten seasons of the English Premier League (i.e., 2010/11 until 2019/20). Every goal scored equates to one point in the final table.

2. Which defensive and midfield states should the defensive game plan suppress movement to in order to decrease the general danger created by the attacking team during the build-up phase?
3. How effective does the identified movement suppression tactic remain after the opposing team adapts to it?

For these analyses, we use the *custom grid* (Figure 6) as transient states.

### 9.1 Shot Suppression

In this first analysis, we will use the methodology outlined in Section 5.2 to identify which states a defending team should pay attention to in order to *directly* and *indirectly* decrease their opponents' shots.

Figure 14 visualizes the percent decrease in the probability of taking a shot from and reaching the favored *shot\_locations* (see Figure 4) for each state and for each of the four considered teams. For Chelsea, when directly suppressing their shots, the most important states lie centrally with a slightly higher decrease on their left side versus their right side. Their most important states shift entirely to their left side when focused on indirectly suppressing their shots. An opponent that prevents them from entering their most important state would decrease Chelsea's chance of shooting by almost 10% and their chance of reaching the favored shot locations by almost 17%. A similar tendency in terms of locations can be found for Real Madrid. Both teams predominantly tend to use the left flank to get the ball centrally and create chances. An opponent that prevents Real Madrid from entering their most important state would decrease their chance of shooting by a little over 9% and their chance of reaching the favored shot locations by almost 17%. One possible explanation for this similar tendency in favored locations for both teams is Eden Hazard who played as a left-winger for both teams during the two considered seasons (i.e., during the 2018/19 season he played for Chelsea, he transferred to Real Madrid for the 2019/20 season). He is known for his impeccable ball-handling skills which allow him to pass many



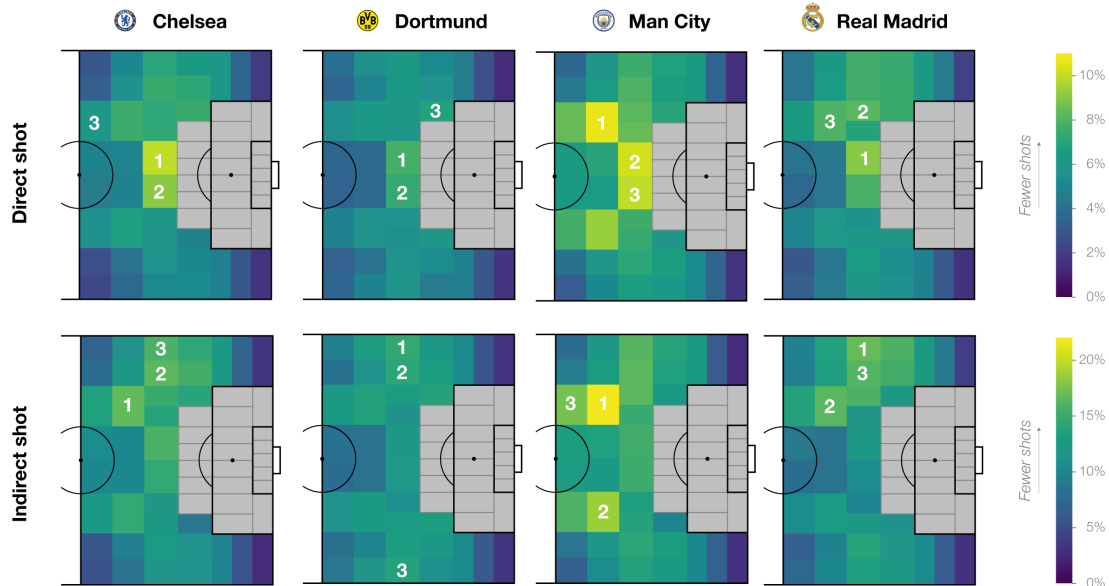


Figure 14: The percent decrease in shooting from (top row) and reaching (bottom row) the common shot locations for Chelsea, Dortmund, Manchester City, and Real Madrid. Yellow shading indicates states that result in a large decrease when the given team is forced to avoid them. Dark blue shading indicates states that result in a smaller decrease. The top three states with the largest impact are labeled in each figure.

defenders and create danger on his left side of the pitch. For Dortmund, their most important states lie centrally when focused on directly suppressing their shots, but on either flank when focused on indirectly suppressing their shots. This is most likely due to their wingers Thorgan Hazard and Jadon Sancho who provide width on the flanks and are very adept at creating assists. An opponent that can prevent them from entering their most important state would decrease Dortmund’s chance of shooting by almost 8% and their chance of reaching the favored shot locations by almost 15%. Finally, for Manchester City, their most important states lie centrally in both cases, with the flanks having a smaller effect for them. This indicates that they predominantly use the center to create chances, whereas the other teams tend to use the flanks more. Tactically, Manchester City is a possession-based team that gradually builds up their attack. Their creative midfielders Silva and De Bruyne, who are very adept at reading and steering the game, can often be found in the center of the pitch. If an opponent can prevent them from entering their most important states, we see the biggest decreases of all four teams (i.e., a decrease of almost 11% in the chance of shooting and a decrease of almost 22% in the chance of reaching the favored shot locations).

## 9.2 Movement Suppression

As a second analysis, we take a step back and analyze which sets of states the defending team should force the opponent to avoid in order to decrease the danger created by the opponent during their build-up phase. Doing so will essentially restrict their movement

Prevent the team to reach the **highlighted states**  
 to decrease their chances of scoring  
 from **final third entries** by 10% and  
 from **the middle third** by 1%

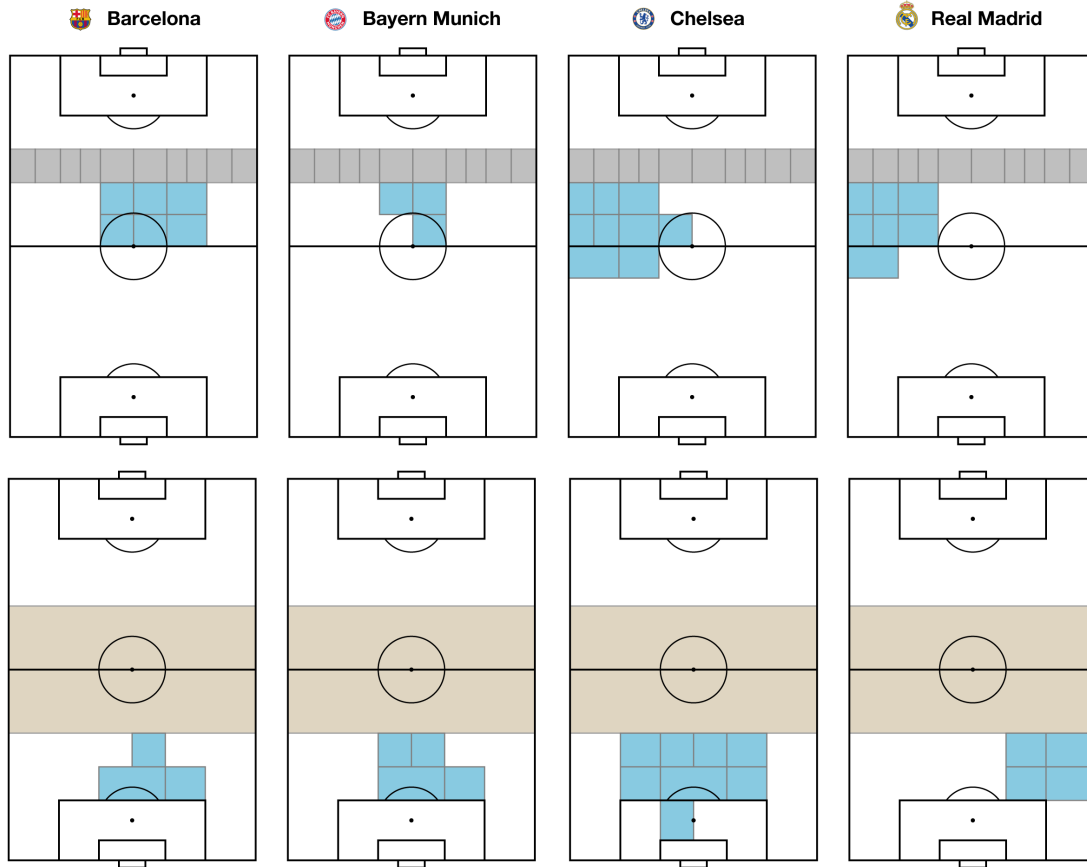


Figure 15: Illustrates for Barcelona, Bayern Munich, Chelsea, and Real Madrid, two areas (blue) to prevent them from reaching in order to decrease their chances of scoring in each final third entry state by at least 10% and in each middle third state by at least 1%.

with the ball and will decrease the danger created by the attacking team. Here, we will use the methodology outlined in Section 5.2 and focus on decreasing the attacking team's danger in the final third entry and middle third regions of the pitch.

Figure 15 shows the results for Barcelona, Bayern Munich, Chelsea, and Real Madrid. To reduce Barcelona's chance of scoring from each of the final third entry states by at least 10%, a crucial area to avoid lies around the center-right of the pitch which is where play-makers Ivan Rakitić and Frenkie de Jong, and winger-stringer Lionel Messi often operate. Decreasing their chance of scoring by at least 1% in each state in the middle third of the field can also be done by forcing them to avoid the center-right of their defensive third. Similarly, for Bayern, it is also best to force them to avoid the center of the pitch in both cases. In the defensive part of the pitch, this is the playing field of Jérôme Boateng and

Table 4: Percent decrease in each team’s probability of scoring when it is forced to avoid its respective region in the top row of Figure 15. The results are shown both before and after adapting to the defensive strategy, and for both approaches of adapting the policy (PR = proportional approach, SA = same aggressiveness approach).

Team	Without adapting	PR	SA
Barcelona	15.7	2.8	3.5
Bayern Munich	7.7	1.5	1.9
Chelsea	12.1	4.2	3.2
Real Madrid	12.2	4.1	3.4

David Alaba, who both have good ball-handling skills and can act as playmakers. In the offensive part of the field, this is the area of midfielder Thiago Alcântara, who is involved in most of Bayern’s possession sequences. To reduce Chelsea’s chance of scoring from each of the final third entry states by at least 10%, a crucial area to avoid lies to their left side of the pitch. This area corresponds to the best locations for indirectly suppressing their shots, and a similar observation can be made for Real Madrid. When looking at their defensive half, the similarities between the two teams disappear. Decreasing Chelsea’s chance of scoring by at least 1% in each state in the middle third can be done by forcing them to avoid the center region of the defensive third. On the other hand, for Real Madrid, it is better to force them to avoid the right side of their defensive third.

### 9.3 Effect of Defensive Strategies Once Team Adapts

Until now, we have explored the effect of forcing a team to avoid certain regions (*area*) on its chance of scoring. In practice, if an opponent enacts such a strategy, a team will eventually react and adapt their old policy  $\pi$  towards a new one  $\pi'$ . The new policy  $\pi'$  will stop trying to reach locations in *area*, so  $\pi'(move\_to(s') \mid s) = 0$  for all  $s' \in area$ . The lost probability mass will then be redistributed using either the *proportional* or *same aggressiveness* approach (Section 6.1). We illustrate the results of both approaches for Barcelona, Bayern Munich, Chelsea, and Real Madrid when they are forced to avoid their respective area in the top row of Figure 15.

Table 4 shows the resulting decrease in each team’s chances of scoring before and after adapting using both approaches. If the team does not adapt, the incurred decrease in the chances of scoring can be quite large, depending on the size of the region. For example, forcing Barcelona to avoid their respective region in the middle of the field decreases their probability of scoring by 15.7%, which is almost double the decrease that Bayern Munich incurs. On the other hand, forcing Chelsea and Real Madrid to avoid their similar regions on the left side of the field incurs a similar decrease for both teams. When the teams adapt, using either strategy, the reduction becomes much smaller. Interestingly, for Barcelona and Bayern Munich, adjusting their policy using the *proportional* approach results in the least decrease, whereas for Chelsea and Real Madrid, using the *same aggressiveness* approach is best. This possibly indicates that for Barcelona and Bayern Munich it might be best to also move the ball backward to later find other ways of reaching a goal, whereas for Chelsea and Real Madrid, moving the ball not through the left side and more through the middle and

right side is also quite effective. While the decrease is less impressive when teams adapt, it still represents a reasonable reduction, certainly given that adapting one’s strategy is hard.

## 10. Related Work

Markov models and reinforcement learning paradigms have been widely used for analyzing sports such as soccer (Fernández et al., 2021; Hirotsu & Wright, 2002; Liu & Hohmann, 2013; Rudd, 2011; Singh, 2019; Van Roy, Robberechts, Decroos, & Davis, 2020; Yam, 2019), American football (Goldman & Rao, 2014), basketball (Cervone et al., 2016; Sandholtz & Bornn, 2018, 2020; Wang, Fox, Skaza, Linck, Singh, & Wiens, 2018), table tennis (Pfeiffer et al., 2010; Wenninger & Lames, 2016), and ice hockey (Routley & Schulte, 2015; Schulte, Khademi, Gholami, Zhao, Javan, & Desaulniers, 2017). The most prominent use case of these techniques is to objectively quantify a player’s contributions during a match. The intuition is that Markov models and reinforcement learning techniques enable assessing how much a player’s action increases his team’s chance of scoring in the near future. Such techniques have begun to have a significant impact in professional soccer, where clubs (e.g., Liverpool<sup>12</sup>, Barcelona (Fernández et al., 2021)) and companies are employing them to help in areas such as player acquisition and match analysis.

There has been less attention on using these models and techniques for reasoning about the strategies teams (could) employ to obtain insights and aid their in-game decision-making. Some examples of existing works in this setting are the analysis of set-pieces such as corners, free-kicks, and throw-ins (Rudd, 2011), identifying where teams create value from by using the values of players’ actions (Fernández et al., 2021; Singh, 2019), analyzing the performance relevance of certain actions (Liu & Hohmann, 2013; Pfeiffer et al., 2010; Wenninger & Lames, 2016), and the analysis of the optimal or an alternative policy (Sandholtz & Bornn, 2018, 2020; Wang et al., 2018).

Our proposed methods also correspond to this setting, with the most closely related work being that of Sandholtz and Bornn (2018, 2020) for basketball. There are several differences with our work, beyond looking at different sports. First, we consider a more expansive action space that allows us to model the intended end location of each action and reason about changes to the movement policy, whereas their methods cannot alter the movement policy and transition function separately. Second, when analysing the effects of modifying a team’s shot policy, we adjust both the frequency and efficiency of shots whereas their work did not take the frequency-efficiency trade-off into account. We also compute the expected number of goals teams would score via the MDP’s fundamental matrix as opposed to simulating the season by sampling from the various distributions. The idea of using the fundamental matrix to compute the probability of a sequence resulting in a goal has also been used by Yam (2019). Third, we analyze event data and focus on the analysis of teams, whereas Sandholtz and Bornn have access to tracking data which allows them to build a more fine-grained model in terms of player analysis. Finally, we propose to use probabilistic model checking techniques to reason about the merits of both offensive *and* defensive strategies. Applying model checking techniques to sports models has not been extensively explored. Dong et al. (2015) have applied probabilistic model checking to a tennis MDP to predict the win probability and identify a player’s best action to improve.

12. <https://freakonomics.com/podcast/london-live/>

To the best of our knowledge, the use of probabilistic model checking to compare the chances of scoring for different action sequences and assess the effect of defensive strategies is novel within sports analytics.

In contrast to real-life soccer, MDP’s have been extensively used in research on strategic reasoning and planning in (simulated) robot and humanoid soccer. For example, an approach for providing coaching advice in simulated robot soccer has been developed by Riley and Veloso (2004) and uses Q-learning to reason about a learned MDP. Liu et al. (2022) make use of multi-agent reinforcement learning in combination with imitation learning to train teams of physically simulated humanoids to play soccer. To train the humanoids, they specify multiple training tasks corresponding to soccer drills, which they model as stochastic games. Additionally, Ahmadi and Stone (2008) and Bai et al. (2012) investigate different automated action planning strategies for in-game decision-making. However, such automated planning strategies cannot immediately be used in real-life soccer as it is compounded by high stakes and the incredibly short tenure of most managers, which discourages experimentation. In addition, removing the randomness from a team’s behavior would make it very predictable for an opponent. Thus, instead of finding the optimal plan, our setting uses passive model-based reinforcement learning, where we aim to analyze a team’s current and counterfactual strategies to propose possible improvements and aid their tactical planning. The use of reinforcement learning for this task was also recently mentioned by Tuyls et al. (2021) in which they provide an overview of how the integration of reinforcement learning and statistical learning could in the future aid soccer teams in their decision-making.

Regarding addressing sparsity issues in Markovian sports models, various approaches have been employed in the literature. While including fine-grained location and context information would allow for a more detailed analysis, it also increases the sparsity of the data. Therefore, most approaches trade off usefulness and sparsity by applying a hand-crafted grid over the field and only sometimes include context such as players (Goldner, 2012; Luo, Schulte, & Poupart, 2020; Routley & Schulte, 2015; Rudd, 2011; Schulte et al., 2017; Singh, 2019; Van Roy, Robberechts, Yang, De Raedt, & Davis, 2021; Yam, 2019). Recent work in basketball goes one step further and applies Bayesian approaches on top of a fine-grained and contextualized state space to estimate the transition probabilities and thereby mitigates any remaining sparsity issues (Cervone et al., 2016; Sandholtz & Bornn, 2018, 2020). This approach has also been proposed in soccer, but without any validation on real-life data and with a rudimentary model containing only four states (Hirotzu & Wright, 2002). To the best of our knowledge, these approaches have not been applied to a more fine-grained Markov model for soccer in order to resolve sparsity issues.

Lastly, evaluation and validation of these models is often not straightforward and the approach taken typically depends on the use case (Davis, Bransen, Devos, Meert, Robberechts, Van Haaren, & Van Roy, 2022). In the case of valuing player actions, some approaches validate their ratings against other metrics such as salaries, goals, and assists (Liu & Schulte, 2018; Liu, Luo, Schulte, & Kharrat, 2020; Luo et al., 2020; Routley & Schulte, 2015; Schulte et al., 2017). While correlation with these metrics might indicate correctness, a perfect correlation indicates that the proposed approaches do not provide any new insights. As there are no ground truth labels available, it is hard to quantitatively validate the correctness and usefulness of these metrics and one often resorts to domain knowledge. In the case of vali-

dating the learned playing style, validation against the team’s actual playing style is equally challenging and mainly done with visual inspection or against domain knowledge (Schulte et al., 2017; Singh, 2019). Quantitative validation of the complete Markov model is often left out. However, in the case of intermediate models (e.g., Bayesian approaches), evaluation of those is often done to yield insight into the correctness of the final model (Cervone et al., 2016; Fernández et al., 2021; Sandholtz & Bornn, 2018, 2020).

## 11. Conclusion

This paper proposed a novel framework to learn and reason about the strategies used in professional soccer by combining techniques from machine learning and artificial intelligence. The proposed framework combines (1) learning a Markov decision process of a team’s offensive behavior from their observed data with (2) probabilistic model checking techniques to reason about the modeled policies. In order to learn the model, we showed how the missing intended end locations of failed actions can be learned using a combination of predictive modeling and domain knowledge. Second, we showed that a hierarchical Bayesian approach to learning the transition model can be used to mitigate sparsity issues. Third, we showed how the framework can be used to reason about both the current and counterfactual policies of teams. For this, we proposed (1) approaches to reason about the effect of both offensive and defensive strategies using probabilistic model checking, and (2) approaches to alter the policy of a team in a principled manner and evaluate its effects. This work is a successful illustration of how learning and reasoning can be combined.

We applied this framework to two use cases. First, we evaluated the current shot policy of teams and the effects of employing a different long-distance shot policy. We found that there are several team-specific areas where shooting is better than attempting to move the ball in the hope of generating a better shot down the line. If teams were to increase their frequency of shooting from distance from a limited number of specific areas, they would be expected to score about one additional goal over the course of a season. Given that each goal roughly equates to one point in the final league table, this could have very important implications regarding relegation from the league or qualification for the Champions League. Second, we analyzed the current offensive strategies of teams to optimize a defending team’s game plan. For each team, we found specific areas an opposing team should force them to avoid in order to decrease the opponent’s chances of conceding a goal, even if the attacking team were to adapt to this.

While there are no strong guarantees about the model’s correctness as would be required in a verification context, it clearly supports reasoning about strategies and policies in the complex dynamic environment of professional soccer. Furthermore, visualizing the results of the queries can help human soccer experts better understand the effects of potential strategies, which in turn contributes to trustworthy AI. From an application perspective, the proposed approaches can form a basis for future tactical analysis in soccer. The resulting insights can be used to better coach players during training sessions, when preparing for a specific opponent, or when setting a team’s tactics. We believe that our approach is also applicable to other environments and other sports. While using domain knowledge is clearly an important factor for successfully constructing the models, we believe the dependence on this domain knowledge is not a problem when transferring it to other sports as such domain

knowledge does exist in most other team sports (e.g., the various lines of research in the field of sports science).

Currently, our approach contains some limitations. The MDP is currently agnostic over its opponent, which means the resulting insights are quite general for each team. To obtain specialized insights when two specific teams play each other, the behavior of each team should be explicitly modeled (e.g., by using a stochastic game). Additionally, if the data allows, the MDP’s state space could be expanded to more contextual and realistic states, which would allow for a more in-depth team or player analysis. Finally, the approaches to evaluate the effects of alternative policies could be expanded by altering the assumption that the initial distribution of possession sequences remains unchanged. However, how to accurately estimate how this distribution should change in relation to the alterations in the policy remains an open question.

## Acknowledgments

This work has received funding from the European Union’s Horizon Europe Research and Innovation program under the grant agreement TUPLES No 101070149, the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” program, the Research Foundation – Flanders under EOS No. 30992574, and the KU Leuven Research Fund (C14/17/070, C14/18/062). Luc De Raedt is also funded by the Wallenberg AI, Autonomous Systems and Software Program (WASP) of the Knut and Alice Wallenberg Foundation.

## Appendix A. Experimental Setup

In this section, we provide all details of our experimental setup. All experiments in this paper were performed using Python v3.9.

### A.1 Gradient Boosted Trees Ensembles

To train the XGBoost models, the training data of each action type (i.e., all successful examples of passes or crosses) was split into a train and test set using a 70-30 split. We performed 5-fold cross-validation randomized search on each action type’s train set to optimize the hyperparameters of the XGBoost models. Therefore, 50 parameter settings were sampled from the following distributions:

```

‘max_depth’: random integer [2, 10],
‘min_child_weight’: random integer [1,11],
‘gamma’: uniform [0,1],
‘reg_alpha’: uniform [0,1],
‘reg_lambda’: uniform [0,10],
‘base_score’: uniform [0.1,1],
‘subsample’: uniform [0.5,1],
‘colsample_bytree’: uniform [0.5,1],
‘colsample_bylevel’: uniform [0.5,1]

```

Table 5: Training time for the transition function and policy models averaged over all 20 teams that were considered. The results are shown for each considered state space (OFF = offensive grid, CUSTOM = custom grid).

State Space	Model	Time (h)
OFF	Transition	6.3
	Policy	2.9
CUSTOM	Transition	1.5
	Policy	0.6

All other hyperparameters were set to their default values.

The final models were trained on a computing server running Ubuntu 20.04 with 16GB RAM and an Intel(R) Core(TM) i7-2600 CPU @ 3.40GHz. Training the ensembles for both passes and crosses for one team takes on average three minutes when using the offensive grid and one minute when using the custom grid.

## A.2 Hierarchical Bayesian Models

All parameters were set to their default values, except the number of iterations needed to fit and sample the models, as explained in the main part of this paper.

The models were trained on a computing server running Ubuntu 20.04 with 32GB RAM and an Intel(R) Xeon(R) CPU E3-1225 v3 @ 3.20GHz. The models were sampled on a computing server running Ubuntu 20.04 with 128GB RAM and an Intel(R) Xeon(R) Silver 4214 CPU @ 2.20GHz. Table 5 shows the training time needed for each of the considered state spaces and models.

## Appendix B. Policy Modifications

In this section, we provide the additional formulas for altering the movement policy and an evaluation of the approach for estimating the expected number of goals.

### B.1 Decreasing the Movement Probability

Suppose that, given a start state  $s$ , the objective is to move to a set of states *area*  $x$  percent less than the team normally would. Decreasing the probability of moving from state  $s$  to each state  $s' \in \textit{area}$  by  $x$  percent can be done as follows:

$$\pi'(\textit{move\_to}(s') \mid s) = (1 - x) * \pi(\textit{move\_to}(s') \mid s) \quad (28)$$

To maintain a probability distribution, the probability of moving from  $s$  to all other states  $s'' \in \mathcal{L} \setminus \textit{area}$  must increase.

Using the **proportional approach**, this is done as follows:

$$\pi'(\textit{move\_to}(s'') \mid s) = \frac{\pi(\textit{move\_to}(s'') \mid s)}{\textit{total\_other}} * (\textit{total\_other} + x * \textit{total\_area}) \quad (29)$$



Using the **same aggressiveness approach**, the probability of choosing to move from  $s$  to  $s'' \notin area$  is increased using the following approach. If  $s' \in area$  is a state that lies higher up the pitch than state  $s$ , the lost probability of choosing to move from  $s$  to  $s'$  is redistributed over all states  $s'' \in s^h \subseteq \mathcal{L} \setminus area$  that also lie higher up the pitch than state  $s$  as follow:

$$\pi'(move\_to(s'') | s) = \frac{\pi(move\_to(s'') | s)}{total} * (total + x * \pi(move\_to(s') | s)) \quad (30)$$

with  $total$  defined as:

$$total = \sum_{s'' \in s^h} \pi(move\_to(s'') | s) \quad (31)$$

For all other states  $s'' \in \mathcal{L} \setminus area$ , the policy remains unchanged. Similarly, if  $s' \in area$  lies at an equal height or lower on the pitch than state  $s$ , the lost probability of choosing to move from  $s$  to  $s'$  is redistributed over all other states  $s'' \in s^l \subset \mathcal{L} \setminus area$  that also lie at an equal height or lower on the pitch than state  $s$  using the same formulas. To keep the probability of moving forward equal, the policy of moving to the other states (i.e., higher up the pitch) then remains unchanged.

## B.2 Estimating the Expected Number of Goals

To assess the accuracy of our estimated expected number of goals, we calculate the relative error between (1) the estimated number of goals scored over a season according to our approach of Section 6 using the observed policy and (2) the actual average number of goals recorded per season in the event stream data of each team. The resulting average relative error over all constructed models using the 2018/19 and 2019/20 English Premier League, German Bundesliga, and Spanish LaLiga data is 14.02%. Given that an average team scores around 45 goals per season, this corresponds to roughly six goals. This indicates that our method produces fairly good estimates of the number of goals for the observed policy.

## Appendix C. Ablation Study

In this section, we provide an ablation study on the features used to predict the intended end location of actions. Therefore, we assess whether each feature provides an improvement on the proposed model’s performance. Specifically, the ablation models are created by removing one feature at a time, yielding four variations: (1) start state, direction, body part, start states of the three preceding actions (ALL); (2) start state, direction, body part (NO\_HISTORY); (3) start state, direction, start states of the three preceding actions (NO\_BODYPART); (4) start state, body part, start states of the three preceding actions (NO\_DIRECTION).

We train the models for all teams that played in both the 2018/19 and 2019/20 German Bundesliga using the same experimental setup as in Appendix A.1. We do not include the domain knowledge step and report the performance of the raw XGBoost models in order to assess the importance of each feature. Table 6 reports the average AUC and Brier scores. When removing the history features, the models incur a performance drop, especially when using the offensive state space. These features seem to be less important when using the

Table 6: AUROC and Brier scores for the ablation study on the end location prediction models. Results are averaged ( $\pm 1$  std) over all teams that played during both 2018/19 and 2019/20 seasons in the German Bundesliga and shown for each state space (OFF = offensive grid, CUSTOM = custom grid).

State Space	Action	Ablation Model	AUROC	Brier Score
OFF	Pass	ALL	0.714 ( $\pm 0.033$ )	0.535 ( $\pm 0.036$ )
		NO_HISTORY	0.683 ( $\pm 0.069$ )	0.594 ( $\pm 0.157$ )
		NO_BODYPART	0.711 ( $\pm 0.030$ )	0.534 ( $\pm 0.037$ )
		NO_DIRECTION	0.537 ( $\pm 0.019$ )	0.891 ( $\pm 0.119$ )
	Cross	ALL	0.511 ( $\pm 0.025$ )	0.961 ( $\pm 0.040$ )
		NO_HISTORY	0.510 ( $\pm 0.043$ )	0.965 ( $\pm 0.034$ )
		NO_BODYPART	0.521 ( $\pm 0.020$ )	0.969 ( $\pm 0.039$ )
		NO_DIRECTION	0.515 ( $\pm 0.026$ )	0.966 ( $\pm 0.037$ )
CUSTOM	Pass	ALL	0.925 ( $\pm 0.011$ )	0.799 ( $\pm 0.019$ )
		NO_HISTORY	0.922 ( $\pm 0.010$ )	0.785 ( $\pm 0.017$ )
		NO_BODYPART	0.926 ( $\pm 0.011$ )	0.800 ( $\pm 0.019$ )
		NO_DIRECTION	0.840 ( $\pm 0.019$ )	0.945 ( $\pm 0.006$ )
	Cross	ALL	0.511 ( $\pm 0.039$ )	0.882 ( $\pm 0.027$ )
		NO_HISTORY	0.516 ( $\pm 0.025$ )	0.867 ( $\pm 0.023$ )
		NO_BODYPART	0.514 ( $\pm 0.037$ )	0.878 ( $\pm 0.023$ )
		NO_DIRECTION	0.502 ( $\pm 0.022$ )	0.889 ( $\pm 0.038$ )

custom state space. Most likely, the historical information is useful when trying to pinpoint a fine-grained location which is necessary when using the offensive state space. The direction of the action seems to be especially important for passes, regardless of the state space. As all crosses typically end within a similar region in front of the goal, the direction information is less important. Finally, using the body part as a feature seems to have no pronounced effect. This is likely due to the fact that most actions are executed with a player’s feet and only a small number are executed with another body part. Overall, we find that the relative importance of the features depends on the type of action and state space. To cover most of the situations, we have used all features for all models in this paper.

## References

- Ahmadi, M., & Stone, P. (2008). Instance-based action models for fast action planning. In *RoboCup 2007: Robot Soccer World Cup XI*, Vol. 5001 of *Lecture Notes in Computer Science*, pp. 1–16.
- Bai, A., Wu, F., & Chen, X. (2012). Towards a principled solution to simulated robot soccer. In *RoboCup 2012: Robot Soccer World Cup XVI*, Vol. 7500 of *Lecture Notes in Computer Science*, pp. 141–153.
- Baier, C. (1998). *On algorithmic verification methods for probabilistic systems*. Habilitation thesis, Faculty of Mathematics and Informatics, University of Mannheim.
- Baier, C., & Katoen, J.-P. (2008). *Principles of Model Checking*. The MIT Press.
- Bellman, R. (1966). Dynamic programming. *Science*, 153(3731), 34–37.

- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1), 1–3.
- Caley, M. (2015). Premier league projections and new expected goals. Available at <https://cartilagefreecaptain.sbnation.com/2015/10/19/9295905/premier-league-projections-and-new-expected-goals>.
- Cervone, D., D'Amour, A., Bornn, L., & Goldsberry, K. (2016). A multiresolution stochastic process model for predicting basketball possession outcomes. *Journal of the American Statistical Association*, 111(514), 585–599.
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794.
- Davis, J., Bransen, L., Devos, L., Meert, W., Robberechts, P., Van Haaren, J., & Van Roy, M. (2022). Evaluating sports analytics models: Challenges, approaches, and lessons learned. In *Proceedings of the AI Evaluation beyond Metrics Workshop at IJCAI 2022*, EBEM, p. 1–11, Vienna, Austria.
- Decroos, T., Bransen, L., Van Haaren, J., & Davis, J. (2019). Actions speak louder than goals: Valuing player actions in soccer. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1851–1861.
- Decroos, T., & Davis, J. (2020). Interpretable prediction of goals in soccer. In *Proceedings of the 2020 AAAI Workshop on AI in Team Sports*, pp. 1–8.
- Dong, J. S., Shi, L., Chuong, L. V. N., Jiang, K., & Sun, J. (2015). Sports strategy analytics using probabilistic reasoning. In *Proceedings of the 20th International Conference on Engineering of Complex Computer Systems (ICECCS)*, pp. 182–185.
- Fernández, J., Bornn, L., & Cervone, D. (2021). A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. *Machine Learning*, 110(6), 1389–1427.
- Forejt, V., Kwiatkowska, M., Norman, G., & Parker, D. (2011). Automated verification techniques for probabilistic systems. In *Proceedings of Formal Methods for Eternal Networked Software Systems: 11th International School on Formal Methods for the Design of Computer, Communication and Software Systems*, pp. 53–113.
- Goldman, M., & Rao, J. M. (2014). Misperception of risk and incentives by experienced agents. Available at SSRN: <https://ssrn.com/abstract=2435551>.
- Goldner, K. (2012). A Markov model of football: Using stochastic processes to model a football drive. *Journal of Quantitative Analysis in Sports*, 8, 1–18.
- Green, S. (2012). Assessing the performance of Premier League goalscorers. Available at <https://www.statsperform.com/resource/assessing-the-performance-of-premier-league-goalscorers/>.
- Hensel, C., Junges, S., Katoen, J.-P., Quatmann, T., & Volk, M. (2022). The probabilistic model checker Storm. *International Journal on Software Tools for Technology Transfer*, 24(4), 589–610.

- Hirotsu, N., & Wright, M. (2002). Using a Markov process model of an association football match to determine the optimal timing of substitution and tactical decisions. *Journal of the Operational Research Society*, 53(1), 88–96.
- Ijtsma, S. (2015). A close look at my new expected goals model. Available at <http://11tegen11.net/2015/08/14/a-close-look-at-my-new-expected-goals-model/> via web.archive.org.
- Knutson, T. (2020). StatsBomb release expected goals with shot impact height. Available at <https://statsbomb.com/2020/07/statsbomb-release-expected-goals-with-shot-impact-height/>.
- Kwiatkowska, M., Norman, G., & Parker, D. (2011). PRISM 4.0: Verification of probabilistic real-time systems. In *Proceedings of the 23rd International Conference on Computer Aided Verification (CAV'11)*, Vol. 6806 of *LNCS*, pp. 585–591.
- Liu, G., Luo, Y., Schulte, O., & Kharrat, T. (2020). Deep soccer analytics: Learning an action-value function for evaluating soccer players. *Data Mining and Knowledge Discovery*, 34(5), 1531–1559.
- Liu, G., & Schulte, O. (2018). Deep reinforcement learning in ice hockey for context-aware player evaluation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pp. 3442–3448.
- Liu, S., Lever, G., Wang, Z., Merel, J., Eslami, S. M. A., Hennes, D., Czarnecki, W. M., Tassa, Y., Omidshafiei, S., Abdolmaleki, A., Siegel, N. Y., Hasenclever, L., Marris, L., Tunyasuvunakool, S., Song, H. F., Wulfmeier, M., Muller, P., Haarnoja, T., Tracey, B., Tuyls, K., Graepel, T., & Heess, N. (2022). From motor control to team play in simulated humanoid football. *Science Robotics*, 7(69), 1–16.
- Liu, T., & Hohmann, A. (2013). Applying the Markov chain theory to analyze the attacking actions between FC Barcelona and Manchester United in the European Champions League finale. *International Journal of Sports Science and Engineering*, 7(2), 79–86.
- Lucey, P., Bialkowski, A., Mofort, M., Carr, P., & Matthews, I. (2015). “Quality vs Quantity”: Improved shot prediction in soccer using strategic features from spatiotemporal data. In *Proceedings of the 9th MIT Sloan Sports Analytics Conference*, pp. 1–9.
- Luo, Y., Schulte, O., & Poupart, P. (2020). Inverse reinforcement learning for team sports: Valuing actions and players. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, pp. 3356–3363.
- Pfeiffer, M., Zhang, H., & Hohmann, A. (2010). A Markov chain model of elite table tennis competition. *International Journal of Sports Science & Coaching*, 5(2), 205–222.
- Power, P., Ruiz, H., Wei, X., & Lucey, P. (2017). Not all passes are created equal: Objectively measuring the risk and reward of passes in soccer from tracking data. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1605–1613.
- Riley, P., & Veloso, M. (2004). Advice generation from observed execution: Abstract Markov decision process learning. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI)*, pp. 631–636.

- Robberechts, P., & Davis, J. (2020). How data availability affects the ability to learn good xG models. In *International Workshop on Machine Learning and Data Mining for Sports Analytics*, pp. 17–27.
- Routley, K., & Schulte, O. (2015). A Markov game model for valuing player actions in ice hockey. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence*, pp. 782–791.
- Rudd, S. (2011). A framework for tactical analysis and individual offensive production assessment in soccer using Markov chains. In *New England Symposium on Statistics in Sports*.
- Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. (2016). Probabilistic programming in Python using PyMC3. *PeerJ Computer Science*, 2, e55.
- Sandholtz, N., & Bornn, L. (2018). Replaying the NBA. In *Proceedings of the 12th MIT Sloan Sports Analytics Conference*, pp. 1–13.
- Sandholtz, N., & Bornn, L. (2020). Markov decision processes with dynamic transition probabilities: An analysis of shooting strategies in basketball. *The Annals of Applied Statistics*, 14(3), 1122–1145.
- Schulte, O., Khademi, M., Gholami, S., Zhao, Z., Javan, M., & Desaulniers, P. (2017). A Markov game model for valuing actions, locations, and team performance in ice hockey. *Data Mining and Knowledge Discovery*, 31(6), 1735–1757.
- Singh, K. (2019). Introducing expected threat. Available at <https://karun.in/blog/expected-threat.html>.
- Spearman, W. (2018). Beyond expected goals. In *Proceedings of the 12th MIT Sloan Sports Analytics Conference*, pp. 1–17.
- Tuyts, K., Omidshafiei, S., Muller, P., Wang, Z., Connor, J., Hennes, D., Graham, I., Spearman, W., Waskett, T., Steel, D., Luc, P., Recasens, A., Galashov, A., Thornton, G., Elie, R., Sprechmann, P., Moreno, P., Cao, K., Garnelo, M., Dutta, P., Valko, M., Heess, N., Bridgland, A., Pérolat, J., De Vylder, B., Eslami, S. M. A., Rowland, M., Jaegle, A., Munos, R., Back, T., Ahamed, R., Bouton, S., Beauguerlange, N., Broshear, J., Graepel, T., & Hassabis, D. (2021). Game plan: What AI can do for football, and what football can do for AI. *Journal of Artificial Intelligence Research*, 71, 41–88.
- Van Roy, M., Robberechts, P., Decroos, T., & Davis, J. (2020). Valuing on-the-ball actions in soccer: A critical comparison of xT and VAEP. In *Proceedings of the 2020 AAAI Workshop on AI in Team Sports*, pp. 1–8.
- Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., & Davis, J. (2021). Leaving goals on the pitch: Evaluating decision making in soccer. In *Proceedings of the 15th MIT Sloan Sports Analytics Conference*, pp. 1–25.
- Wang, J., Fox, I., Skaza, J., Linck, N., Singh, S., & Wiens, J. (2018). The advantage of doubling: A deep reinforcement learning approach to studying the double team in the NBA. In *Proceedings of the 12th MIT Sloan Sports Analytics Conference*, pp. 1–12.
- Wenninger, S., & Lames, M. (2016). Performance analysis in table tennis - stochastic simulation by numerical derivation. *International Journal of Computer Science in Sport*, 15(1), 22–36.

Yam, D. (2019). Attacking contributions: Markov models for football. Available at <https://statsbomb.com/2019/02/attacking-contributions-markov-models-for-football/>.