KU LEUVEN

ARENBERG DOCTORAL SCHOOL
Faculty of Engineering Science

Soccer Through the Lens of Artificial Intelligence

Analyzing Actions and Tactics by Learning and Reasoning

Maaike Van Roy

Dissertation presented in partial fulfillment
of the requirements for the degree of
Doctor of Engineering Science (PhD):
Computer Science

Supervisor:
Prof. dr. Jesse Davis

January 2025

SOCCER THROUGH THE LENS OF ARTIFICIAL INTELLIGENCE

ANALYZING ACTIONS AND TACTICS BY LEARNING AND REASONING

Maaike VAN ROY

Supervisor:
Prof. dr. Jesse Davis

Members of the
Examination Committee:
Em. Prof. dr. ir. Jean Berlamont, chair
Prof. dr. Luc De Raedt
Prof. dr. Tom Holvoet
Prof. dr. Albrecht Zimmermann
(Université de Caen Normandie)
Dr. Sigrid Olthof
(Liverpool John Moores University)
Dr. ir. Jan Van Haaren
(Club Brugge)

Dissertation presented in partial fulfillment of the requirements for the degree of Doctor of Engineering Science (PhD): Computer Science

January 2025

© 2025 KU Leuven – Faculty of Engineering Science
Uitgegeven in eigen beheer, Maaike Van Roy, Boutersem (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher.

Acknowledgements

During the past five years, I had the pleasure of working on multiple projects that combine both my love for sports with my interest in artificial intelligence (AI). Looking back on this journey, I could not have done this without the support of many people, whom I wish to thank.

My supervisor, Jesse, thank you for giving me the opportunity to start a PhD under your supervision five years ago. I had but one condition for you, that my PhD project could entail "something sports-related". And oh boy it did! Enter: this ∼200-page dissertation that is all about sports and AI. Thank you for guiding and mentoring me along the way, for teaching me how to write about my work in a clear and structured way, for providing feedback (again and again and again) whenever I needed it, for teaching me how to present my work clearly to any audience (and because it's "good practice"), for teaching me how to "sell my work" to others (because this Belgian is way too shy at doing this), and highlighting the positive points whenever I was doubting the work (or myself).

My jury members, Luc, Tom, Albrecht, Sigrid, Jan, and Jean. Luc and Tom, thank you for the valuable feedback during these past five years. It really helped shape the work from the artificial intelligence side and greatly improved the final manuscript. Jan, thank you for all the fun discussions, feedback, collaborations, and opportunities over the past years (and of course also for saving the cake!). Albrecht and Sigrid, thank you for the great discussions and valuable feedback at Dagstuhl and the defense. Jean, thank you for chairing the defenses.

I wish to thank Statsbomb for providing the data sets used in this research, and the various projects and agencies that provided funding opportunities (Research Foundation-Flanders under EOS No. 30992574 and G0D8819N, KU Leuven Research Fund under iBOF/21/075, and the Flemish Government under the "Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen" programme).

Of course, this journey would not have been the same without the wonderful

people at the DTAI research group. The “old” 02.144 office, Vincent, Tom, Dries, and Jessa (yes, you also belong here even though I “stole” your desk). I especially loved our work-banter combination, which involved working hard to meet deadlines and random talks about literally any topic which provided some much needed breaks. The “new” 02.144 office, Luca and Loren. You brought a lot of fresh energy to the office. I love our dynamic of joking around, the fact that Luca always knows when there’s sandwiches, candy, or fruit in the cafeteria, and even the slow shift from Belgian lunch time to a more Italian lunch time (yes Luca, I see what you’re doing). Tim, for mentoring me during my master’s thesis and transferring your love for technology and sports. Jonas and Laurens, for being the best (coffee) buddies throughout all these years. Kilian, for always extending a “quick 5 minute coffee break” to “shoot, I’ve been gone for 30”.

I wish to thank my friends and family for their support. My parents, Ronny and Sonja, without whom I would not have been able to pursue these studies. The Drieskes, Peeterskes, and “aanhangsels”, for your continued support and interest in all that I am doing. The pink fluffy unicorns, Jonas, Charlie, PJ, Daan, Mister Michel, Thor, and Tarunika, for brightening up many days with not-always-pasta evenings and (sometimes very very very) competitive board game nights. My high school friends, Jozefien, Amber, Ky, Marjolijn, and Jef, for always being there and making it so easy to pick up conversations where they left off, even though we only see each other a few times a year.

Above all, Bram, for your unwavering love and support throughout this rollercoaster that is a PhD (and life). Thank you for all the love, the laughter, and the joy, and for being my partner in crime throughout all of it. Lof joe!

To our son Lasse, during this past year, you have amazed me with what you have already learned and I so look forward to exploring the world with you. Lassiebol, keep being your cheeky little self!

Popularized Abstract

Over the past few decades, the amount of data collected in professional soccer has skyrocketed, sparking the interest of clubs, media, researchers, and fans alike. This data offers insights into crucial aspects like player scouting, match preparation, and monitoring of physical performance. For example, it can help answer questions such as, "How much does this player contribute to creating scoring opportunities?" or "How would shooting more often from outside the penalty box affect a team's goal tally?"

Extracting meaningful insights from this data is far from simple. The data is complex and incomplete, due to the fast-paced nature of soccer, the many different events with varying characteristics that are spread out all over the pitch, and the possible errors during data collection. Added to that is the unpredictability of a game with human participants, which yields high levels of uncertainty and randomness. Many of the interesting questions also require counterfactual thinking — asking "what if" — which goes beyond traditional statistics like goals scored or passes completed.

We specifically focus on how artificial intelligence (AI) techniques can support player recruitment and match preparation. For these tasks, it is crucial to understand the in-game behavior and decision-making of players and teams. From an AI perspective, this can be viewed as agents that interact with an environment. A policy determines the decisions each agent makes, and each decision changes the environment, either immediately or further into the future. Understanding the policies and their implications boils down to two aspects: learning and reasoning. First, we need to learn these policies from match data. Then, we must reason about these policies, including counterfactual ones, to evaluate their implications and answer counterfactual questions.

Even though it is important to combine learning and reasoning, current approaches have focused mostly on learning-based techniques. These techniques can help with identifying the exact policies teams use and can give insights into

individual actions. The latter is commonly done by combining the learning-based techniques with the simplest form of reasoning, i.e., reasoning about how an action changes the environment in one time step. This combination of learning and basic reasoning is already useful, but there are two significant shortcomings that prevent us from providing insights into policies. First, the subsequent nature of the actions and their effects is not explicitly considered. Thus, these techniques cannot reason about which sequences of actions are better in certain situations. Second, because these techniques focus on the current trends in the data, it is also not possible to ask questions about counterfactual examples or decisions that were not seen during training.

This dissertation aims to fill this gap by developing a novel framework for learning and reasoning about policies of agents. Its core contribution is a Markov model-based framework for reasoning about the implications of learned and counterfactual policies of a team using probabilistic model checking techniques. Beyond this core contribution, this dissertation makes two other advances involving the combination of learning and reasoning. First, it provides a novel adaptation of an approach that can flag untrustworthy predictions or abstain from making predictions for data debugging. Second, it provides the first in-depth comparative study of two popular learning-based frameworks that can perform basic reasoning about the effect of a player's actions on the chances of scoring a goal.

Although the work is grounded in sports, its methodologies and findings contribute to broader AI research on combining learning with reasoning. The insights can be transferred to any field where it is crucial to understand and reason about policies of agents that interact with an environment.

Gepopulariseerde Samenvatting

Het afgelopen decennium is de hoeveelheid data die verzameld wordt over professioneel voetbal enorm toegenomen. Dit heeft ook de interesse van clubs, nieuwsdiensten, onderzoekers, en fans aangewakkerd om gebruik te maken van deze data om zo verschillende vragen te beantwoorden in verband met het aanwerven van spelers, voorbereiden van wedstrijden, en het opvolgen van de fysieke prestaties van de spelers. Bijvoorbeeld, "Hoeveel draagt een speler bij tot de doelkansen van zijn team?" en "Wat zou de impact op het aantal doelpunten zijn als men meer van buiten het strafschopgebied zou schieten?".

Inzichten uit de verzamelde data verkrijgen is niet makkelijk. De data is zeer complex, bevat vele componenten, en is daarnaast ook heel schaars omdat voetbal een dynamische sport is waar maar weinig acties op exact dezelfde manier en plaats worden uitgevoerd. Daarbovenop wordt de data verzameld door mensen en sensoren waardoor ze ook veel ruis en fouten kan bevatten. Door het menselijke aspect van de sport is het ook moeilijk te voorspellen wat er zal gebeuren en de data zal dus ook veel onzekerheid en willekeurigheid bevatten. Tot slot zijn ruwe statistieken zoals het aantal doelpunten niet voldoende om inzichten te verwerven. Veel interessante vragen bevatten namelijk iets tegenfeitelijk, "wat als", waarvoor een dieper inzicht in de data nodig is.

In dit proefschrift focussen we op hoe technieken uit het veld van artificiële intelligentie (AI) kunnen helpen bij het aanwerven van spelers en voorbereiden van wedstrijden. Hiervoor moeten we begrijpen hoe spelers en teams zich tijdens een wedstrijd gedragen en welke beslissingen ze maken. Vanuit een AI-perspectief kan dit gezien worden als agenten die interageren met hun omgeving. Het beleid van de agenten bepaalt welke beslissingen ze maken en elke beslissing heeft een impact op hun omgeving. Om inzichten te verwerven over hun beleid en de impact ervan zijn er twee belangrijke aspecten nodig: leren en redeneren.

We moeten hun beleid en de effecten ervan kunnen “leren” uit de beschikbare data, en daarnaast moeten we kunnen “redeneren” over het geleerde, of mogelijks tegenfeitelijke, beleid om het van context te voorzien en de gevolgen ervan in te schatten.

Hoewel het belangrijk is om zowel leren als redeneren te combineren, focussen huidige technieken vooral op leren. Met deze technieken kunnen we het beleid van agenten leren uit de beschikbare data en kunnen we inzichten verwerven over individuele acties door technieken om te leren te combineren met de meest eenvoudige vorm van redeneren. Namelijk, redeneren over hoe een individuele actie de omgeving verandert in één enkele tijdstap. Deze combinatie van leren en de simpele vorm van redeneren is al zeer nuttig, maar er zijn twee belangrijke tekortkomingen die het moeilijk maken om inzichten te verwerven over het beleid van agenten. Ten eerste wordt er niet uitdrukkelijk rekening gehouden met het sequentiële aspect van acties en hun effecten. Het laat ons dus niet toe om te redeneren over de voordelen en de nadelen van verschillende opeenvolgingen van acties. Ten tweede, omdat deze methoden focussen op de huidige trend in de data, is het niet mogelijk om vragen te stellen over tegenfeitelijke voorbeelden die nog nooit gezien zijn.

Dit proefschrift heeft als doel deze gebreken op te lossen door een nieuw raamwerk te ontwikkelen voor het leren van en redeneren over het beleid van agenten. De belangrijkste bijdrage is een raamwerk om te redeneren over de gevolgen van zowel een geleerd als een tegenfeitelijk beleid van een team door gebruik te maken van Markov modellen en probabilistische modelverificatie technieken. Hiernaast bevat dit proefschrift nog twee andere bijdragen in verband met de combinatie van leren en redeneren. Ten eerste stelt het een manier voor om fouten in de data te ontdekken door een nieuwe toepassing van methoden die, in het geval van ongeloofwaardige voorspellingen, de gebruiker kunnen waarschuwen of ervoor kunnen zorgen dat het model geen voorspelling maakt. Ten tweede stelt het de eerste diepgaande vergelijkende studie voor van twee populaire raamwerken die technieken om te leren toepassen om op de meest eenvoudige manier te redeneren over het effect van een speler zijn acties op de kans op scoren.

Hoewel het werk gebaseerd is op sport, dragen de methoden en inzichten bij tot het bredere AI-onderzoek rond het combineren van leren en redeneren. Het werk bevat inzichten die toepasbaar zijn in eender welk veld waar het noodzakelijk is om te redeneren over het beleid van agenten die interageren met een omgeving.

Abstract

There has been a tremendous increase in the amount of data collected from professional soccer over the past decades. This has sparked the interest of clubs, news media, researchers, and fans to make use of this data to answer many of the questions they have related to player scouting, match preparation, and monitoring of physical performance. For example, it enables answering questions such as "How much does the prospective player contribute towards goal-scoring opportunities?" and "How would shooting more often from outside the penalty box affect a team's number of goals scored?".

However, deriving meaningful insights from soccer data presents several challenges. First, the data is high-dimensional and contains many heterogeneous components that must be considered in order to capture the full context. Second, the data is inherently sparse due to the dynamic nature of soccer. Third, as the data is observational, collected by humans and sensors, it is prone to noise and may be incomplete. Fourth, soccer's inherent dynamism and human involvement introduce uncertainty and randomness. Fifth, many of the interesting questions that stakeholders wish to address involve counterfactuals, i.e., answering "what if" questions. Addressing counterfactual questions requires a deeper understanding of the data than what raw statistics can provide.

Concretely, we focus on how artificial intelligence (AI) techniques can support player recruitment and match preparation. Therefore, understanding the in-game behavior and decision-making of players and teams is crucial. From an AI perspective, this can be viewed as agents that interact with an environment. A policy determines the decisions each agent makes, and each decision alters the environment, either in the short or long term. Deriving insight into policies and their implications requires both the ability to "learn" the policies and effects from the available data, and "reason" about these learned, or possibly counterfactual, policies to contextualize them and assess their implications.

Despite the importance of combining learning and reasoning, current approaches

focus mostly on learning-based techniques. These techniques are used to (1) learn the current policies from data, and (2) derive insights into individual actions by combining the learning-based techniques with the simplest form of reasoning, i.e., reasoning about how an action changes the environment in one time step. While the combination of learning and this simple form of reasoning is already useful, there are two significant shortcomings that prevent us from providing insights into policies. First, the subsequent nature of the actions and their effects is not explicitly taken into account. Hence, these techniques cannot reason about the merits and pitfalls of various sequences of actions. Second, because these techniques focus on the current trends in the data, it is also not possible to ask questions about counterfactual examples or decisions that were not seen during training.

This dissertation aims to fill this gap by developing a novel framework for learning and reasoning about policies of agents. Its core contribution is a Markov model-based framework for reasoning about the implications of learned and counterfactual policies of a team using probabilistic model checking techniques. Beyond this core contribution, this dissertation makes two other advances involving the combination of learning and reasoning. First, it provides a novel adaptation of learning with abstention for data debugging. Second, it provides the first in-depth comparative study of two cornerstone learning-based frameworks that can perform basic reasoning about the effect of a player's actions on the likelihood of scoring a goal.

Although the work is grounded in sports, its methodologies and findings contribute to broader AI research on combining learning with reasoning, and offers transferable insights for any field where it is crucial to understand and reason about policies of agents that interact with an environment.

Beknopte Samenvatting

De hoeveelheid data over professioneel voetbal is het afgelopen decennium enorm toegenomen, wat de interesse van clubs, media, onderzoekers en fans heeft aangewakkerd. Deze data wordt gebruikt om vragen te beantwoorden over het aanwerven van spelers, voorbereiden van wedstrijden, en het opvolgen van fysieke prestaties. Bijvoorbeeld, “Hoeveel draagt een speler bij tot de doelkansen van zijn team?” en “Wat zou de impact op het aantal doelpunten zijn indien men meer van buiten het strafschopgebied zou schieten?”.

Het verkrijgen van inzichten uit deze data is niet eenvoudig en er zijn een aantal uitdagingen aan verbonden. De data is hoog-dimensioneel en bevat verschillende componenten (bijv., temporele en spatiale componenten). Het dynamisch karakter van voetbal zorgt ervoor dat er weinig acties op exact dezelfde manier en plaats worden uitgevoerd. Hierdoor is de data ook heel schaars. Daarbovenop kan de data ook veel ruis en fouten bevatten, aangezien het om observationele data gaat die verzameld wordt door mensen en sensoren. Omdat voetbal een sport is met menselijke deelnemers, is het ook moeilijk te voorspellen wat er zal gebeuren en zal er in de data ook veel onzekerheid en willekeurigheid zitten. Tot slot is het berekenen van ruwe statistieken uit de data niet voldoende want veel interessante vragen bevatten iets tegenfeitelijk, bijv., “wat als”, waarvoor een dieper inzicht in de data nodig is.

Meer specifiek focussen we op hoe technieken uit het veld van artificiële intelligentie (AI) kunnen helpen bij het aanwerven van spelers en voorbereiden van wedstrijden. Hiervoor is het noodzakelijk om te begrijpen hoe spelers en teams zich tijdens een wedstrijd gedragen en welke beslissingen ze maken. Vanuit een AI-perspectief kan dit gezien worden als agenten die interageren met hun omgeving. Het beleid van de agenten bepaalt welke beslissingen ze maken en elke beslissing heeft een impact op hun omgeving. Om inzichten te verwerven over hun beleid en de impact ervan is het noodzakelijk om zowel hun beleid en de effecten ervan te “leren” uit de beschikbare data, als te “redeneren” over het geleerde, of mogelijks tegenfeitelijke, beleid om het van context te voorzien en

de implicaties ervan in te schatten.

Hoewel het belangrijk is om zowel leren als redeneren te combineren, focussen huidige technieken vooral op leren. Deze technieken laten ons toe om (1) het beleid van agenten te leren, en (2) inzichten te verwerven over individuele acties door technieken om te leren te combineren met de meest eenvoudige vorm van redeneren. Namelijk, redeneren over hoe een individuele actie de omgeving verandert in één enkele tijdstap. Deze combinatie van leren en de simpele vorm van redeneren is al zeer nuttig, maar er zijn twee belangrijke tekortkomingen die ons verhinderen om inzichten te verwerven over het beleid van agenten. Ten eerste wordt er niet expliciet rekening gehouden met het sequentiële aspect van acties en hun effecten. Hierdoor is het niet mogelijk om te redeneren over de voordelen en de nadelen van verschillende sequenties van acties. Ten tweede, omdat deze methoden focussen op de huidige trend in de data, is het niet mogelijk om vragen te stellen over tegenfeitelijke voorbeelden die nog nooit gezien zijn.

Dit proefschrift heeft als doel deze gebreken op te lossen door een nieuw raamwerk te ontwikkelen voor het leren van en redeneren over het beleid van agenten. De belangrijkste bijdrage is een raamwerk gebaseerd op Markov modellen om te redeneren over de implicaties van zowel een geleerd als een tegenfeitelijk beleid van een team door gebruik te maken van probabilistische modelverificatie technieken. Hiernaast bevat dit proefschrift nog twee andere bijdragen in verband met de combinatie van leren en redeneren. Ten eerste stelt het een nieuwe manier voor om methoden voor het leren met onthouding te gebruiken voor het ontdekken van fouten in de data. Ten tweede stelt het de eerste diepgaande vergelijkende studie voor van twee populaire raamwerken die technieken om te leren toepassen om op de meest eenvoudige manier te redeneren over het effect van een speler zijn acties op de kans op scoren.

Hoewel het werk gebaseerd is op sport, dragen de methoden en inzichten bij tot het bredere AI-onderzoek rond het combineren van leren en redeneren. Het werk biedt overdraagbare inzichten die toepasbaar zijn in eender welk veld waar het noodzakelijk is om te redeneren over het beleid van agenten die interageren met een omgeving.

List of Abbreviations

AI Artificial Intelligence.

AUROC Area Under the Receiver Operating Characteristic Curve.

DTMC Discrete-Time Markov Chain.

EPV Expected Possession Value.

IFAB International Football Association Board.

LTL Linear Temporal Logic.

MDP Markov Decision Process.

MRP Markov Reward Process.

OC Output Configuration.

PCTL Probabilistic Computational Tree Logic.

SPADL Soccer Player Action Description Language.

t-SNE t-Distributed Stochastic Neighbor Embedding.

UMAP Uniform Manifold Approximation and Projection.

VAEP Valuing Actions by Estimating Probabilities.

xG Expected Goals.

xT Expected Threat.

Contents

Popularized Abstract iii

Gepopulariseerde Samenvatting v

Abstract vii

Beknopte Samenvatting ix

List of Abbreviations xi

Contents xiii

List of Figures xvii

List of Tables xxiii

1 Introduction 1
- 1.1 Dissertation Statement . . . 5
- 1.2 Contributions and Structure of the Dissertation . . . 7
  - 1.2.1 Background . . . 8
  - 1.2.2 Contribution 1: A Critical Comparison of Action-Value Frameworks . . . 8
  - 1.2.3 Contribution 2: Contextualizing Action-Values by Reasoning . . . 9
  - 1.2.4 Contribution 3: A Framework for Learning and Reasoning About Policies . . . 10
  - 1.2.5 Conclusion . . . 11
- 1.3 Other Research Contributions . . . 11

2 Background 19
- 2.1 Soccer . . . 20

2.2 Data Sources . . . 21
2.2.1 Physical Data . . . 22
2.2.2 Observational Match Data . . . 22
2.3 Artificial Intelligence Techniques . . . 26
2.3.1 Probabilistic Classification . . . 26
2.3.2 Markov Models . . . 26
2.3.3 Bayesian Modeling and Probabilistic Programming . . . 28
2.3.4 Probabilistic Model Checking . . . 29
2.4 Related Work . . . 30
2.4.1 Quantifying Player Performance . . . 30
2.4.2 Discovering and Analyzing Tactics . . . 34

3 Valuing On-the-Ball Actions: A Critical Comparison of xT and VAEP 37
3.1 Expected Threat . . . 39
3.2 VAEP . . . 40
3.3 Conceptual Comparison of xT and VAEP . . . 41
3.3.1 Location-Based vs Feature-Based . . . 42
3.3.2 Possession-Based vs Window-Based . . . 43
3.4 Experimental Comparison of xT and VAEP . . . 44
3.4.1 Data . . . 44
3.4.2 Experimental Comparison of Action-Values . . . 44
3.4.3 Experimental Comparison of Player Ratings . . . 47
3.5 Conclusion . . . 52

4 DataDebugging: Enhancing Trust in Action-Value Frameworks by Contextualization 55
4.1 Detecting Abnormal Examples in Action-Value Frameworks . . . 56
4.1.1 Action-Value Framework . . . 56
4.1.2 Detecting Abnormal Examples in Tree Ensembles . . . 56
4.2 Experiments . . . 57
4.2.1 Data . . . 57
4.2.2 Detecting Abnormal Examples . . . 57
4.2.3 Analysis Using `TRACE` . . . 59
4.3 Conclusion . . . 65

5 A Markov Framework for Learning and Reasoning About Tactics 67
5.1 Modeling Offensive Behavior as a Markov Model . . . 69
5.1.1 State Space . . . 69
5.1.2 Action Space . . . 69
5.1.3 Transition Function . . . 70
5.1.4 Policy . . . 70
5.1.5 Reward Function, Discount Factor and Value Function . . . 71
5.2 Learning the Markov Model from Event Stream Data . . . 71

5.2.1 Predicting the Intended End Locations of Actions . . . 72
5.2.2 A Bayesian Approach to Learning the Transition Model and Policy . . . 74
5.3 Reasoning About Learned Policies . . . 75
5.3.1 Reasoning About the Effect of Offensive Tactics . . . 76
5.3.2 Reasoning About the Effect of Defensive Tactics . . . 77
5.4 Reasoning About Alternative Policies . . . 80
5.4.1 Counterfactual Movement Actions . . . 81
5.4.2 Counterfactual Shot Actions . . . 83
5.5 Experimental Setup and Evaluation . . . 86
5.5.1 Data Sets . . . 87
5.5.2 Transient State Spaces . . . 87
5.5.3 Learning the MDP . . . 88
5.5.4 Evaluation of the Intended End Location Predictions . . . 89
5.5.5 Evaluation of the Bayesian Approach . . . 90
5.5.6 Evaluation of the Complete Model . . . 91
5.6 Use Case: Reasoning About Shot Policies . . . 92
5.6.1 Shoot Immediately or Move Before Shooting? . . . 93
5.6.2 Probability of Generating a Better Shot Later in the Possession . . . 96
5.6.3 Uniformly Shooting More or Less From Distance . . . 97
5.6.4 Targeted Increases of Shots From Distance . . . 99
5.7 Use Case: Reasoning About Defensive Tactics . . . 100
5.7.1 Shot Suppression . . . 101
5.7.2 Movement Suppression . . . 102
5.7.3 Effect of Defensive Tactics Once Team Adapts . . . 104
5.8 Related Work . . . 105
5.9 Conclusion . . . 107

6 A Markov Framework for Analyzing Build-Up Effectiveness and Defensive Setups 111
6.1 Shortcomings of Raw Data Querying Techniques . . . 112
6.2 A Markov Model of Build-Ups . . . 113
6.2.1 Defining Build-Ups and Defensive Setups . . . 113
6.2.2 Defining the Markov Model . . . 115
6.2.3 Learning the Markov Model . . . 118
6.3 Reasoning About Build-Up Effectiveness . . . 119
6.4 Experiments . . . 121
6.4.1 Data . . . 122
6.4.2 Experimental Settings . . . 123
6.4.3 Identifying Offensive Build-Up Patterns . . . 124
6.4.4 Optimizing Defensive Tactics . . . 125
6.4.5 Evaluating Stability . . . 129

6.4.6 Evaluating Design Choices . . . 129
6.5 Related Work . . . 131
6.6 Conclusion . . . 132

7 Conclusion **133**
7.1 Research Gaps and Data Challenges . . . 133
7.2 Contributions . . . 135
7.2.1 Contribution 1: A Critical Comparison of Action-Value Frameworks . . . 135
7.2.2 Contribution 2: Contextualizing Action-Values by Reasoning . . . 136
7.2.3 Contribution 3: A Framework for Learning and Reasoning about Policies . . . 137
7.3 Future Work . . . 137

A A Markov Framework for Learning and Reasoning About Tactics **141**
A.1 Experimental Setup . . . 141
A.1.1 Gradient Boosted Trees Ensembles . . . 141
A.1.2 Hierarchical Bayesian Models . . . 142
A.2 Policy Modifications . . . 143
A.2.1 Decreasing the Movement Probability . . . 143
A.2.2 Estimating the Expected Number of Goals . . . 144
A.3 Ablation Study . . . 144

B A Markov Framework for Analyzing Build-Up Effectiveness and Defensive Setups **147**
B.1 Where Do Build-Ups End? . . . 147
B.2 Compute Infrastructure and Timing . . . 147
B.3 Frequent Patterns Using a Sequential Pattern Miner . . . 148
B.4 Stability Experiment . . . 149

Bibliography **155**

Curriculum Vitae **173**

Publications **175**

List of Figures

2.1 Illustration of a soccer pitch. 21
2.2 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. 24
2.3 Illustration of one action from the Belgium versus Portugal game on June 27, 2018, as recorded in the contextual event stream data format. 24
2.4 Illustration of (a) a sequence of tracking data, and (b) a sequence of broadcast tracking data where the grey area denotes the field of view of the camera. 25
2.5 Generalized illustration of the approach taken for valuing game states. The various methods mainly differ in how they assign values to the game states, which context is taken into account to describe the game states, and which machine learning model is used to learn the relationship between the game states and their associated values. 33

3.1 A heatmap of the xT values obtained after training the xT model on the training data set until convergence. Darker colors denote higher xT values. 40
3.2 Histograms of the VAEP and xT values for a set of actions that are rated differently by both frameworks. 45
3.3 The Jaccard similarity coefficient between the top-k rankings created by the xT and VAEP frameworks. The plot shows an increase in similarity for the top-25 players of both models and then a drop when adding the next 25 players before steadily increasing. 50

3.4 Scatter plots of the player ratings produced on both samples. The linear least squares regression line is shown in red. The plot for xT displays a high correlation between both ratings. For VAEP, the plots are more random than the one for the xT framework, but a correlation becomes visible when only looking at the offensive values for passes, dribbles and crosses of VAEP. 52

4.1 Two identified annotation errors in the FAWSL: (a) the start and end location of the two dribbles are incorrectly annotated, and (b) the y-coordinate of the first action should be near the penalty box border instead of near the sideline. 58
4.2 Four representative examples of rare sequences in the EURO 2020 (a, b) and FAWSL (c, d). Figures (a) and (c) are two rare sequences with long successful passes and dribbles, either along the sideline or through the middle of the defense. Figures (b) and (d) are two rare successful shots, one from the halfway line and one from a tight angle outside the penalty box. 59
4.3 Four representative examples of rather normal sequences in the FAWSL that are falsely identified as abnormal. 60
4.4 Visualization of a 2-dimensional UMAP and t-SNE embedding of the output configurations of "no goal" predicted test examples in the EURO 2020. Red means that the example is more abnormal according to its OC-Score while blue means it is more normal. . . 61
4.5 Visualization of a 2-dimensional UMAP and t-SNE embedding of the output configurations of "no goal" predicted test examples in the FAWSL. Red means that the example is more abnormal according to its OC-Score while blue means it is more normal. . 62
4.6 The selected abnormal cluster in the top left of the UMAP embedding can be tracked and mapped with `TRACE` to the bottom cluster in the t-SNE embedding, indicating that there is an underlying structure present in the output configurations of "no goal" predicted test examples in the EURO 2020. The shift in the distribution of OC-Scores of the selected examples also indicates that this cluster is indeed more abnormal. 63
4.7 UMAP embeddings of "no goal" predicted test examples in the FAWSL and EURO 2020 with a selected number of (a) normal examples and (b) abnormal examples and their high-dimensional neighbors. 64

5.1 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. . . 73
5.2 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. 77
5.3 Possible expected goals (xG) distribution in one state. Orange indicates the average xG value (μ_s), green indicates the average of the below-average xG values (μ_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 (μ_s^{high}). 86
5.4 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 3m × 3m. This state space is more suited for analyzing shot policies 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 policies. 88
5.5 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%. 93
5.6 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. 94
5.7 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. 95

5.8 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. 97
5.9 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 5.6 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. 98
5.10 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 5.6 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. 99
5.11 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. 100
5.12 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. 102
5.13 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%. 103

6.1 Illustration of the defensive setups with dots depicting the defending team's players. 114
6.2 The location-based states (left) and extended states (left × right) with patches created relative to the ball carrier. 115
6.3 Illustration of the final third of the pitch (gray), and the left (red) and right (blue) flank during build-up. 120

6.4 Top 10 most likely build-up sequences against a high/low block. The most likely sequence for each team is in red. 124
6.5 Percentage increase in build-up effectiveness when starting build-ups from the left/right side, both against a high/low block. Bars show a team's most efficient side. 126
6.6 Percentage decrease in build-up effectiveness when forced to play the ball towards a certain side. 127
6.7 Percentage decrease in build-up effectiveness when all players are blocked on the center back's left/right side, both against a high/low block. 128
6.8 Illustrates the definition of "blocking all players on a certain side of the center back". The cross-hatched patch indicates the patch in which no reachable teammate is allowed. 129
6.9 Stability of the build-up effectiveness against a left- and right-forcing block, compared split-wise over 10 runs. 130

B.1 For each state, the percentage of build-up sequences that ends in it. 148
B.2 The top 10 most frequent patterns involving three actions for (a) Barcelona and (b) RB Leipzig according to PrefixSpan on the build-up sequences against a high block for the respective teams. Barcelona and RB Leipzig are playing left to right. 149
B.3 Stability of the build-up effectiveness against a left- and right-forcing block for all Bundesliga teams, compared split-wise over 10 runs. 149
B.4 Stability of the build-up effectiveness against a high and low block when starting the build-up from the left and right side for all LaLiga teams, compared split-wise over 10 runs. 151
B.5 Stability of the build-up effectiveness against a high and low block when starting the build-up from the left and right side for all Bundesliga teams, compared split-wise over 10 runs. 152
B.6 Stability of the build-up effectiveness against a high and low block when blocking the players on the center back's left and right side for all considered LaLiga teams, compared split-wise over 10 runs. 153
B.7 Stability of the build-up effectiveness against a high and low block when blocking the players on the center back's left and right side for all considered Bundesliga teams, compared split-wise over 10 runs. 154

List of Tables

3.1 The top-25 players who played at least 900 minutes in the 2018/2019 English Premier League season with ratings according to the xT framework. Goals per 90 minutes (G/90) and assists per 90 minutes (A/90) are also shown for each player. 48
3.2 The top-25 players who played at least 900 minutes in the 2018/2019 English Premier League season with ratings according to the VAEP framework. Goals per 90 minutes (G/90) and assists per 90 minutes (A/90) are also shown for each player. 49

4.1 Random triplet accuracy for both dimension reduction techniques and all three random seeds on both data sets. 62

5.1 Average number of actions that are passes, dribbles, crosses, or shots per team in each season for each of the considered competitions. 87
5.2 AUROC and Brier scores for the end location prediction models averaged (± 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 (see Section 5.2.1). For AUROC, higher values are better. For the Brier score, lower values are better. 90

5.3 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 (± 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. 91
5.4 Percent decrease in each team's probability of scoring when it is forced to avoid its respective region in the top row of Figure 5.13. The results are shown both before and after adapting to the defensive tactics, and for both approaches of adapting the policy (PR = proportional approach, SA = same aggressiveness approach). 104

6.1 The average number of actions (± 1 std) per team during the extracted build-up situations in each data set. 123
6.2 Original Brier scores (± 1 std) for the intended state prediction with the original (O) and teammate information (TM) approach, averaged over all teams. The baseline predicts the distribution of observed end locations in the train set. 130

A.1 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). 142
A.2 AUROC and Brier scores for the ablation study on the end location prediction models. Results are averaged (± 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). 145

Chapter 1

Introduction

The amount of data collected from professional soccer has increased tremendously over the past decades. This data varies from simple statistics (e.g., the number of shots taken by a player, or passes made by a team), to more complex data that captures the actions that players perform, the context of these actions (e.g., positions of the other players), as well as the physical state of the players (e.g., their heart rate over time). Clubs, news media, researchers, and fans have increasingly been interested in analyzing this data to gain insights from it. On a high level, this data is typically used to aid human decision-making in three domains.

Player Recruitment. Squads are optimized multiple times per year, with players being sold to and recruited from various other teams. Before recruiting new players, one needs to perform an extensive analysis of the players' skills and capabilities. That is, it is necessary to assess the players on the technical, tactical, physical, and mental level, and to ensure that he/she will fit within the team. Data is increasingly used to support this process [10, 17, 36, 92].

Typical questions the recruitment staff may be confronted with during this process are:

1. How much does the prospective player contribute towards goal scoring opportunities, and towards defending against opponent goals?
2. How creative is the prospective player?
3. What is the prospective player's playing profile (i.e., where does he/she tend to perform his/her actions, where/who are they aimed at, and how well are they executed)?

Match Preparation. Preparing for the upcoming match typically involves scouting the opponent and analysing their tactics, both as a team and for the individual players. Typically, video analysts spend hours watching and annotating video footage to understand the used tactics. As more data has become available, there is increasing interest in automated data-based approaches to aid analysts with this process [25, 40, 54, 156, 161].

Typical questions a coach and analyst may be confronted with are:

1. In which areas of the pitch is an attacking team most effective at reaching 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?
4. Which defensive tactics are more effective to prevent the opponent from building up?

Monitoring of Physical Performance. Ensuring that players stay fit and free of injuries is of paramount importance. Given that players take part in multiple training sessions and matches each week, which put a lot of stress on the players' bodies, it is necessary to continuously monitor their fitness and injury risk [123, 153].

Typical questions the coach and medical staff may be confronted with are:

1. How should the training schedule be adapted to further optimize a player's fitness?
2. Should a certain player be left out of the selection to minimize their risk of injury?

In this dissertation, we focus on how artificial intelligence (AI) techniques can support player recruitment and match preparation. From an AI perspective, data about soccer matches poses five main challenges.

Challenge 1: Spatial and Temporal Information. The data contains multiple heterogeneous components that need to be combined in order to capture the full context. For example, it contains both spatial (i.e., continuous location

on the pitch) and temporal (i.e., order or timestamp, time series) information. Automated techniques must be able to handle many heterogeneous components (e.g., both spatial and temporal).

Challenge 2: Data Sparsity. The data in itself is inherently sparse. Soccer is highly dynamic and teams rarely perform the exact same sequence of actions in the same order at the same locations. Including many of the spatial and temporal components in an analysis further compounds the sparsity. Furthermore, the most important reward in soccer (i.e., goals) amounts to $< 1\%$ of all events during a match.[1] Combining data of other players or previous seasons to combat this sparsity is not always beneficial, as the characteristics of the players may be different and changes in players and management happen regularly, amounting to different behavior.

Challenge 3: Observational Data. The data is observational, meaning that it is collected from what can be observed from a match, either by humans or sensors. This means that the data can be incomplete and noisy. For example, when actions fail, their intended end location is unknown (e.g., if a player attempts a cross that the opposition clears, we are unsure of where the player was aiming for). Additionally, sensors can fail, measurement errors can occur, and data providers that collect the data do not always agree on terminology (e.g., certain actions might be defined differently) which contributes to the already noisy nature of the data.

Challenge 4: Uncertainty. As soccer is a highly dynamic and real-world sport that involves human participants, it also contains a lot of uncertainty and randomness. For example, the specific type of action players will choose at a certain moment will not always be the same and the outcomes of these actions is also subject to some randomness (e.g., passes may be deflected). Thus, automated methods must be able to capture and handle this uncertainty.

Challenge 5: Counterfactuals. Many of the above questions involve counterfactuals. That is, they require understanding what would happen if certain aspects were performed differently (e.g., actions are performed under different circumstances, a team's behavior or style changes over time or in response to certain in-game factors) [112]. Therefore, these questions require a deeper understanding of the underlying dynamics and interactions, and an ability to reason about the consequences of decisions, actions, and tactics.

Supporting player recruitment and match preparation requires understanding the in-game behavior and decision-making of players and teams. From an AI perspective, how the players and teams behave, make decisions, and interact

[1] A top flight soccer team will perform between 1500 and 3000 on-the-ball actions during a game. On average, less than three goals are scored in each game of a top league.

with each other and the ball can be viewed as agents that interact with an environment. A policy describes how each agent behaves, i.e., which decisions they make in each situation in the environment, and each decision alters the environment. These effects are both immediate (e.g., whether a pass succeeds or fails) and longer term (e.g., whether the team is likely to create danger later on). Deriving insights into the policies requires both the ability to "learn" the current policies and their effects on the environment from the available data, and "reason" about these learned, or possibly counterfactual, policies to contextualize them and assess their implications.

Despite the importance of combining learning and reasoning, current approaches focus mostly on learning-based techniques. Typically, these techniques are used to (1) learn the current policies that teams use from data (e.g., identifying recurrent patterns [40, 156, 161]), and (2) derive insights into individual actions by combining the learning-based techniques with the simplest form of reasoning, i.e., reasoning about how an action changes the environment in one time step. In particular, this form of reasoning is often used to reason about the impact of a single action on the likelihood of scoring a goal in the future [36, 59, 133, 136].

While the combination of learning and this simple form of reasoning is already useful, current learning and reasoning techniques have two significant shortcomings that prevent us from providing insights into the various policies used by agents. First, this simple form of reasoning does not explicitly take the subsequent actions and their effects into account, which is required to reason about the sequential behavior of actions. Thus, it does not allow for reasoning about the different merits and pitfalls of various sequences of actions, and hence cannot give insights into the learned policies. Second, learning-based techniques only focus on the current trends in the data. Therefore, it is not possible to ask questions and reason about counterfactual examples or entirely different decisions, i.e., what was not seen during training. Yet, counterfactuals are at the heart of many of the stakeholders' questions.

This dissertation aims to fill this gap by developing a novel framework for learning and reasoning about policies of agents. Specifically, its core contribution is a Markov model-based framework for reasoning about the implications of learned and counterfactual policies of a team using probabilistic model checking techniques. In general, we propose a novel way of using probabilistic model checking techniques to reason about policies of agents in learned Markov models. The framework allows for reasoning about the learned policies and their implications, both in the short and long term, and can be used for trading off the merits and pitfalls of various sequences of actions. Additionally, the framework makes it possible to reason in a counterfactual way about policies to understand the implications of alternative decisions. For example, we show that while learning-based methods indicate that shots closer to goal are more

beneficial, which has influenced the current trend in soccer to shoot less from long distance, teams could score more goals by going against the conventional wisdom and shooting more frequently from distance.

Beyond this core contribution, this dissertation makes two other advances involving the combination of learning and reasoning. First, we provide a novel adaptation of learning with abstention for data debugging. This allows us to reason about the predictions of a learning-based model, and hence flag unusual examples, such as annotation errors, for which the predictions cannot be trusted. Second, we provide the first in-depth comparative study of two cornerstone learning-based frameworks that can perform basic reasoning about the effect of a player's actions on the likelihood of scoring a goal. While there exist numerous such frameworks, they are never directly compared, either conceptually or empirically, and insights into their advantages and disadvantages are lacking.

Although this work is grounded in sports, its methodologies and findings contribute to broader AI research on combining learning with reasoning, offering transferable insights for any field where agents interact with dynamic environments, and where understanding and reasoning about policies is crucial.

1.1 Dissertation Statement

In this dissertation, we focus on the novel research direction of combining learning and reasoning techniques from the field of artificial intelligence to derive insights into learned and counterfactual policies. We illustrate the novel frameworks in the setting of soccer.

The primary focus is on frameworks that can reason about learned and counterfactual policies of agents in dynamic environments. In particular, we focus on reasoning about a soccer team's and their opponent's in-game behavior. Therefore, we investigate (1) how one can capture a team's offensive behavior and the opponent's defensive structure in a Markov model that is learned from data, and (2) a novel way of reasoning about various tactical aspects, both offensive and defensive, in the observed and counterfactual policies in the Markov model using probabilistic model checking techniques.

From an AI perspective, we aim to answer the following two research questions:

RQ1 Can probabilistic model checking, a technique used for verifying the correctness of a model, be used to reason about the learned policies in a Markov model?

RQ2 To what extent does this approach allow for counterfactual reasoning about the policies in a Markov model?

From a soccer perspective, we aim to answer the following six research questions:

RQ3 How should one trade off the sparsity of the data with the context to be included in the Markov model?

RQ4 How should one ensure that the obtained results remain interpretable to the end user?

RQ5 Can reasoning about the offensive actions in the model give us insight into which actions are preferred in certain situations?

RQ6 Can reasoning about the learned offensive behavior of a team give us insight into which areas an opponent should target defensively?

RQ7 Can counterfactual reasoning about how often certain offensive actions are performed in the model give us insight into the long-term effect on a team's outcome (e.g., goals)?

RQ8 Can reasoning about the opponent's defensive structure give us insight into the effect on a team's offensive effectiveness?

The secondary focus is on a framework that enables flagging unusual examples for which the predictions of learning-based approaches might not be trustworthy. Learned models struggle to extrapolate beyond what they have seen in the data. Sports continuously evolve and the examples the models were trained on might deviate from the examples for which predictions need to be made. In such cases, it would be safer to abstain from making a prediction or flag the examples that are different to the end user. In soccer, learning-based approaches are often combined with basic reasoning to estimate the impact of actions on the likelihood of scoring a goal. These approaches are often called "action-value frameworks", as they attempt to estimate how valuable in-game actions are. We investigate how a reasoning approach for learning with abstention can be used to debug the data provided to an action-value framework, and hence flag deviating examples for which the predictions should be treated with caution.

From an AI perspective, we aim to answer the following research question:

RQ9 Does an approach to reason about the predictions of a learned model allow us to identify deviating examples for which the model is at risk of mispredicting?

From a soccer perspective, we aim to answer the following two research questions:

RQ10 Given that players occasionally perform unusual actions, which might not have a well-calibrated prediction, is it possible to detect these by reasoning about the predictions of an action-value framework?

RQ11 Given that the data is noisy and might contain errors due to manual annotation, is it possible to detect these by reasoning about the predictions of an action-value framework?

The tertiary focus is on a comparative study of two popular action-value frameworks. While there are numerous action-value frameworks available, they are rarely, if ever, directly compared. Therefore, we perform the first in-depth comparative study of two cornerstone action-value frameworks (i.e., Expected Threat (xT) [133] and VAEP [36]).

From an AI perspective, we aim to answer the following research question:

RQ12 How does a Markov model-based approach (i.e., Expected Threat [133]) conceptually and experimentally compare with a probabilistic classifier-based approach (i.e., VAEP [36]) for reasoning about the impact of actions?

From a soccer perspective, we aim to answer the following four research questions:

RQ13 What are the implications of the different design choices in xT and VAEP for the predicted values of specific actions?

RQ14 How do the rankings of top players differ when using xT versus VAEP?

RQ15 How do xT and VAEP compare with traditional player metrics such as goals and assists?

RQ16 How robust are the ratings produced by xT and VAEP?

1.2 Contributions and Structure of the Dissertation

This section provides an overview of the structure of this dissertation and summarizes its three main contributions in order of appearance.

1.2.1 Background

Chapter 2 provides the necessary background on the game of soccer and its collected data sources, the main techniques from artificial intelligence that are used in this dissertation, and the relevant related work in soccer analytics.

The content of this chapter is based on the following publications [34, 146, 148, 151].

Davis, J., Bransen, L., Devos, L., Jaspers, A., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Methodology and Evaluation in Sports Analytics: Challenges, Approaches, and Lessons Learned. *Machine Learning 113* (2024), 6977–7010

Van Roy, M., Robberechts, P., and Davis, J. A Markov Model Approach to Analyzing Build-Up Efficiency and Defensive Setups in Soccer. *Submitted to Data Mining and Knowledge Discovery* (2024)

Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. In *Artificial Intelligence in Team Sports Workshop at AAAI* (2020), pp. 1–8

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. A Markov Framework for Learning and Reasoning About Strategies in Professional Soccer. *Journal of Artificial Intelligence Research 77* (2023), 517–562

1.2.2 Contribution 1: A Critical Comparison of Action-Value Frameworks

Chapter 3 provides the first in-depth comparative study of two cornerstone frameworks for the most prominent use of learning techniques in soccer analytics: valuing actions. That is, it provides a comparative study of two action-value frameworks, Expected Threat (xT) [133] and VAEP [36]. These frameworks

are able to provide an estimate of how valuable different in-game actions are with respect to scoring goals. That is, they allow for basic reasoning about how an action changes the state of the game and thereby changes the team's (or the opponent's) chances of scoring a goal. Such frameworks can be used for scouting purposes, evaluating players' decisions, and creating game reports. In the literature, these frameworks are never directly compared and a deeper understanding of the implications of their design choices is currently lacking.

This contribution is based on the following publication [148].

Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. In *Artificial Intelligence in Team Sports Workshop at AAAI* (2020), pp. 1–8

1.2.3 Contribution 2: Contextualizing Action-Values by Reasoning

Chapter 4 discusses a novel approach to reason about the predictions of such action-value frameworks. There are two reasons why such reasoning approaches are necessary. First, the collected data might include erroneous measurements. Second, players sometimes perform unique actions that are dissimilar to what is mostly observed. However, learning techniques struggle to extrapolate beyond what is seen during training. In both cases, the predicted values might not be trustworthy, and it is beneficial to flag this to the end user. Especially because these models are often used in high-stakes decision-making settings (e.g., scouting) and it could be costly to base the decisions on erroneous information. Currently, approaches that can contextualize the predictions of these models and provide an understanding of their predictive behavior are lacking. This chapter outlines an approach to reason about the predictions made by a particular type of action-value frameworks, namely tree ensemble-based action-value frameworks such as VAEP [36].

This contribution is based on the following publication [144].

Van Roy, M., and Davis, J. DataDebugging: Enhancing Trust in Soccer Action-Value Models by Contextualization. In *Proceedings of 13th World Congress of Performance Analysis of Sport and 13th International Symposium*

on Computer Science in Sport (2023), vol. 1448 of *Advances in Intelligent Systems and Computing*, Springer Nature Switzerland, pp. 193–196

1.2.4 Contribution 3: A Framework for Learning and Reasoning About Policies

Next, this dissertation moves on from reasoning about predictions of individual actions to reasoning about overall tactics that are used during a game. Like the Expected Threat framework [133], we take a Markov perspective and utilize Markov models to learn a behavioral model of a team. We propose a novel way of reasoning about and analyzing various tactical aspects, both offensive and defensive, using probabilistic model checking techniques. Currently, no approaches exist that allow for in-depth reasoning about various tactical aspects.

Chapter 5 introduces our general learning+reasoning framework. It outlines how the offensive behavior of a team can be captured as a Markov model, how it can be learned from the collected data despite the many challenges, as well as how probabilistic model checking techniques can be used to reason about the effect of various offensive and defensive tactics.

Chapter 6 builds on this framework and provides an approach to integrate the defensive structure of the opponent and to reason about the efficacy of this defensive setup on the offensive output of the attacking team.

This contribution is based on the following publications [145, 146, 151].

Van Roy, M., Robberechts, P., and Davis, J. Optimally Disrupting Opponent Build-ups. In *StatsBomb Conference* (2021), pp. 1–16

Van Roy, M., Robberechts, P., and Davis, J. A Markov Model Approach to Analyzing Build-Up Efficiency and Defensive Setups in Soccer. *Submitted to Data Mining and Knowledge Discovery* (2024)

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. A Markov Framework for Learning and Reasoning About Strategies

in Professional Soccer. *Journal of Artificial Intelligence Research 77* (2023), 517–562

1.2.5 Conclusion

Chapter 7 summarizes the current research gaps, data challenges, and contributions of this dissertation, and provides possible directions for future research.

1.3 Other Research Contributions

This dissertation focuses on four first-author publications that concern learning and reasoning about actions or tactics in the context of soccer. However, this is only a subset of the research conducted during my doctoral studies. This section gives a short overview of the other research contributions.

SoccerMix: Representing Soccer Actions with Mixture Models

In this paper, we focus on learning the style of soccer teams and players, i.e., their prototypical behaviors during and across matches. Typically, capturing the style of teams and players is done by representing them in terms of the on-the-ball actions they perform on the pitch. To this end, the on-the-ball actions are divided into groups (i.e., clusters) of similar actions and the frequency of the performed actions within each cluster is computed for a team or player. This yields a vectorized summary of the prototypical actions that a team or player performs.

However, assessing the similarity of actions to decide which actions belong to the same cluster is challenging. First, teams and players rarely perform the same action in the same location multiple times. Second, these actions are described by various discrete (e.g., type of the action) and continuous (e.g., location, direction) attributes. One must design an approach that can overcome the sparsity of the data and generalize over the different actions while taking into account both discrete and continuous attributes. Existing approaches either (1) discretize the continuous attributes by e.g., discretizing the pitch into zones, or (2) do not consider the continuous attributes at all. By discretizing the continuous attributes, these approaches introduce arbitrary and

abrupt boundaries, which can cause spatially similar actions to appear dissimilar. Additionally, these approaches could even further increase the sparsity when the resolution of the grid is too fine-grained or additional attributes need to be taken into account. By not considering the continuous attributes, these approaches ignore a very important aspect of play, i.e., the location and direction of the actions.

In contrast to these hard clustering approaches, we apply a mixture models-based approach to soft cluster the actions based on their type, location, and direction. This approach addresses the shortcomings of the grid-based approaches. We illustrate `SoccerMix` on a number of use cases, which show that our approach cannot only give insight into a team's or player's offensive style but also give insight into a team's defensive style.

This contribution has been published at a peer-reviewed conference [42]. The implementation is publicly available at `https://github.com/ML-KULeuven/soccermix`.

Decroos, T., Van Roy, M., and Davis, J. SoccerMix: Representing Soccer Actions with Mixture Models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)* (2020), vol. 12461 of *Lecture Notes in Computer Science*, Springer, pp. 459–474

Leaving Goals on the Pitch: Evaluating Decision Making in Soccer

In this paper, we focus on reasoning about various aspects of long-distance shooting behavior (i.e., shots resulting from outside the penalty box). This has been inspired by the analytical-driven trend that long-distance shots have steadily decreased over the past decade in all major leagues, although they might still be valuable.

In fact, there is a risk-reward trade-off to forgoing a shot. If a player shoots immediately when in possession outside the penalty box, this ensures at least the chance of scoring. The potential payoff of not shooting is that an even better shot may arise down the line, but there is no guarantee of this happening as the team may lose the ball prior to another shooting chance arising. To answer various questions related to which options are better and what the effects are of shooting more or less often from outside the penalty box, reasoning approaches are required. On the one hand, one must reason about the number of ways a possession could play out as well as the inherent uncertainty of actions

succeeding (e.g., a player mishitting a ball). On the other hand, understanding the effect of different shooting policies requires counterfactual reasoning: one needs to reason about what could have happened if a different policy had been followed.

In this paper, we propose a novel generic framework to reason about such questions by combining a learning and reasoning approach. In particular we model the offensive behavior of a team by learning a Markov decision process from their event stream data. Next, we apply both statistical and probabilistic model checking techniques to analyze a team's chance of scoring from particular situations as well as the effects of modifying their shooting policy.

This contribution has been published at the MIT Sloan Sports Analytics Conference [150]. It has received some media attention such as an interview by John Muller for FiveThirtyEight[2] and an article in the New York Times[3]. It is also the basis for our work in Chapters 5 and 6.

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. Leaving Goals on the Pitch: Evaluating Decision Making in Soccer. In *15th MIT Sloan Sports Analytics Conference* (2021), pp. 1–25

Looking Beyond the Past: Analyzing the Intrinsic Playing Style of Soccer Teams

In this paper, we focus on learning the offensive style of soccer teams. Previous approaches, including our `SoccerMix` approach, tackle this by looking at individual on-the-ball actions that are performed during a match in isolation. Thus, these approaches ignore the sequential aspect of the game, which is important to model patterns of play. Additionally, these approaches are not able to generalize over the limited amount of data in order to capture slight variations of play that a team may use in the future. In summary, these existing approaches are not able to capture the *intrinsic* playing style of teams.

We propose a novel approach to playing style analysis by learning a model that captures a team's intrinsic offensive behavior. That is, we model the team's offensive behavior as a discrete-time Markov chain that is learned from event stream data. This allows us to incorporate the sequential nature of the game and generalize better over the limited amount of data. Afterward, we characterize

[2] `https://fivethirtyeight.com/features/is-soccer-wrong-about-long-shots/`

[3] `https://www.nytimes.com/2021/04/17/sports/soccer/erling-haaland-transfer.html`

playing style by defining various features based on domain knowledge, and extract values for the features from the learned model by reasoning on the model using probabilistic model checking and analytical techniques. We illustrate our approach for a number of teams and use cases.

Although this paper focuses on the combination of learning and reasoning in the context of playing style analysis in soccer, this paper is not included in this dissertation because it is based on the master thesis of Jeroen Clijmans, for whom I was the daily advisor during the 2020-2021 academic year.

This contribution has been published at a peer-reviewed conference [33]. The implementation is publicly available at `https://github.com/JeroenClijmans/MarkovSoccer`.

Clijmans, J., Van Roy, M., and Davis, J. Looking Beyond The Past: Analyzing The Intrinsic Playing Style Of Soccer Teams. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)* (2022), vol. 2035 of *Communications in Computer and Information Science*, Springer, pp. 370–385

un-xPass: Measuring Soccer Player's Creativity

In this paper, we propose an approach to measure a player's passing creativity. Creativity is a highly valued skill as it can help overcome defensive tactics of the opponent and create goal-scoring opportunities by introducing unpredictability in the game. In addition, it also makes the game more exciting to watch. Therefore, being able to quantify a player's creativity is especially important for the recruitment process.

Current metrics to value players typically consider the result of the performed actions to decide upon how valuable each action, and in aggregate, a player is. However, creativity goes beyond the result of the performed actions. Being creative means to do something that is both valuable as well as novel or unique.

We propose to capture creativity by defining creative actions as those that are (1) different from the typical action that other players would chose in a given game situation, and (2) more valuable than the typical choice. We employ and compare various machine learning methods to learn predictive models for each of these components from contextual event stream data. With our approach, we inspect (1) which players are most creative in their passing, (2) which pairs

of players are creative together, and (3) how creativity is impacted by the game state.

This contribution has been presented at the soccer analytics-oriented StatsBomb Conference [121] and has been published at a peer-reviewed conference [122]. This work has also been covered by an article on ESPN.[4] The implementation is publicly available at `https://github.com/ML-KULeuven/un-xPass`.

Robberechts, P., Van Roy, M., and Davis, J. un-xPass: Measuring Soccer Player's Creativity. In *StatsBomb Conference* (2022), pp. 1–21

Robberechts, P., Van Roy, M., and Davis, J. un-xPass: Measuring Soccer Player's Creativity. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2023), Association for Computing Machinery, p. 4768–4777

Methodology and Evaluation in Sports Analytics: Challenges, Approaches, and Lessons Learned

As sports data are being increasingly collected and machine learning methods frequently applied to this data, in this paper, we aim to provide a structured framework of both methodological and evaluation approaches and challenges. We start by giving an overview of the different types of data that are being collected and common sports analytics tasks they are used for. Next, we highlight various challenges that arise when working with this data, both from a methodological and evaluation perspective. Finally, we discuss how the machine learning models and their derived indicators should be evaluated, and summarize the lessons learned for future research.

This contribution has been published as a workshop paper [35], and has been extended and published as a peer-reviewed journal paper [34].

Davis, J., Bransen, L., Devos, L., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Evaluating Sports Analytics Models:

[4] `https://www.espn.com/sports/soccer/insider/story/_/id/34720311/kevin-de-bruyne-tariq-lamptey-most-creative-new-research`

Challenges, Approaches, and Lessons Learned. In *AI Evaluation beyond Metrics Workshop at IJCAI* (2022), vol. 3169, CEUR-WS, pp. 1–11

Davis, J., Bransen, L., Devos, L., Jaspers, A., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Methodology and Evaluation in Sports Analytics: Challenges, Approaches, and Lessons Learned. *Machine Learning 113* (2024), 6977–7010

ETSY: A Rule-Based Approach to Event and Tracking Data Synchronization

In this paper, we tackle a major engineering challenge that is rarely covered in the literature. Namely, there are two widely used data sources in soccer, event stream data and tracking data, which capture complementary information and which need to be combined to allow one to use the full context during the analyses. On the one hand, event stream data captures the semantic information of the actions during a game, but lacks broader context information on the positioning of the other players. On the other hand, tracking data does capture this broader context, but lacks the semantic information on which actions are performed. To perform contextualized analyses, one must take both data sources into account. However, these are often not time-synchronized with each other. That is, the event at a specific timestamp does not correspond to the tracking frame with the same timestamp. While a human might be able to synchronize both streams based on what he sees, this is very time-consuming. In this paper, we hypothesize that by reasoning about various features of the tracking frames and events using a rule-based approach, mimicking the human-like behavior, it would be possible to find similarities between them so that a matching tracking frame can be found for each event. To this end, we introduce `ETSY`, a rule-based approach that is able to synchronize event and tracking data.

This contribution has been published as a peer-reviewed workshop paper [143]. The implementation is publicly available at `https://github.com/ML-KULeuven/ETSY`.

Van Roy, M., Cascioli, L., and Davis, J. ETSY: A Rule-Based Approach to Event and Tracking Data SYnchronization. In *Machine Learning and Data*

Mining for Sports Analytics Workshop at ECML/PKDD (2023), vol. 2035 of *Communications in Computer and Information Science*, Springer Nature Switzerland, pp. 11–23

Chapter 2

Background

In this chapter, we provide the necessary background for the contributions of this dissertation. First, we provide a brief outline of the sport of soccer (Section 2.1). Second, we give an overview of the main data sources collected from soccer (Section 2.2). Third, we discuss the main artificial intelligence techniques that will be used in this dissertation (Section 2.3). Finally, we provide an overview of the relevant related work in soccer analytics that is crucial to this dissertation (Section 2.4).

The content of this chapter is based on the following publications [34, 146, 148, 151].

Davis, J., Bransen, L., Devos, L., Jaspers, A., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Methodology and Evaluation in Sports Analytics: Challenges, Approaches, and Lessons Learned. *Machine Learning 113* (2024), 6977–7010

Van Roy, M., Robberechts, P., and Davis, J. A Markov Model Approach to Analyzing Build-Up Efficiency and Defensive Setups in Soccer. *Submitted to Data Mining and Knowledge Discovery* (2024)

Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. In *Artificial Intelligence in Team Sports Workshop at AAAI* (2020), pp. 1–8

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. A Markov Framework for Learning and Reasoning About Strategies in Professional Soccer. *Journal of Artificial Intelligence Research 77* (2023), 517–562

2.1 Soccer

Soccer is a ball sport that is played between two teams on a grass pitch. Pitches can have a varying size, although the standard dimensions defined by IFAB (i.e., International Football Association Board) is 105 meters by 68 meters. A pitch is illustrated in Figure 2.1. 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 2.1). 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.

Shot A shot is any type of action that deliberately tries to move the ball into the opponent's goal.

Pass 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 pass).

Cross 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.

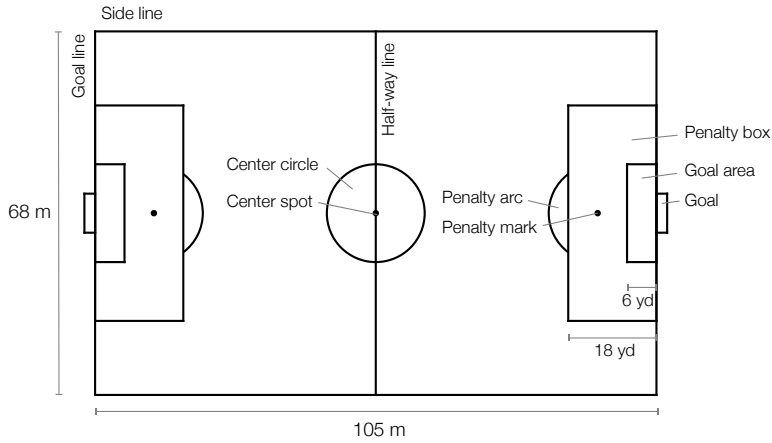

Figure 2.1: Illustration of a soccer pitch. Each team has their own goal (drawn outside the pitch) at opposite sides of the pitch. The two rectangles near each team's goal are mainly important for goalkeepers: from within the six-yard box (small rectangle) the goalkeepers can perform goal kicks, within 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 kickoffs.

Dribble 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. The definition of dribbles used in this dissertation 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 Data Sources

There are a variety of data sources collected about soccer matches and training sessions. Such data is not only available for soccer, but also for a number of other sports (e.g., American football, basketball, and ice hockey). These

data sources can broadly be divided into two categories: physical data and observational match data.

2.2.1 Physical Data

Physical data mainly concerns the physical parameters of players, and can be collected during both training sessions and matches. Athletes often wear a GNSS (Global Navigation Satellite System) tracker with inertial measurement unit (IMU) technology (e.g., from Catapult Sports[1]). These systems measure various physical parameters such as distance covered, number of high-speed sprints, high-intensity accelerations, decelerations, and impacts. In addition, physiological data such as heart rate, blood oxygen level, and skin temperature may also be measured (e.g., using Polar[2] wearable devices). For a training session, these parameters are often augmented with questionnaire data [24] to obtain subjective measurements such as the rating of perceived exertion (RPE) [15] which can be used to capture the fatigue of a player.

This data is typically used to monitor and eventually optimize a player's fitness level [16, 75], and to predict and reduce their injury risk [2, 53], with the overarching aim being to ensure the player's availability and ability to compete. As this dissertation focuses on supporting player recruitment and match preparation by learning and reasoning about in-game actions and tactics, this type of data will not be used.

2.2.2 Observational Match Data

Observational match data concerns information that can be collected by visually observing a match. For soccer, it can be divided into three broad categories of data sources: match sheet data, event stream data, and tracking data. These data sources differ from each other in two main aspects: availability and included context. The more context is included in a data source, the less widely available it is and vice versa [142]. Next, we list and discuss these data sources from least contextualized to most contextualized.

[1] `https://www.catapult.com`
[2] `https://www.polar.com`

Match sheet data

Match sheet data is the most common and widely available source of data (e.g., it is reported on websites such as FBRef[3] and WhoScored[4]). These include aggregate statistics for the team and/or the players on a per-match basis. For example, this could contain information such as line-ups, substitutions, cards, goals and assists. More advanced details such as the number of passes attempted and completed per player or team may also be included. Moreover, for a number of sports, this data has been collected for decades.

Event stream data

Another widely available source of data collected about soccer matches is event stream or play-by-play data with vendors (e.g., StatsBomb[5], Wyscout[6], and Stats Perform[7]) collecting this data for 100s of leagues worldwide. Typically, these vendors employ human annotators, possibly assisted by computer vision techniques, to collect the data by watching broadcast video [13, 154]. The inclusion of human annotators is needed because event stream data aims to capture the semantics of the game by describing all on-the-ball actions that occur. For each on-the-ball action, event stream data records a number of features such as the type of 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. In total, each match typically has around 1500 to 3000 on-the-ball actions.[8] Figure 2.2 illustrates a sequence of four actions from a game between Manchester City and Liverpool as they were recorded in the event stream data format.

Recently, a more contextualized version of event stream data was introduced by StatsBomb.[9] In addition to the above event stream data features, it also includes a snapshot at the moment of each annotated event with information about off-the-ball players that are visible on camera such as their location and relationship to the ball carrier (i.e., teammate or opponent). Even though the moments between actions are not captured, this contextualized event stream data can

[3] `https://fbref.com`
[4] `https://www.whoscored.com`
[5] `https://www.statsbomb.com`
[6] `https://footballdata.wyscout.com`
[7] `https://www.statsperform.com`
[8] The number of on-the-ball actions depends on what is annotated in the data (e.g., pressure events, separate actions for pass and receival).
[9] `https://statsbomb.com/what-we-do/soccer-data/360-2/`

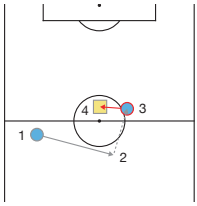

	Time	Action	Start	End	Result	Player	Team
1	37:16	pass	(57, 11)	(64, 38)	success	L. Sané	Man City
2	37:19	dribble	(64, 38)	(48, 43)	success	J. Stones	Man City
3	37:22	pass	(48, 43)	(48, 33)	fail	J. Stones	Man City
4	37:24	interception	(48, 33)	(48, 33)	success	R. Firmino	Liverpool

Figure 2.2: 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.

	Time	Action	Result	Bodypart	Start	End	Player	Team
1	9m7s	pass	success	foot	(64.8, 14.0)	(62.4, 20.1)	Toby Alderweireld	Belgium

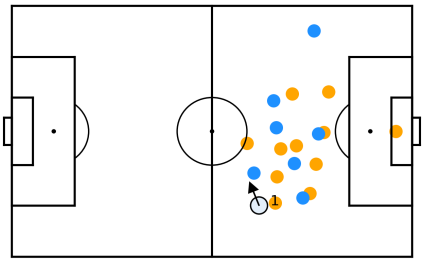

Figure 2.3: Illustration of one action from the Belgium versus Portugal game on June 27, 2018, as recorded in the contextual event stream data format.

already provide a broader view on how the different players are positioned during each action compared to standard event stream data. Figure 2.3 illustrates one action from a game between Belgium and Portugal as it was recorded in the contextual event stream data format.

As event stream data is widely available, all chapters of this dissertation will make use of it to learn and reason about soccer.

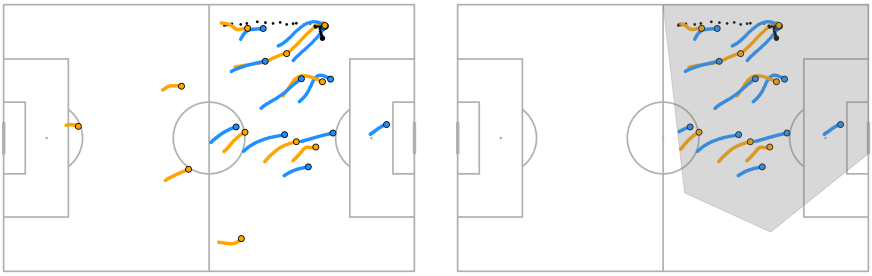

(a) Tracking data (b) Broadcast tracking data

Figure 2.4: Illustration of (a) a sequence of tracking data, and (b) a sequence of broadcast tracking data where the grey area denotes the field of view of the camera.

Tracking data

There are two types of tracking data, also called positional data, available: in-stadium and broadcast tracking data.[10] In-stadium tracking data is collected using a fixed installation in a team's stadium using high-resolution cameras (e.g., Stats Perform's SportVU system[11]). It reports the locations of all the players and the ball multiple times per second (typically between 10 and 25 Hz). Such an in-stadium setup is expensive and typically only used in top leagues. Broadcast tracking data on the other hand aims to solve this by extracting the tracking data from broadcast video shown on television (e.g., SkillCorner[12] takes this approach). However, in broadcast video, the entire pitch and the players on it are not always visible. Throughout the game, replays can be shown and the cameras can zoom in on specific players, spectators, or parts of the field. Therefore, the obtained tracking data is not complete and might miss the coordinates of all or a few players, or even the ball, for a number of seconds. It is also not straightforward to obtain because one must deal with, amongst other challenges, low-resolution images, occlusions of players, and a non-trivial mapping between video and pitch coordinates. Nevertheless, broadcast tracking data is a cost-effective solution for many teams. Thus, there is extensive work on tracking solutions based on broadcast video for a variety of sports [4, 74, 76]. Figure 2.4 shows a short sequence of (broadcast) tracking data.

[10]The global positioning data from the GNSS sensor, used to collect various physical parameters, could also be seen as another type of tracking data. However, this is not done at a large scale.

[11]`https://www.statsperform.com/team-performance/football-performance/sportvu`

[12]`https://skillcorner.com/tracking-data`

2.3 Artificial Intelligence Techniques

This section contains an overview of the main artificial intelligence techniques that will be used in the following chapters to learn and reason about data collected from soccer matches.

2.3.1 Probabilistic Classification

Probabilistic classification is a common learning task in artificial intelligence. Formally, given a data set with examples (x, y) with $x \in X$ the features that describe the example and $y \in Y$ the class (i.e., label) of the example, probabilistic classification models aim to learn the conditional distribution $P(Y|X)$ that assigns for a given $x \in X$ a probability distribution over all possible classes $y \in Y$, with $\sum_{y \in Y} P(Y = y|X = x) = 1$ [126]. In the case of two classes $Y = \{0, 1\}$, which typically denote a negative and positive class respectively, this is called binary probabilistic classification. A binary probabilistic classifier thus learns a function $f : X \to [0, 1]$ which assigns for each example $x \in X$ the probability of the example being positive [126].

There are various models for probabilistic classification, such as logistic regression [126], neural networks [58], and tree-based approaches (e.g., gradient boosted decision trees [31], random forests [22]). Tree-based approaches are among the most popular approaches in sports due to their (near) state-of-the-art predictive performance, interpretability, and ease of use. Typically, these are used to predict the probability of a goal being scored (conceded) in the near future, given the current state of the game. We will also use them for this task in Chapters 3, 4, and 5. In addition, Chapters 5 and 6 make use of tree-based probabilistic classification models to estimate a probability distribution over the likely destination of passes, dribbles, and crosses.

2.3.2 Markov Models

Markov models are a core framework in artificial intelligence that is ubiquitously used to describe, control, and analyze the behavior of complex dynamic systems. It captures the system as a state-transition model in which the states describe information of the system at a certain point in time and the transitions specify how it can transition from one state to another. There exist multiple types of Markov models. In this dissertation, the following two popular types are used.

Discrete-time Markov chain

A discrete-time Markov chain (DTMC) models a stochastic process as one that can transition between states in discrete time steps. Formally, a DTMC is a tuple $M = \langle \mathcal{S}, P \rangle$ where $\mathcal{S}$ is a finite set of states and $P : \mathcal{S} \times \mathcal{S} \to [0, 1]$ is the transition function. The probability of transitioning from a state s to a state s' (i.e., $P(s, s')$) only depends on the current state, also called the Markov property. The transition function P must denote a proper probability distribution for all states, that is, $\sum_{s' \in \mathcal{S}} P(s, s') = 1$. A state can be *transient* or *absorbing*. A transient state can be passed through infinitely often. An absorbing state cannot be escaped once entered, and it denotes the end of the episodic task (either success or failure) [77].

Markov decision process

A Markov decision process (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, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ is the transition function, $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}$ is the reward function, and γ is the discount factor. The transition function P must denote a proper probability distribution for all state-action pairs (s, a), that is, $\sum_{s' \in \mathcal{S}} P(s, a, s') = 1$ [12].

Given an MDP, a path of length n is denoted by $\rho_n = s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} ... \xrightarrow{a_{n-1}} s_n$ 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 $\rho = s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} ....$

A policy $\pi : \mathcal{S} \times \mathcal{A} \to [0, 1]$ specifies, for each state, a probability distribution over all actions. Given an MDP and a policy π, the value function $V_\pi : \mathcal{S} \to \mathbb{R}$ represents the value of being in a state. The value function is denoted as V [12].

$$V_\pi(s) = \sum_{a \in \mathcal{A}} \pi(a \mid s) \sum_{s' \in \mathcal{S}} P(s, a, s')(R(s, a, s') + \gamma V_\pi(s')) \tag{2.1}$$

A Markov reward process (MRP) is a specialization of an MDP which allows for exactly one fixed policy and a reward function. Like an MRP, a DTMC is a specialization of an MDP which allows for exactly one fixed policy, but without an associated reward function [7].

Markov models can be used to learn a representation of the underlying dynamic system given data that is generated by the system. The Markov model, in combination with additional reasoning techniques, can be used to reason about the modeled behavior and extract insights. We will use Markov models in this manner in this dissertation. Chapter 3 discusses an existing approach to value

actions in soccer on the basis of Markov models. Chapter 5 and 6 introduce a novel Markov model approach to model the offensive behavior of teams and the defensive structure of the opponent.

2.3.3 Bayesian Modeling and Probabilistic Programming

Bayesian models are statistical models that apply Bayesian inference to learn the models' parameters from data. The models' parameters are treated as random variables, about which prior beliefs can be expressed in terms of probability distributions, and which are updated on the basis of the provided data using Bayes' theorem.

$$P(\Theta|D) = \frac{P(D|\Theta)P(\Theta)}{P(D)} \tag{2.2}$$

Here, $P(\Theta|D)$ is the posterior probability (i.e., the probability of the parameters Θ given the data D), $P(D|\Theta)$ is the likelihood (i.e., the probability of observing the data D given the parameters Θ), $P(\Theta)$ is the prior probability (i.e., the prior beliefs about the parameters' values), and $P(D)$ is the marginal likelihood or model evidence (i.e., the probability of the data, regardless of the parameters' values). This approach allows one to incorporate the domain knowledge one has about certain parameters into the model, which can be beneficial in a sports setting. It also allows one to handle uncertainty, both about the parameters as well as the predictions, which is required in a highly dynamic setting such as soccer.

As an example, one could model the expected number of goals that a team scores $y_{>t}$ after time t as a Poisson distribution $y_{>t} \sim \text{Poisson}((T-t) * \theta_t)$, with the parameters θ being the scoring intensities of the team at time t. Prior beliefs, such as a linear relationship between the scoring intensity and the input data, $\theta_t = \text{invlogit}(\vec{\alpha}_t * x_t + \beta)$, can be formulated about the parameters. This in itself is not yet a Bayesian approach, due to the deterministic linear relationship. However, the parameters $\vec{\alpha}_t$ and β can also be specified with prior distributions, depending on the prior knowledge about these parameters (e.g., a uniform distribution can be used in case no prior knowledge is known), yielding a Bayesian approach that incorporates uncertainty about these parameters [120].

Bayesian models can be expressed in probabilistic programming languages such as Stan [135] and PyMC [128] which allow for estimating the posterior distributions of the models' parameters from data. Typically, the parameters of the models are inferred from the given data using Markov Chain Monte Carlo (MCMC) algorithms. However, these algorithms are computationally expensive

and do not scale well with large data sets, which can be problematic in sports applications. Approximate inference methods, such as Automatic Differentiation Variational Inference (ADVI) [81], can be used to provide approximations of the models' parameters while maintaining scalability to large data sets.

Bayesian models can thus be used to model the probabilities of certain events occurring in a given data set, allow for specifying prior knowledge, and incorporating uncertainty. Hence, they can be used to estimate the parameters of Markov models, as will be done in this dissertation in Chapter 5.

2.3.4 Probabilistic Model Checking

Probabilistic model checking is a formal verification technique that determines whether a stochastic system satisfies a given desired stochastic property [68, 83]. 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 [52]. We consider MRPs, MDPs, and PCTL* (which subsumes Probabilistic Computational Tree Logic (PCTL) and Probabilistic Linear Temporal Logic (LTL)) [52].

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 ϕ) and a path formula (denoted by ψ) over $\mathcal{AP}$ can be constructed by the grammar below. A PCTL* property is always a state formula ϕ.

$$\text{state formula } \phi ::= \mathit{true} \mid c \mid \phi \wedge \phi \mid \neg\phi \mid \mathtt{P}_{\bowtie \mathtt{p}}[\psi] \tag{2.3}$$

$$\text{path formula } \psi ::= \phi \mid \mathtt{X}\ \psi \mid \psi\ \mathtt{U}^{\leq \mathtt{k}}\ \psi \mid \psi\ \mathtt{U}\ \psi \mid \neg\psi \tag{2.4}$$

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

Intuitively, $\mathtt{X}\ \alpha$ means α holds in the next step, $\alpha\ \mathtt{U}\ \beta$ means α holds until β holds, $\alpha\ \mathtt{U}^{\leq \mathtt{k}}\ \beta$ means α holds until β holds and β will hold within $\mathtt{k}$ steps, and $\mathtt{P}_{\bowtie \mathtt{p}}[\alpha]$ means α holds with a probability $\bowtie \mathtt{p}$. Additionally, the commonly used temporal component $\mathtt{F}\ \alpha$ (α will eventually hold) can be defined by the until operator: $\mathtt{F}\ \alpha \equiv \mathtt{true}\ \mathtt{U}\ \alpha$. General examples of PCTL* properties are $\mathtt{P}[\mathtt{F}\ A]$ ("probability of finally reaching a state A"), or $\mathtt{P}[\neg A\ \mathtt{U}\ B]$ ("probability of never entering A before reaching B"). A specific example for soccer could be $\mathtt{P}[\mathtt{F}\ \mathit{penalty_box}]$ ("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 to Baier [6] for more details of PCTL*.

Probabilistic model checking techniques will be used in Chapters 5 and 6 of this dissertation to reason about the modeled behavior of teams. It will allow us to extract insights regarding a team's offensive patterns, their effectiveness, as well as the efficacy of defensive tactics employed by the opponent.

2.4 Related Work

This section summarizes the related work of two main soccer analytics areas that are relevant for this dissertation: quantifying player performance, and discovering and analyzing tactics.

2.4.1 Quantifying Player Performance

A fundamental task in soccer analytics is to objectively quantify a player's performance during a match. Typically, the goal is to summarize a player's contribution to the team's performance using one or a handful of numbers. This can help inform a variety of different decisions that a club must make in areas such as team selection, opponent scouting, and player acquisition. For example, it allows for assessing a player's creativity [122], their mutual chemistry [19], as well as how they perform under mental pressure [17]. Furthermore, these approaches facilitate fan engagement as they provide fodder for debating the relative merits of different players [36, 99] or they can help tell the story of a match [39]. Metrics that quantify a player's performance can be obtained from the available data in two ways. They can be (1) directly derived from the data, or (2) estimated through machine-learned models [80].

Derived metrics

Derived metrics, also called count-based statistics or box scores, typically give insight into *what* players do on the pitch and *how often* they do it. This can be extracted directly from the data by counting how often a player performs a certain type of action (e.g., pass, shot) or by aggregating them over action attributes (e.g., pass success ratio, backward pass proportion) [159]. In addition, some derived metrics attempt to give insight into *how valuable* the performed actions are, which is commonly done by assigning a fixed value to each specific action type, and aggregating these values based on the count-based statistics to further quantify a player's performance [98, 99, 109].

As these metrics mainly concern the actions players perform, this is typically derived from match sheet data or event stream data. While these metrics can be useful to help distinguish between different player profiles (e.g., an offensive vs. a holding midfielder), a main limitation is that they largely do not take the context in which players perform these actions into account. For example, two midfielders might perform a similar number of forward passes, but one might perform them more often in a challenging (and possibly more valuable) setting with high pressure from the opponent whereas the other performs them more often in an easy setting with limited to no opponent pressure. However, accounting for this additional context with count-based statistics is increasingly difficult due to the sparsity of the data. The more context is included (e.g., specific action type, location, time, setup of players, pressure), the fewer actions the data will contain that match the specific context. Therefore, all actions of the same type, regardless of any additional context, are usually treated equivalently, and it becomes impossible to separate the two aforementioned midfielders.

Model-based metrics

Model-based metrics, on the other hand, are estimated through statistical or machine learned models using historical data. This means they can learn the impact of different contexts on the specific aspect of the game they aim to quantify. Within the model-based metrics, we can make a distinction based on the following design choices:

- What aspect of the game is quantified, or *valued*?
- What defines the value?
- Which contextual features are taken into account to estimate the value?
- What type of learning model is used?

The majority of model-based metrics focuses on measuring the quality of *one specific action type* in a variety of specific game situations. The most prominent example is the expected goals (xG) metric [59] which aims to quantify the quality of each scoring opportunity. It does so by assigning a probability to each shot that represents its chance of directly resulting in a goal. 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 [27, 38, 73, 79, 93, 119]. The context can range from relatively straightforward such as just a location 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). Depending on the context that is included, either event stream data or a combination of (contextual) event stream data and tracking data are used. Learning these models is typically done using standard probabilistic classifiers such as logistic regression or tree ensembles using large historical data sets of shots. Some approaches build on top of this methodology and measure the impact a player has on a specific target metric while they are on the pitch (e.g., their impact on the team's xG value or total number of goals) [72, 78, 127, 137].

xG has become a well-accepted metric that appears in mainstream media outlets and is mentioned by managers. However, shots represent less than 1% of all on-the-ball actions. This means that only a very small fraction of a player's performance can be quantified this way. Therefore, various other approaches have been developed that focus on other action types, such as passes and dribbles [18, 20, 41, 64, 85, 97, 104, 116]. They typically value these actions based on the likelihood of scoring a goal, either immediately (e.g., by treating each non-shot action as if it were a shot and assigning xG values to these [85, 97, 104]) or later on in the possession (e.g., by explicitly collecting similar possessions and calculating the proportion of these that result in a goal [18, 20, 41, 64]).[13] Because of their focus on specific actions, these approaches typically make use of event stream data. Similar learning models as for xG are employed, with in addition approaches such as nearest neighbors search for identifying similar possession sequences.

Instead of designing a separate model for each action type, more recent approaches have attempted to join these models together in a unifying framework that can value *a wide range of action types* in varying game scenarios [36, 50, 51, 82, 87, 102, 125, 133, 136, 158, 162]. In general, they all use the same type of reasoning. A soccer match can be viewed as a sequence of n consecutive actions $a_1, a_2, \dots a_n$, and each on-the-ball action changes the state of the game. Thus, the effect of an action a_i is to move the game from state $S_{i-1} = \{a_1, \dots, a_{i-1}\}$ to state $S_i = \{a_1, \dots, a_{i-1}, a_i\}$. Conceptually, the vast majority of these approaches can be seen as modeling the game using a Markov model perspective. At a high-level, these approaches all value actions according to the following equation:

$$V(a_i) = Q(S_i) - Q(S_{i-1}) \tag{2.5}$$

where Q captures the value or quality of a particular game state. These approaches are also referred to as Possession State Value models. Other

[13] What exactly constitutes a possession may differ. Most approaches consider a possession to be a sequence of ball control by one team until a turnover or a goal. Other approaches consider possessions to start and end with kickoffs and goals, and the ball may switch teams throughout the possession.

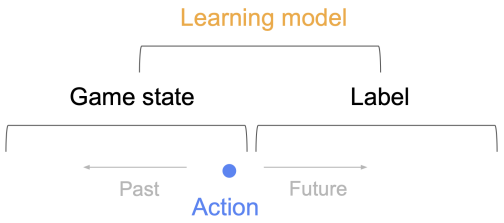

Figure 2.5: Generalized illustration of the approach taken for valuing game states. The various methods mainly differ in how they assign values to the game states, which context is taken into account to describe the game states, and which machine learning model is used to learn the relationship between the game states and their associated values.

approaches go beyond valuing the on-the-ball actions, and instead attempt to value *every moment* in the game. That is, these approaches attempt to measure how valuable it is to have possession (i.e., control) of the ball in any given game state, and are often referred to as Expected Possession Value (EPV) models [45, 46, 50, 51, 86]. The main differences among all the various approaches arise in how they assign values to the game states, which context is taken into account to describe the game states, and which machine learning model is used to learn the relationship between the game states and their associated values. This is summarized in Figure 2.5.

There are two common approaches to assign values to game states: scoring-based and win-based. With scoring-based approaches, the value that is assigned to a game state relates to the probability of scoring (or conceding) a goal in the short term (e.g., within a number of actions/seconds [36, 82, 158]) or long term (e.g., at the end of the possession sequence [50, 51, 87, 102, 125, 133, 136, 162]). An action's value thus denotes the impact on the team's chances of scoring (or conceding) a goal. With win-based approaches, the value of a game state is related to the probability of the team in possession eventually winning the game [120]. An action's value then denotes the impact on the team's chances of winning (or losing) the game. In principle, one could assign various other custom rewards as well, depending on what is deemed valuable by practitioners (e.g., moving the ball closer to the opponent's goal [45, 46]).

The context that is taken into account also varies broadly. Given that most of these approaches attempt to value the performed actions, event stream data is needed to extract semantically relevant information about the actions. Like the approach taken for xG models, many features can be

included ranging from simple location-based information about the action under consideration [102, 133] to more complex information (e.g., distance and angle to goal, used body part, game time, score differential, defensive pressure, and set-piece information) [87, 125, 136, 162], not only of the action under consideration but also of previous actions in the same possession [36, 82, 158]. In addition, some approaches make use of tracking data to include extra contextual information (e.g., the locations and velocities of the other offensive and defensive players) [45, 49, 50, 51, 86, 134]. These approaches typically not only value the on-the-ball actions, but attempt to value all moments in the game. However, the vast majority of work focuses on valuing the on-the-ball actions using event stream data as it is more widely available, both in terms of leagues covered and availability to clubs.[14]

The learning approaches used to learn the relationship between the game states and their associated values is akin to the approach taken for xG models. Commonly used models are probabilistic classification approaches such as logistic regression, decision tree ensembles, and neural networks to estimate the probability of a goal being scored in the short or long term given a certain game state [36, 49, 50, 51, 82, 136, 158]. Some approaches explicitly model the possessions of teams as transitions between game states in a Markov model [51, 87, 125, 133, 162], which often results in more interpretable models.

In Chapter 3, we compare two state-of-the art action-value frameworks that are a canonical example of, on the one hand, a Markov model-based approach that models entire possessions, and on the other hand, a pure action-based approach.

2.4.2 Discovering and Analyzing Tactics

Besides quantifying players' and teams' performance, discovering and analyzing the tactics and strategies used by the teams is another major task in soccer analytics. Both terms (i.e., tactics and strategies) are often used interchangeably to denote the same concept, although distinctions with respect to the time scale on which they operate are sometimes made [21, 28, 117]. On the one hand, tactics can be viewed as short-term executions or in-game behavior of teams and players. This is summarized as *which player(s)* perform *which action(s)* in *which area(s) of the pitch*. Tactics can be hierarchically organized into individual tactics (e.g., offensive and defensive actions of a player), group tactics (e.g., the defensive block used by the defenders and midfielders), team tactics (e.g., the team's overall positioning on the pitch), and match tactics (e.g., whether a team

[14] Often, tracking data is not shared across leagues, which makes event stream data valuable for player recruitment purposes. Broadcast tracking data is only recently emerging as a supplementary resource for player recruitment.

is more possession-oriented or focuses on counter attacks). This formulation is often used by soccer practitioners [28, 117]. On the other hand, strategies can be viewed as more long-term plans that are made pre-game which can include player scouting, recruitment, and line-up and substitution selection. The chosen strategies naturally influence the in-game tactics that teams and players use. We follow the definition of tactics in the literature and in the remainder of this dissertation we will mainly focus on the concept of tactics (i.e., the in-game behavior of teams and players) as opposed to strategies.

Discovering and analyzing tactics in soccer has been tackled using multiple different approaches. Clustering techniques are commonly used to detect where players perform certain actions on the pitch. That is, these approaches attempt to identify the prototypical action types and/or the locations where they are often performed. Using these prototypical actions and/or locations, one can count how often a team or player performs these and use this to characterize their tactics [42, 61, 66, 94, 105, 106]. However, as these approaches focus on single actions in isolation, they do not consider the interplay between different actions and how these result in typical patterns. Clustering techniques have also been used to detect team and group tactics such as team formations [14] and player subgroups [55].

Instead of focusing on single actions, other approaches take a possession-based perspective. There are two main lines of work within this perspective. On the one hand, there are approaches that aim to find similar and frequently used possession sequences, as was also done in early approaches to value player contributions [18, 20, 41, 64]. To this end, various clustering techniques have been employed to identify groups of similar possession sequences [40, 103, 156], where similarity can denote e.g., the types of actions and their locations, but also the speed of play within a possession sequence. In addition, pattern mining techniques [40, 62] and inductive logic programming [141] have also been applied to uncover frequently used possession sequences. On the other hand, there are approaches that aim to model the in-game behavior of teams and players by modeling their possession sequences. To this end, Markov model-based approaches are commonly used. Even though these models yield a concrete description of the in-game behavior of teams, the most prominent use cases are however not to analyze the in-game tactics, but to evaluate player performance by valuing their contributions [30, 51, 124, 125, 131, 133, 162], and to assess win probability [47, 48, 91, 160]. A handful of approaches perform tactical analysis based on the constructed Markov models such as pass network analysis in soccer [51], shot analysis in basketball [130], and stroke analysis [115] and analysis of tactical changes [157] in tennis. In addition, Markov models have also been used to evaluate longer term strategies such as substitutions in soccer [70].

A final line of research focuses on the use of graphs (networks) to model the

interactions between players or teams. Here, graphs are often used to analyze ball-movement patterns between players or the (sub-)groups of players that are present, with various applications ranging from safe passing configurations, detecting player patterns, to measuring the stability of a team's style [3, 11, 25, 32, 54, 63, 100, 110, 111, 113, 114]. This approach has many similarities with Markov model-based approaches. For example, the nodes in the graph often denote the players or locations on the pitch (i.e., a commonly used state space in soccer Markov models) and the edges correspond to the transitions between different locations.

In Chapters 5 and 6, we will use a Markov decision process to model the in-game behavior of soccer teams and apply probabilistic verification techniques to extract tactical insights from the learned models.

Chapter 3

Valuing On-the-Ball Actions: A Critical Comparison of xT and VAEP

Action-value frameworks are ubiquitously used in soccer to estimate the effect of a player's chosen actions. These frameworks allow for the simplest form of reasoning, i.e., reasoning about how an action changes the state of the game in one time step. This allows them to provide an estimate of how valuable different in-game actions are with respect to scoring goals. Despite the prevalence and importance of these frameworks, they are rarely, if ever, directly compared either conceptually or empirically. Broadly speaking, there are three styles of approaches to value the actions of players using event stream data:

Count-based approaches. These techniques [98, 99, 109] rate players by (1) assigning a weight to each action type, and (2) calculating a weighting sum of the number of times a player performs each action type (e.g., pass, dribble, cross, tackle) during a match. The weights are typically learned by training a model that correlates these counts with either the match outcome or the number of goals scored (or conceded).

Possession state value approaches. These techniques [41, 97, 125, 133, 162] divide a match into possessions or phases, which are sequences of consecutive on-the-ball actions where the same team possesses the ball. Hence, these models value each action that progresses the ball, typically by seeing how much the action changed the team's chances of producing a

goal scoring attempt. Conceptually, the vast majority of these approaches can be seen as modeling a possession using a Markov model.

Action-based approaches. These approaches go beyond the possession-based ones by trying to value a broader set of actions and by taking the action and game context into account. A recent approach is VAEP [36], which frames the problem as a binary classification task and rates each action by estimating its effect on the short-term probabilities that a team will both score or concede.

In this chapter, we will focus on providing such a comparison between the Expected Threat (xT) framework [133] and the VAEP framework [36]. We select these two frameworks as they are canonical exemplars for the latter two approaches. Because the Markov model-based and action-based approaches both focus on rating each individual action based on properties of the action, they are more closely related in spirit to each other than to the count-based approaches which look at aggregated actions without accounting for any aspect of an action's context.

We highlight the key differences in design choices made by xT and VAEP, which yields different strengths and weaknesses. Qualitatively, we show several illustrative actions where this leads to each formalism producing different valuations for particular actions. Quantitatively, we show that this leads to different rankings of players with xT being slightly more correlated with playmaking whereas VAEP tends to favor shooting. Importantly, both rankings deviate from traditional metrics like goals or assists per 90 minutes, which shows they give novel insights into player performance.

The content of this chapter is based on the following publication and blog post [147, 148].

Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. In *Artificial Intelligence in Team Sports Workshop at AAAI* (2020), pp. 1–8

Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. DTAI Sports Analytics Lab Blog. Available at `https://dtai.cs.kuleuven.be/sports/blog/valuing-on-the-ball-actions-in-soccer:-a-critical-comparison-of-xt-and-vaep`, 2019

An implementation of the Expected Threat framework was made publicly available at `https://github.com/ML-KULeuven/socceraction`.

3.1 Expected Threat

The Expected Threat or xT framework [133] is a possession-based Markov model. This modeling approach implies that soccer games are divided into possessions, which are periods of the game where the same team has control of the ball. Subsequently, each possession can be discretized into a sequence of consecutive ball-progressing actions. The key insight underlying xT and similar models is that players perform these actions with the intention to move the game into a state in which they are more likely to score. These game states directly correspond to the transient states of a Markov model: players transition the game from one state to another by passing or dribbling[1] until absorption (i.e., a goal or possession turnover).

Although the game states can be made arbitrarily complex [125, 162], xT represents each game state S_i by only considering the location of the ball. Therefore, xT overlays a $M \times N$ grid on the pitch in order to divide it into $M \cdot N$ zones. Each zone s is then assigned a value $\mathrm{xT}(s)$ that reflects how threatening teams are at that location, in terms of scoring (Figure 3.1). The value $Q(S_i)$ of game state $S_i = \{a_1, \ldots, a_i\}$ is then simply the value of the zone corresponding to a_i's end location. The Markov model view allows deriving these xT values from historical data by iteratively solving the following equation:

$$\mathrm{xT}(s) = \pi(shot \mid s) \cdot \mathrm{xG}(s) + \pi(move \mid s) \cdot \sum_{s'=1}^{M \times N} P(s, s') \cdot \mathrm{xT}(s'), \qquad (3.1)$$

where $\pi(shot \mid s)$ is the probability that a player will shoot when in zone s, $\mathrm{xG}(s)$ is the probability of a shot from zone s being converted into a goal, $\pi(move \mid s)$ is the probability that a player will move the ball when in zone s, and P is a transition function that defines the probability that the player moves the ball to each of the other zones when in zone s. Intuitively, solving the equation boils down to looking another action ahead with each added iteration. In the first iteration, all $\mathrm{xT}(s)$ values are initialized to zero. After iteration i, $\mathrm{xT}(s)$ then represents the probability of scoring within the next i actions.

[1] Note that event stream data contains various definitions of what constitutes a pass. We regard all open play actions that intentionally move the ball between two players of the same team as passes. This includes regular passes, head passes, long range kicks, crosses, etc. Similarly, the term dribble is quite ambiguous as it can denote trying to move past a player or just carrying the ball from one location to another. We regard both types as a dribble.

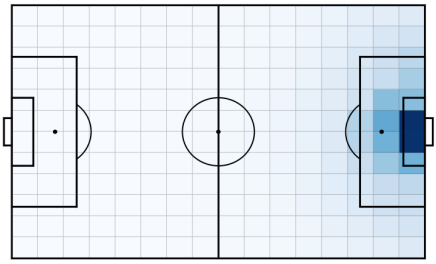

Figure 3.1: A heatmap of the xT values obtained after training the xT model on the training data set until convergence. Darker colors denote higher xT values.

Subsequently, the model values a successful action a_i that moves the ball from zone s to zone s' by computing the difference between the threat value before and after that action:

$$V_{xT}(a_i) = \mathrm{xT}(s') - \mathrm{xT}(s). \tag{3.2}$$

3.2 VAEP

VAEP [36] uses a much more complex game state representation than xT. It considers the three last actions that happened during the game: $S_i = \{a_{i-2},\ a_{i-1},\ a_i\}$. Then each game state is represented using three types of features. The first category of features includes characteristics of the action itself such as its location and type as well as more complex relationships such as the distance and angle to the goal. The second category of features captures the context of the action, such as the current tempo of the game, by comparing the properties of consecutive actions. Examples of this type of feature include the distance covered and time elapsed between consecutive actions. The third category of features captures the current game context by looking at things such as the time remaining in the match and the current score differential.

Like xT (Equation 3.2), VAEP values an action based on how it alters the game state:

$$V_{\mathrm{VAEP}}(a_i) = Q(S_i) - Q(S_{i-1}). \tag{3.3}$$

However, it differs in how it values each game state:

$$Q(S_i) = P^k_{\mathrm{score}}(S_i, t) - P^k_{\mathrm{concede}}(S_i, t). \tag{3.4}$$

where $P^k_{score}(S_i, t)$ and $P^k_{concede}(S_i, t)$ are the probabilities that team t which possesses the ball in state S_i will respectively score or concede in the next k

actions. This valuation arises from the insight that players tend to perform actions not only to increase their team's chance of scoring a goal, but also to decrease their team's chance of conceding a goal in the near future.

Hence, an alternative way to view this is that VAEP estimates the risk-reward trade-off of an action:

$$V_{\text{VAEP}}(a_i) = \Delta P_{\text{score}}(a_i, t) - \Delta P_{\text{concede}}(a_i, t) \tag{3.5}$$

where

$$\Delta P_{\text{score}}(a_i, t) = P^k_{\text{score}}(S_i, t) - P^k_{\text{score}}(S_{i-1}, t) \tag{3.6}$$

$$\Delta P_{\text{concede}}(a_i, t) = P^k_{\text{concede}}(S_i, t) - P^k_{\text{concede}}(S_{i-1}, t). \tag{3.7}$$

Thus, ΔP_{score} is referred to as the "offensive value" of an action, while $\Delta P_{concede}$ can be thought of as the "defensive value" of an action.

VAEP models the scoring and conceding probabilities separately as these effects may be asymmetric in nature and context-dependent. Hence, it trains one gradient boosted tree model to predict each one based on the current game state. For estimating $P^k_{score}(S_i, t)$, each game state is given a positive label ($= 1$) if the team that possesses the ball after action a_i scores a goal in the subsequent k actions. Otherwise, a negative label ($= 0$) is given to the game state. Analogously, for estimating $P^k_{concede}(S_i, t)$, each game state is given a positive label ($= 1$) if the team that possesses the ball after action a_i concedes a goal in the subsequent k actions. If not, a negative label ($= 0$) is given to the game state.

3.3 Conceptual Comparison of xT and VAEP

xT and VAEP are similar approaches in the sense that they both value individual on-the-ball actions of soccer players by evaluating how actions increase or decrease the likeliness of yielding a goal. To do so, both approaches rely on the insight that on-the-ball actions are distinct events that modify the game state. The goal of both frameworks is to measure how valuable an action's resulting change of game state is by computing the differences between the game state values before and after the action.

However, xT and VAEP approach two aspects differently. First, xT uses a very limited game state representation that is purely location-based, while VAEP employs a detailed feature-based representation that captures the action and game context. Second, xT is possession-based, meaning that the xT framework

splits up the game in possessions and estimates the likelihood of any goal occurring within the same possession. VAEP splits up the game in action sequences of a fixed length and looks beyond turnovers. Below, we outline how these design choices impact the generated action-values.

3.3.1 Location-Based vs Feature-Based

VAEP models the game state as an extensive feature-based description of the previous three actions and the game context. xT's game state representation, on the other hand, is purely location-based. It discretizes the pitch into $M \cdot N$ zones by overlaying a grid and encodes the game state as the zone in which the ball is. This enables using simple and elegant dynamic programming approaches to compute each game state's xT value. However, it seriously limits the game dynamics that its values can capture in terms of the action types that can be valued, and the action and game context that is captured.

xT can only value ball-progressing actions Since a game state is fully captured by a zone on the pitch, xT can only value actions that move the ball from one zone to another (i.e., passes, dribbles and crosses). Hence, it ignores defensive actions like tackles and interceptions, as well as valuable offensive actions such as take-ons within the same zone of the pitch.[2] Therefore, xT and related models are often referred to as "ball-progression models" [162].

VAEP captures the action context xT captures only a very limited portion of the context in which actions are performed. Most of the time, the location of an event in space is not sufficient to fully evaluate its potential impact. The probability of scoring in a game state can depend on the type, accuracy, and speed during the previous actions leading up to the current state. For instance, a subsequent shot might be easier when a player is positioned in front of goal via a through ball, compared to when he had to dribble past a couple of defenders first. Therefore, a critical aspect to properly evaluate soccer situations is to have a clear understanding of the ongoing context. VAEP has a much more accurate representation of this context.

VAEP captures the game context VAEP also includes features in its model to capture the game context such as the number of goals scored by each team, the time remaining in the match, and the score difference. This could be

[2] By definition, xT assigns all such actions a value of zero.

valuable as it is known that the chances of scoring vary slightly according to goal difference [38, 120]. Again, xT does not consider these factors.

xT values are interpretable The detailed game state representation used in the VAEP framework has a cost in terms of interpretability. Where each game state in the xT framework is assigned one out of $M \cdot N$ possible values which only depends on the location of the ball, a function approximator (e.g., a gradient boosted tree ensemble) is needed to value game states in the VAEP framework. As such, game state values are derived from complex interactions between a large set of features. Explaining why a particular value is assigned to a specific game state is no longer straightforward in this framework.

3.3.2 Possession-Based vs Window-Based

Both approaches require historical observations of action sequences to estimate the value of a game state. They differ in how they split up the game in those sequences. xT is possession-based. The framework splits up the game in possessions (i.e., sequences of consecutive on-the-ball actions where the same team possesses the ball) and estimates the likelihood of any goal occurring within the same possession. In contrast, VAEP values an action by looking at the probability of a goal being scored within a finite number of actions. This leads to the following two consequences:

VAEP captures the risk involved in an action The xT framework only values an action's offensive contribution, that is, how it changes the team's chance of scoring. In contrast, because VAEP considers what happens after turnovers, it can estimate how an action alters a team's chance of conceding in addition to the action's offensive contribution. Hence, it may better capture the risks associated with taking certain actions. For example, a square pass in the middle of the field enables the other team to quickly launch a counter attack if they would intercept the ball. While VAEP may capture the risks associated with such a pass, xT ignores it.

VAEP can value 'failed' actions accurately Not all losses of possession are equal in soccer. For example, consider a scenario in which a player has the opportunity to clear the ball having no opportunity to reach another teammate. In such a scenario, he can simply kick the ball forward giving the other team the opportunity to recover the ball easily and quickly build up a new attack. Alternatively, he can kick a long ball out of bounds, giving his team the

opportunity to escape the pressing and try to recover the ball by aggressively pressing the subsequent throw-in. The second option is clearly the better one, but cannot be valued by the xT framework.

3.4 Experimental Comparison of xT and VAEP

The conceptual comparison above already gives a good indication of which aspects are tackled differently by both approaches. Next, we perform an experimental comparison to illustrate the impact of these design choices on the action-values and player ratings resulting from both approaches.

3.4.1 Data

The data used for the experiments in this chapter is StatsBomb's event stream data from the English Premier League for the 2017/2018 and 2018/2019 seasons. Both models are trained on the data of the first season. Afterwards, the models are used to rate actions in the second season. Our match event stream data is encoded in the SPADL format [36], which is a language designed for analysis that unifies the representations used by different vendors. This language facilitates analysis by ignoring the optional information (e.g., about weather changes) in the data and by representing all on-the-ball actions using the same fixed set of attributes. Moreover, a publicly available converter is available that translates the data from various providers into this format.[3] We compare both frameworks using their default parameter settings to give insight into the default behavior of both frameworks. For VAEP, we use the previous three actions to represent the game state and use the $k = 10$ subsequent actions to determine the label. The XGBoost [31] algorithm with default parameters was used as the prediction method for the VAEP framework. For xT, the default 16×12 grid was used, and convergence was reached after 6 iterations.

3.4.2 Experimental Comparison of Action-Values

In this section, we compare and contrast how xT and VAEP assess different actions in different contexts. Concretely, we will explore four actions: a (risky) backward pass in the own half, recovering the ball to set up a counter attack, a forward dribble into the opponent's penalty box, and a through ball near the opponent's penalty box.

[3] `https://github.com/ML-KULeuven/socceraction`

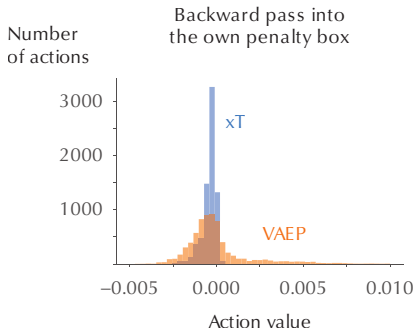

(a) xT does not capture the risk nor reward of backward passes into a team's own penalty box.

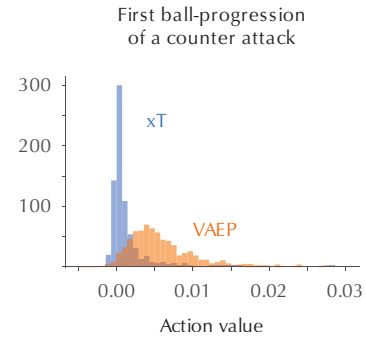

(b) Only VAEP captures initiating a counter attack.

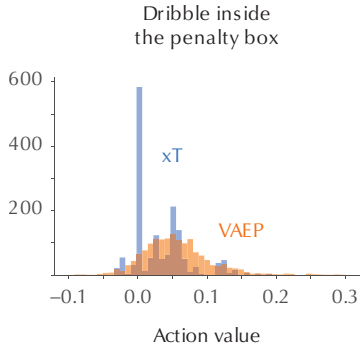

(c) Short dribbles that do not move the ball to a different zone are not valued by xT.

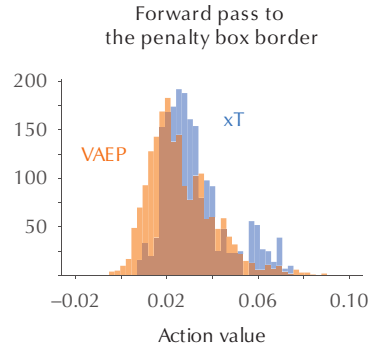

(d) xT values the positional advantage gained by through balls higher than VAEP.

Figure 3.2: Histograms of the VAEP and xT values for a set of actions that are rated differently by both frameworks.

Backward passes into a team's own penalty box

Backward passes have an interesting risk-reward trade-off as they usually open up space (reward), but also move the ball closer to the team's own goal (risk). In particular, a backward pass into your team's penalty box is especially risky as losing the ball in this position may lead to a big scoring chance for the opposing team. These passes happen roughly 19 times per game.

xT neither captures the risk nor reward of these passes, as all zones near a team's own goal are valued close to zero (Figure 3.1). Figure 3.2a shows how VAEP assigns more diverse values to these passes, both positive and negative.

First ball progression of counter attacks

The first ball progression of a counter attack is an action that interrupts the opponent's possession sequence. This action must occur in your own half and be the first action in a sequence of actions that leads to a shot in the next 20 seconds. Such counter attacks occur roughly 2 to 3 times per game.

This ball progression can be a valuable action from both an offensive and a defensive viewpoint. On the one hand, the action interrupts the opponent's attack, thus reducing the odds of conceding a goal. On the other hand, recovering the ball while the opponent is still in an offensive position gives the player's team the opportunity of building a fast counter attack that exploits the opponents' unorganized defensive positioning and thus increasing the odds of scoring a goal. The fact that the ball was only very recently recovered (and how this differs from normal open play) is a contextual clue that can only be leveraged by VAEP's more powerful reasoning on game states. Figure 3.2b illustrates how the distribution of VAEP values of ball recoveries in your own half has a higher mean and variance, whereas the distribution of xT values of ball recoveries in your own half is extremely skewed towards zero.

Forward dribbles inside the penalty box

Next, we consider forward dribbles inside the penalty box that end up in the 23×13 meter rectangle in front of goal. This rectangle corresponds to the rectangle made up of 4×2 cells in front of the goal with the highest xT values. Although our definition of dribbles does not require that they pass a defender, they still considerably raise the odds of scoring a goal for two reasons: (1) the ball got moved a lot closer to the goal, and (2) the player kept control of the ball close to goal. Successful forward dribbles inside the penalty box occur roughly four times per game.

Figure 3.2c shows how xT assigns a value of zero to a major part of these forward dribbles inside the penalty box. Since xT discretizes the pitch into relatively large zones, many short dribbles do not move the ball into a different zone and therefore do not increase the xT value. Yet, these short dribbles may suffice to take on a defender and – when the ball is extremely close to the goal – small differences in location can considerably increase the odds of scoring.

Completed forward passes to the border of the penalty box

Here we explore successful forward passes from the third quarter of the field to the border of the penalty box. More precisely, we consider passes that end within 3 meters of either side of the penalty box border line that is parallel to the goal. Such completed forward passes occur roughly 5 times per game.

A successful forward pass to the border of the penalty box can considerably raise the odds of scoring a goal for two reasons: (1) it moves the ball closer to the goal, and (2) it often bypasses at least one player from the opposing team. Figure 3.2d shows how, on average, xT values the positional advantage gained by a through ball more than VAEP. Due to the detailed game state representation used by VAEP, it is hard to explain why VAEP assigns lower values to these types of actions than xT. Using only the reasoning behind xT, one possible explanation for this is that xT is better than VAEP at capturing the positional advantage. However, it is possible that only using the positions of the actions will overestimate the action-values and more information about the game state is needed to value these actions correctly. Unfortunately, determining the ground truth of these action-values is very difficult, if not impossible. Thus, deciding which method is better suited for valuing these types of actions, or which method makes a better estimation of these action-values is not straightforward.

3.4.3 Experimental Comparison of Player Ratings

The most important application of an action-value framework is summing the action-values of players to construct player ratings. Since spending more time on the pitch offers more opportunities to contribute, player ratings are normalized per 90 minutes of game time [36]. Given a time frame T and player p, a player's rating is computed as

$$\text{rating}(p) = \frac{90}{m} \sum_{a \in A_p^T} V(a), \tag{3.8}$$

where A_p^T is the set of actions the player p performed during time frame T, $V(a_i)$ is the value of an action according to either xT or VAEP, and m is the number of minutes the player played during T.

In this section, we compare and contrast how xT and VAEP rate players. We compare each method on (1) their top player rankings, (2) their correlation to traditional player performance metrics, and (3) their robustness.

Comparison of top-25 player rankings

To some extent, it is possible to subjectively evaluate both rating systems. Tables 3.1 and 3.2 show the top-25 players in the 2018/2019 season of the English Premier League according to xT and VAEP. Soccer fans would agree that the rankings produced by both xT and VAEP feature top players. Yet, there are some major differences between both rankings.

Table 3.1: The top-25 players who played at least 900 minutes in the 2018/2019 English Premier League season with ratings according to the xT framework. Goals per 90 minutes (G/90) and assists per 90 minutes (A/90) are also shown for each player.

R_{xT}	R_{VAEP}	Player	Rating	G/90	A/90
1	1 =	Eden Hazard	0.547	0.49	0.46
2	16 ▼	Adama Traoré	0.490	0.10	0.10
3	24 ▼	Kevin De Bruyne	0.456	0.19	0.19
4	25 ▼	Alex Iwobi	0.441	0.14	0.27
5	6 ▼	Anthony Martial	0.400	0.55	0.11
6	8 ▼	Felipe Anderson	0.399	0.26	0.12
7	106 ▼	Alexis Sánchez	0.395	0.10	0.31
8	2 ▲	Gerard Deulofeu	0.386	0.43	0.21
9	12 ▼	Wilfried Zaha	0.384	0.30	0.15
10	3 ▲	Riyad Mahrez	0.371	0.47	0.27
11	9 ▲	Raheem Sterling	0.340	0.55	0.32
12	67 ▼	Willian	0.338	0.13	0.25
13	30 ▼	Kieran Trippier	0.337	0.04	0.12
14	7 ▲	Mohamed Salah	0.327	0.60	0.22
15	20 ▼	James Milner	0.324	0.25	0.20
16	80 ▼	Nathan Redmond	0.321	0.16	0.11
17	15 ▲	Trent Alexander-Arnold	0.316	0.04	0.44
18	10 ▲	Jonjo Shelvey	0.313	0.10	0.10
19	51 ▼	Benjamin Mendy	0.310	0.00	0.50
20	14 ▲	Ryan Fraser	0.308	0.20	0.40
21	36 ▼	Oleksandr Zinchenko	0.302	0.00	0.23
22	53 ▼	Andrew Robertson	0.298	0.00	0.31
23	21 ▲	David Silva	0.296	0.22	0.30
24	4 ▲	Xherdan Shaqiri	0.289	0.52	0.26
25	110 ▼	Marc Albrighton	0.287	0.11	0.11

A notable omission from the top-25 players according to xT is Manchester City's Sergio Agüero as he is ranked 19th by VAEP. One possible explanation for this is that Agüero is a world class striker and thus probably scores more goals than expected from his shots due to his superior finishing skill. Overperforming on shots is a skill that will be rewarded heftily by the VAEP framework, as this directly influences the scoreline. In summary, the reason why VAEP picks

Table 3.2: The top-25 players who played at least 900 minutes in the 2018/2019 English Premier League season with ratings according to the VAEP framework. Goals per 90 minutes (G/90) and assists per 90 minutes (A/90) are also shown for each player.

R_{VAEP}	R_{xT}	Player	Rating	G/90	A/90
1	1 =	Eden Hazard	0.558	0.49	0.46
2	8 ▼	Gerard Deulofeu	0.535	0.43	0.21
3	10 ▼	Riyad Mahrez	0.525	0.47	0.27
4	24 ▼	Xherdan Shaqiri	0.489	0.52	0.26
5	47 ▼	Son Heung-Min	0.479	0.52	0.26
6	5 ▲	Anthony Martial	0.469	0.55	0.11
7	14 ▼	Mohamed Salah	0.466	0.60	0.22
8	6 ▲	Felipe Anderson	0.466	0.26	0.12
9	11 ▼	Raheem Sterling	0.463	0.55	0.32
10	18 ▼	Jonjo Shelvey	0.442	0.10	0.10
11	210 ▼	Ruben Loftus-Cheek	0.425	0.56	0.19
12	9 ▲	Wilfried Zaha	0.411	0.30	0.15
13	33 ▼	Mesut Özil	0.404	0.26	0.10
14	20 ▼	Ryan Fraser	0.398	0.20	0.40
15	17 ▼	Trent Alexander-Arnold	0.382	0.04	0.44
16	2 ▲	Adama Traoré	0.381	0.10	0.10
17	35 ▼	Dwight McNeil	0.379	0.17	0.28
18	63 ▼	Sadio Mané	0.375	0.64	0.03
19	109 ▼	Sergio Agüero	0.368	0.75	0.29
20	15 ▲	James Milner	0.361	0.25	0.20
21	23 ▼	David Silva	0.351	0.22	0.30
22	55 ▼	Christian Eriksen	0.351	0.26	0.39
23	81 ▼	Pierre-Emerick Aubameyang	0.347	0.72	0.16
24	3 ▲	Kevin De Bruyne	0.345	0.19	0.19
25	4 ▲	Alex Iwobi	0.343	0.14	0.27

up Agüero and xT does not is his average contribution from ball-progressing actions, but superior finishing.

A notable omission from the top-25 players according to VAEP is Manchester United's Alexis Sánchez as he is ranked 7th by xT. One possible reason for this is that it was Sánchez's first full season at Mourinho's defensively organised Manchester United, a playing style he was not used to after four seasons at Arsenal. During this season, his expected goals per 90 minutes (i.e., xG/90) more than halved. Yet, his number of key passes leading to a shot per 90 minutes (i.e., KP/90) stayed roughly the same. These stats indicate that he created roughly the same amount of threat as he did at Arsenal by positioning others in front of goal, but that he had fewer goal-scoring opportunities himself. As xT does not consider shots when ranking players, the fact that Sánchez still managed to complete key passes into high-value zones delivers him a higher xT value. On the other hand, the decrease in shots leads to VAEP valuing him less.

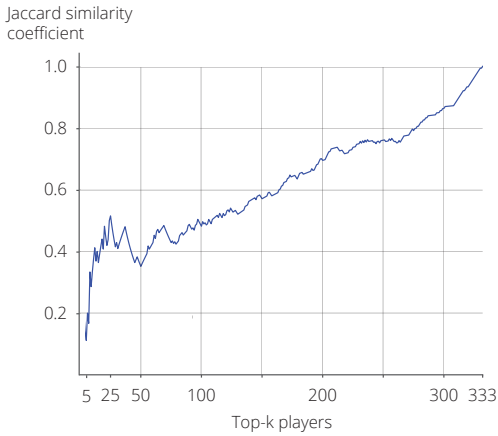

Figure 3.3: The Jaccard similarity coefficient between the top-k rankings created by the xT and VAEP frameworks. The plot shows an increase in similarity for the top-25 players of both models and then a drop when adding the next 25 players before steadily increasing.

In both models, offensive actions have access to higher rewards. It is thus easier for offensive players to get higher ratings than for defensive players. Therefore, the top of the lists mainly contain attacking players, whereas defensive players are ranked lower. For example, top class defender Virgil van Dijk is ranked 81st by VAEP and 142nd by xT.

Finally, Figure 3.3 shows the evolution of the Jaccard similarity coefficient between top-k rankings created by xT and VAEP. This metric essentially measures the similarity between the sets of players in both rankings. The top-25 player rankings of both models have a relatively small similarity coefficient of 0.48, indicating again that xT and VAEP value different qualities. After about the top 25 players, most players get a similar, average rating such that small rating differences can have large effects in the rankings. Hence, the similarity coefficient drops to 0.35 before steadily increasing.

Comparison with traditional performance metrics

Currently, players' offensive contributions are usually quantified by counting goals and assists, as those events directly influence the scoreline. Although they largely fail to account for the circumstances under which the actions were performed, these statistics provide some insights into the performances of

individual soccer players. Therefore, we correlate the ratings of both models to these two baseline metrics, normalized for game time.

For VAEP ($\rho_{G/90} = 0.41$) we obtain a stronger correlation with goals per 90 minutes than xT ($\rho_{G/90} = 0.26$), while assists per 90 minutes is correlated stronger with xT ($\rho_{A/90} = 0.53$) than with VAEP ($\rho_{A/90} = 0.33$). That is because VAEP generally assigns goals high action-values, such that players can boost their VAEP ranking by scoring many goals. In xT, players do not get credit for scoring goals such that the xT-based rankings are biased towards creative players that complete many key passes and dribbles.

Robustness

A good rating system should capture the true quality of all players. Although some fluctuations in performances are possible across games, over the course of a season a few outstanding performances (possibly stemming from a big portion of luck) should not dramatically alter an assessment of a player. To measure which rating system produces the most robust player ratings over the course of a season, we split the data in two random disjoint subsets. Subsequently, we compute each player's average rating separately for both subsets and evaluate the Pearson correlation between both.

Figure 3.4 plots the relation between each player's average rating on the two subsets for xT (Figure 3.4a) and VAEP (Figure 3.4b). As can be observed visually from these figures, the xT framework achieves a much stronger correlation ($\rho = 0.89$) than the VAEP framework ($\rho = 0.25$). The correlation for the VAEP framework improves when only the ball-progressing actions that are valued by the xT framework are included (i.e., passes, dribbles and crosses) and when only the offensive value of actions is considered (Figure 3.4c). In this setting, VAEP achieves a correlation of $\rho = 0.59$.

This ultimately leads to the conclusion that the ratings produced by the xT framework are more robust than the ones from the VAEP framework, even when adjusting for the different actions and risk (i.e., defensive value). These differences can be attributed to two factors. First, because VAEP assigns high action-values to goals, the aggregated ratings can vary significantly based on whether or not these goals are included. Especially for defensive players, a difference in only three goals can double or half their ratings. Second, since players are pretty consistent in what type of actions they perform at which locations [37], it makes sense that a metric purely based on zonal changes gives consistent results. The VAEP ratings add more context and therefore allow more variation.

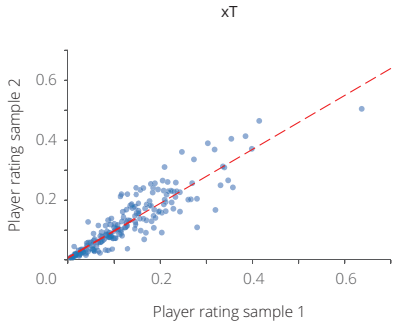

(a) xT ratings for players in both samples.

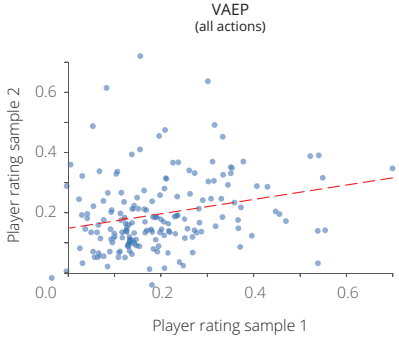

(b) VAEP ratings for players in both samples.

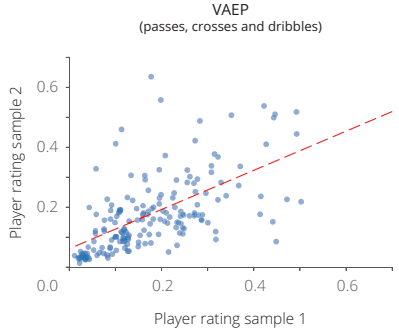

(c) VAEP ratings considering only offensive values for passes, dribbles and crosses in both samples.

Figure 3.4: Scatter plots of the player ratings produced on both samples. The linear least squares regression line is shown in red. The plot for xT displays a high correlation between both ratings. For VAEP, the plots are more random than the one for the xT framework, but a correlation becomes visible when only looking at the offensive values for passes, dribbles and crosses of VAEP.

3.5 Conclusion

xT and VAEP are two prominent approaches for the important task of valuing actions in a soccer match. In this chapter, we performed a critical comparison of these two approaches, conceptually, qualitatively and quantitatively. Key

differences arise in how each approach represents the game state and what actions are valued. These lead to interesting differences such as VAEP better capturing the risk-reward trade-off of actions and xT being more robust. Importantly, both metrics produce rankings that deviate from those produced by considering traditional metrics (goals or assists). Hence, they provide additional insights into player performance.

The current study contains one main limitation, being that it only compares two out of the various existing action-value frameworks. We have decided to compare two of the most popular frameworks which are canonical exemplars for (1) a Markov model-based approach that models entire possessions, and (2) a pure action-based approach that uses a probabilistic classifier to estimate the effects of actions. While there exist many more frameworks, they often differ only in one or a few of the design choices (e.g., game state representation, used learning approach) involved in the discussed frameworks. Therefore, by comparing two representative frameworks of each approach, this study provides a main overview of the differences and implications one can expect. Future work could investigate a more elaborate comparison that includes other variants as well.

Chapter 4

DataDebugging: Enhancing Trust in Action-Value Frameworks by Contextualization

Tree ensembles (e.g., random forests, gradient boosted trees) are very popular as they achieve state-of-the-art performance on many tasks. Specifically, in soccer, tree ensembles form the basis of many approaches for valuing players' actions [36, 119, 136, 140]. While such learned models perform quite well, there are two situations in which the predicted values cannot always be trusted. First, these models are applied to data that is annotated by humans while watching videos of matches. These annotators sometimes make mistakes. When applied to incorrectly annotated data, the model's predictions will be based on incorrect information and will inherently be wrong. Second, the model's predictions should be treated with caution when it is confronted with examples that are highly dissimilar to the examples used during training because it is well known that learned models struggle to extrapolate beyond the training data [71].

This chapter discusses how to automatically detect these potentially abnormal situations when using the VAEP framework discussed in the previous chapter. The approach assesses how dissimilar the test examples are to examples the model was trained on. The flagged abnormal examples can then be corrected or removed before moving on to the downstream use cases. We evaluate our approach on two public event stream data sets and show that it can detect

annotation errors and rare situations (e.g., extraordinary shots).

The content of this chapter is based on the following publication [144].

Van Roy, M., and Davis, J. DataDebugging: Enhancing Trust in Soccer Action-Value Models by Contextualization. In *Proceedings of 13th World Congress of Performance Analysis of Sport and 13th International Symposium on Computer Science in Sport* (2023), vol. 1448 of *Advances in Intelligent Systems and Computing*, Springer Nature Switzerland, pp. 193–196

4.1 Detecting Abnormal Examples in Action-Value Frameworks

Detecting if an unseen example is abnormal requires two things: (1) a trained model underlying an action-value framework, which can be queried, and (2) an approach to score the given example based on how abnormal it is with respect to the model. Next, we describe both parts.

4.1.1 Action-Value Framework

As action-value framework, we use the state-of-the-art VAEP framework [36], as explained in Chapter 3. VAEP is a simple tree-based action-value framework that, given a sequence of the three latest actions, estimates the probability that the team currently possessing the ball will score a goal in the next 10 actions. We use the default features of VAEP to describe a sequence, which contain e.g., the start and end location of the actions, their action type, and the body part used to execute them. To obtain an optimized learned model that can be used in practice, we use a 5-fold cross validation grid search during training to optimize the depth (i.e., 3, 5, 8) and the number of trees included in the underlying tree ensemble model (i.e., 50, 75, 100).

4.1.2 Detecting Abnormal Examples in Tree Ensembles

To detect abnormal examples in the learned tree ensemble model, we use Devos et al.'s [44] OC-Score approach that is designed to identify dissimilar test examples when using tree ensembles. Instead of using an example's feature representation, it encodes each test example as the ordered set of leaf nodes

that are reached when executing the ensemble, which is called the example's output configuration (OC). Even though two examples may seem similar when comparing their feature representations, small changes in these features can yield very different output configurations and hence final predictions. By comparing the output configuration of a test example to those of the examples the ensemble was trained on, the approach can detect dissimilar test examples for which the ensemble's prediction may not be trusted:

$$OC{-}score(x) = min\{Hamming(OC(x), OC(x')) \,|\, x' \in \mathcal{R}\} \tag{4.1}$$

where $\mathcal{R}$ is a reference set containing all correctly classified training examples with the same label that is predicted for the test example x (i.e., we have a reference set of correctly classified training examples for each class). For each data set, we apply this method and flag a test example as abnormal when its OC-score is greater than the corresponding reference set's average OC-score plus one standard deviation.

4.2 Experiments

In this section, we illustrate our approach on two event stream data sets and show how it can detect abnormal examples in both. Next, we perform a complementary analysis on the overall space of output configurations of the test examples with `TRACE` [67].

4.2.1 Data

We use publicly available event stream data from the UEFA EURO 2020 and the FA Women's Super League (FAWSL) 2020/21.[1] Each data set is split into a train set used to learn the tree ensemble models, and a test set in which abnormal examples will be detected. For each data set, we use an 80-20 game-wise split. This yields approximately 90,000 and 200,000 train examples, and 20,000 and 50,000 test examples, for the respective data sets.

4.2.2 Detecting Abnormal Examples

To validate whether our approach can detect both annotation errors and rare examples, we compare the flagged examples with their corresponding video.

[1] `https://github.com/statsbomb/open-data`

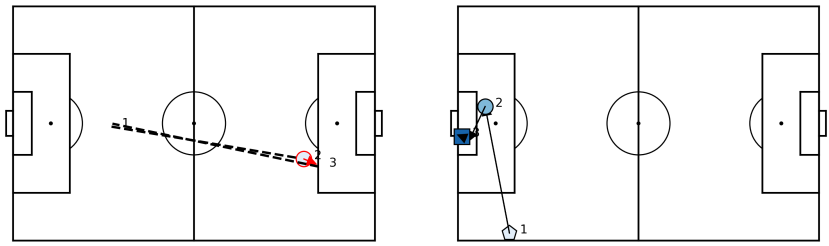

(a) Error in both dribbles (b) Error in first action

Figure 4.1: Two identified annotation errors in the FAWSL: (a) the start and end location of the two dribbles are incorrectly annotated, and (b) the y-coordinate of the first action should be near the penalty box border instead of near the sideline.

However, it is not straightforward to obtain videos of each example. Therefore, for each data set, we selected 10 examples out of the top 100 flagged abnormal test examples for which videos could be obtained. Based on the videos, two out of the 20 flagged sequences were found to be annotation errors, 11 could be considered infrequent ones, and 7 were found to be rather normal sequences.

Figure 4.1 shows the two annotation errors identified by the approach. It identifies the incorrectly annotated dribbles of Faye Bryson and Caroline Weir during the Bristol - Manchester City match (Figure 4.1a), and the incorrectly annotated cross by Ellie Brazil during the Brighton & Hove Albion - Bristol City match (Figure 4.1b).

Figure 4.2 shows four representative examples of rare action sequences identified by the approach. For example, it identifies infrequent sequences such as the well-executed goal-sequence starting with a long high pass by Italy against Switzerland during the EURO 2020 (Figure 4.2a), and it picks up on exceptional goals that rarely occur, such as Patrick Schick's goal from the halfway line during the EURO 2020 (Figure 4.2b).

Finally, Figure 4.3 shows four representative examples of rather normal action sequences falsely identified by the approach. For example, it identifies the sequence involving the corner-kick of Katie Zelem during the Manchester United - Aston Villa match as abnormal (Figure 4.3a). However, from a sports perspective, this situation is not uncommon. This example might have been flagged due to e.g., corner-kicks being underrepresented or the gap in time and location between the corner-kick and the action prior to it.

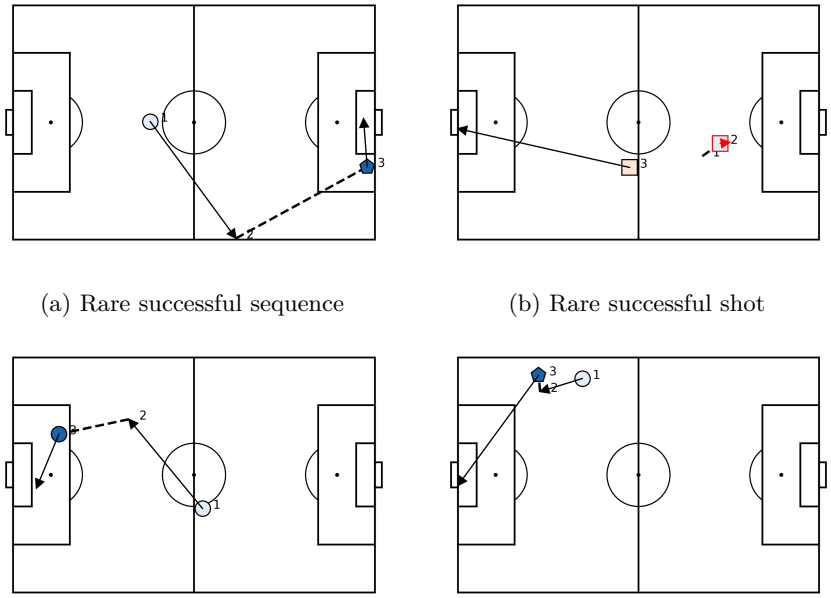

(a) Rare successful sequence

(b) Rare successful shot

(c) Rare successful sequence

(d) Rare successful shot

Figure 4.2: Four representative examples of rare sequences in the EURO 2020 (a, b) and FAWSL (c, d). Figures (a) and (c) are two rare sequences with long successful passes and dribbles, either along the sideline or through the middle of the defense. Figures (b) and (d) are two rare successful shots, one from the halfway line and one from a tight angle outside the penalty box.

4.2.3 Analysis Using TRACE

As a complementary analysis, we perform a visual inspection of the space of all output configurations of the test examples. The main idea behind the OC-Score algorithm is that the output configurations of abnormal examples (i.e., the ordered set of leaf nodes that are reached in the ensemble) are different from those covered by the normal examples of the same class in the training data. Thus, our hypothesis is that, when visualizing the space of all output configurations of the test examples of a particular class,[2] the normal test examples will be grouped together, the abnormal test examples will be grouped together, and both groups will be mostly disjoint from each other.

[2]In this analysis, we visualize the class "no goal" as the predictions for this class cover 99% of the test examples.

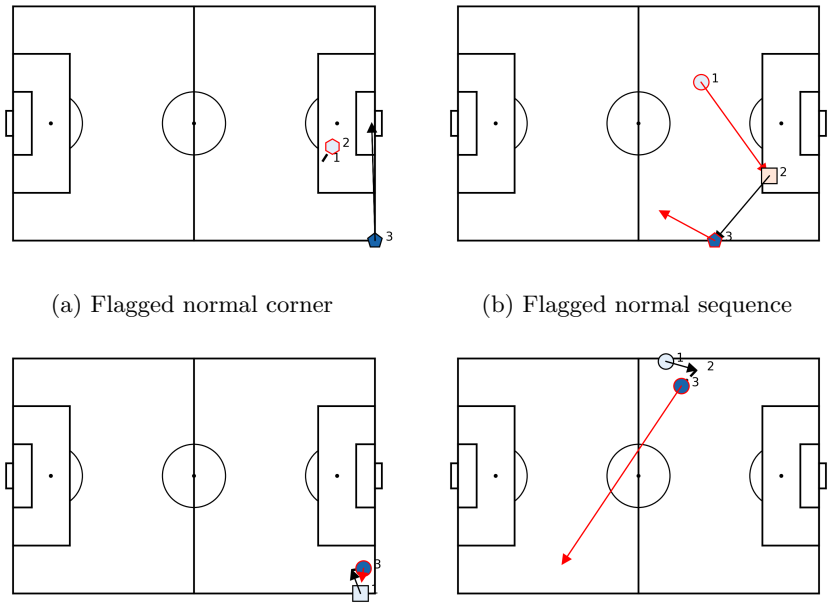

(a) Flagged normal corner

(b) Flagged normal sequence

(c) Flagged play near corner

(d) Flagged failed pass

Figure 4.3: Four representative examples of rather normal sequences in the FAWSL that are falsely identified as abnormal.

We will use the interactive tool `TRACE` [67] to analyze the overall inherent structure within the space of output configurations. This tool allows for visualizing and interactively exploring and comparing different embeddings that can be created from a given data set. In our case, our data sets consist of the output configurations of the test examples for each considered event stream data set. To create 2-dimensional embeddings that can be visualized with `TRACE`, we use t-Distributed Stochastic Neighbor Embedding (t-SNE) [139] and Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) [101]. For both dimension reduction techniques, we use the Manhattan distance as distance metric, given that the OC-Score algorithm also relies on this metric. All other parameters are set to their default values and we use three different random seeds for each dimension reduction technique to investigate the effect of the randomness that is present in these dimension reduction techniques.

Figures 4.4 and 4.5 each show one of the resulting embeddings for both dimension reduction techniques and data sets. As can be seen in Figure 4.4, there is one

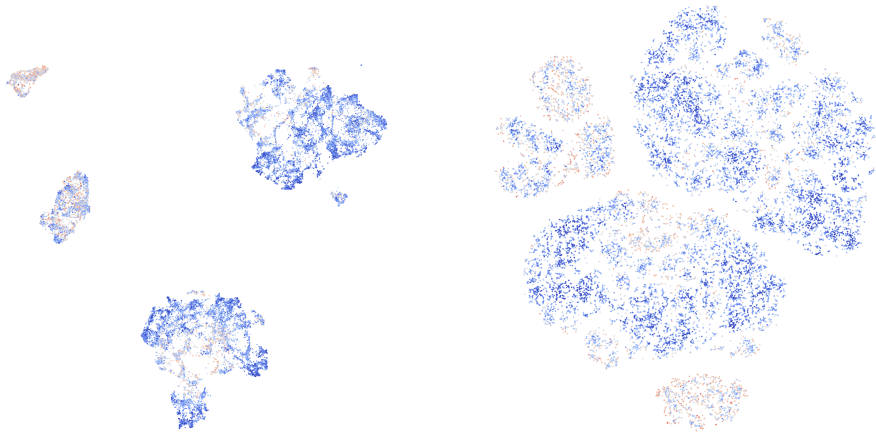

(a) UMAP (b) t-SNE

Figure 4.4: Visualization of a 2-dimensional UMAP and t-SNE embedding of the output configurations of "no goal" predicted test examples in the EURO 2020. Red means that the example is more abnormal according to its OC-Score while blue means it is more normal.

prominent cluster containing more abnormal than normal examples (i.e., the top left cluster in Figure 4.4a and the bottom cluster in Figure 4.4b). In Figure 4.5 using the FAWSL data, this is less pronounced, especially using t-SNE. However, a slight clustering of abnormal examples is still visible in the left most cluster of Figure 4.5a and in the top right of Figure 4.5b.

Besides visual analysis, `TRACE` allows for computing quality metrics over multiple embeddings. This allows us to investigate how representative the embeddings are of the underlying data. Table 4.1 shows the random triplet accuracy metric for both dimension reduction techniques and all three random seeds. This metric quantifies the ratio of random triplets (i, j, k) where the relative order of j and k with respect to i in the high-dimensional space (i.e., the output configurations) is preserved in the embedding. The results show that the order of the examples is roughly maintained within the different embeddings with no large differences between the different seeds. In the remainder of this analysis, we will continue with one embedding of each type for each data set.

Next, the tool allows for selecting certain examples and tracking them throughout multiple embeddings. For example, in Figure 4.6a the top left abnormal cluster is selected in the UMAP embedding, which can be seen to map entirely to the bottom abnormal cluster in the t-SNE embedding in Figure 4.6b.

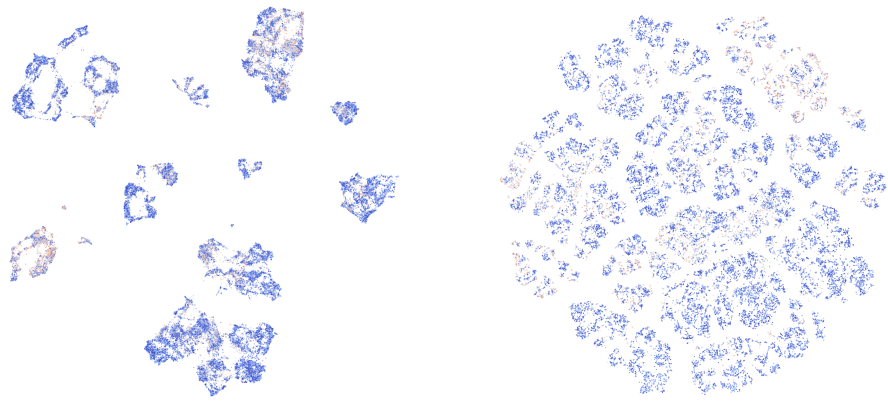

(a) UMAP (b) t-SNE

Figure 4.5: Visualization of a 2-dimensional UMAP and t-SNE embedding of the output configurations of "no goal" predicted test examples in the FAWSL. Red means that the example is more abnormal according to its OC-Score while blue means it is more normal.

Table 4.1: Random triplet accuracy for both dimension reduction techniques and all three random seeds on both data sets.

Embedding	**EURO 2020**	**FAWSL**
t-SNE1	0.686	0.691
t-SNE2	0.704	0.704
t-SNE3	0.680	0.704
UMAP1	0.738	0.697
UMAP2	0.717	0.687
UMAP3	0.729	0.690

Thus, the abnormal examples stay clustered throughout multiple embeddings, indicating that there is an inherent underlying structure present in the output configurations. Figure 4.6c shows the distribution of OC-Score of the selected examples compared to the distribution of the other examples. As can be seen from the shift in the distribution, the selected cluster is more abnormal.

Finally, we can select a number of normal (or abnormal) examples and let `TRACE` visualize the 50 closest high-dimensional neighbors for each selected example. High-dimensional neighbors are in this case the closest examples according to their distance in their output configurations. This allows us to investigate whether the examples that are similar in their output configurations are also similar in the lower dimensional embeddings. Figure 4.7a visualizes a



(a) UMAP

(b) t-SNE

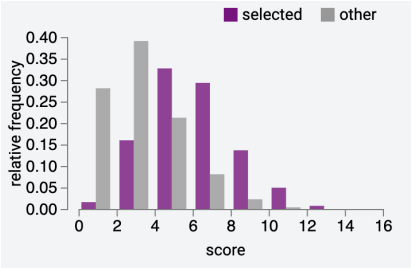

(c) Distribution of OC-Scores

Figure 4.6: The selected abnormal cluster in the top left of the UMAP embedding can be tracked and mapped with `TRACE` to the bottom cluster in the t-SNE embedding, indicating that there is an underlying structure present in the output configurations of "no goal" predicted test examples in the EURO 2020. The shift in the distribution of OC-Scores of the selected examples also indicates that this cluster is indeed more abnormal.



(a) Normal - FAWSL (b) Abnormal - EURO 2020

Figure 4.7: UMAP embeddings of "no goal" predicted test examples in the FAWSL and EURO 2020 with a selected number of (a) normal examples and (b) abnormal examples and their high-dimensional neighbors.

selected cluster of normal examples and their high-dimensional neighbors for the FAWSL. Figure 4.7b visualizes the same experiment using a selected cluster of abnormal examples in the EURO 2020. As can be seen, the highlighted set of high-dimensional neighbors for the normal examples mainly contains normal examples, and vice versa for the abnormal examples. This confirms the hypothesis that normal examples are similar to each other but dissimilar from abnormal examples, and similarly for abnormal examples. Additionally, for the normal examples we can see that the high-dimensional neighbors can be found both close to the selected normal examples but also more dispersed throughout the embedding. It is likely that the various different normal examples do not differ much from each other with respect to their output configuration, yet the dimension reduction algorithm attempts to cluster them based on these small differences. For the abnormal examples, we can observe that the abnormal examples in other parts of the embedding are not highlighted as high-dimensional neighbors. Even though the different examples are labeled as abnormal by the OC-Score algorithm, the way in which they are abnormal is probably highly different (e.g. an abnormal passing sequence vs. an abnormal shot).

4.3 Conclusion

This chapter introduced an automatic approach to detect potentially abnormal situations in tree-based action-value frameworks. By inspecting the output configuration of such tree-based action-value framework, the method can detect examples for which the underlying model's predictions may not be trustworthy. On the one hand, the approach can identify annotation errors, for which the model's predictions are inherently wrong. This could help data providers automatically verify their data before release. On the other hand, it can identify infrequent examples that are dissimilar to the training set. The model's value should be treated with caution in these situations. By visually and interactively inspecting the space of output configurations of the test examples with `TRACE`, we could confirm that using the output configuration of examples is indeed an appropriate approach to identify abnormal and normal examples.

Currently, the outlined approach contains two main limitations. First, in our experiments, we were only able to validate the flagged examples against video for a limited number of game situations, due to the public availability of the data and video. However, to assess the proportion of (correctly) flagged abnormal vs. (falsely) identified normal examples, and thus to determine how many examples of each would be rejected at particular thresholds of the OC-Score, the approach should be evaluated on all actions of multiple games for which one could obtain videos. Second, one could wonder how the OC-Score approach, which reasons about the predictions of a model, compares against an outlier detection approach, which reasons about the features of the examples themselves. Our hypothesis is that an outlier detection approach will (1) be able to pick up examples for which there is a clear deviation from the training data (e.g., Figure 4.2b) but will fail to pick up examples that only slightly deviate in certain features (e.g., Figure 4.1b), and (2) that an outlier detection approach might be more computationally expensive, given the rather straightforward distance-based approach of OC-Score. Future work could investigate a comparison of the proposed approach with various outlier detection approaches.

Chapter 5

A Markov Framework for Learning and Reasoning About Tactics

Optimizing tactics is a fundamental element of dynamic and complex team sports such as soccer, American football, and basketball. If one wants to gain insights into the effectiveness of specific tactics, or if alternative options need to be compared with each other, a data-driven approach can uncover insights that are not available from qualitative analysis. However, this is not straightforward, and a combination of learning and reasoning is warranted. 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 sequential decisions, their effects, and what would happen if the team were to behave differently.

This chapter presents a novel framework that can support a professional soccer coach's tactical planning by using the combination of learning and reasoning. The learning component entails learning the model underlying a team-specific opponent-agnostic Markov decision process (MDP) from data about their matches. By averaging over, e.g., one or two seasons of data, the MDP captures a team's general behavior. The reasoning component operates on two different levels. First, we demonstrate a novel application of probabilistic model checking [83] 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 tactics 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.[1] 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 tactics 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.

The content of this chapter is based on the following publication [151].

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. A Markov Framework for Learning and Reasoning About Strategies in Professional Soccer. *Journal of Artificial Intelligence Research 77* (2023), 517–562

An interactive tool is provided on the following website: `https://dtai.cs.kuleuven.be/sports/ssac21/`.

[1] `https://x.com/devinpleuler/status/1226919308762193920`

5.1 Modeling Offensive Behavior as a Markov Model

To reason about a team's observed behavior, we model each team's offensive behavior using a Markov decision process (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, or (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 2.3.2).

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

5.1.1 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 and will be discussed later.

5.1.2 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 [125, 133, 162].

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.

5.1.3 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:

- $P(s, move_to(s'), s')$ is the probability of successfully moving to state $s' \in \mathcal{L}$ from state s;
- $P(s, move_to(s'), lost_possession)$ is the probability of unsuccessfully moving to state $s' \in \mathcal{L}$ from state s, this is equal to $1-P(s, move_to(s'), s')$;
- $P(s, shoot, goal)$ is the probability of scoring a goal from state s, and this can be viewed as a location-based xG value;
- $P(s, shoot, no_goal)$ is the probability of failing to score a goal when shooting from state s, this is equal to $1 - P(s, shoot, goal)$;
- $P(s, a, s') = 0$ in all other cases.

5.1.4 Policy

The policy π defines the probability distribution over actions for each state: $\pi(a \mid s) = Pr[A = a \mid S = s]$. 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 [130]. 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 tactics. Furthermore, by adapting the policy to certain counterfactual scenarios and estimating the effects of these changes, further insights into what would happen if different tactics were to be employed can be provided to practitioners.

5.1.5 Reward Function, Discount Factor and Value Function

The reward function $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}$ 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., $R(s, shoot, goal) = 1$. All other actions do not receive a reward.

The discount factor γ 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.

When learning the value function of a team, we learn it for a fixed policy and use the above-mentioned reward function and discount factor. Thus, the MDP becomes an MRP. Our setting roughly corresponds to a passive model-based reinforcement learning setting, albeit one where the reward is known.

5.2 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. 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 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.

5.2.1 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. A movement action can be a pass, a cross, or a dribble. Each has characterizing dynamics, thus 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 six 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.[2] These features are commonly used in sports analytics to describe the context of actions [27, 36, 116, 119]. Appendix A 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

[2]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.

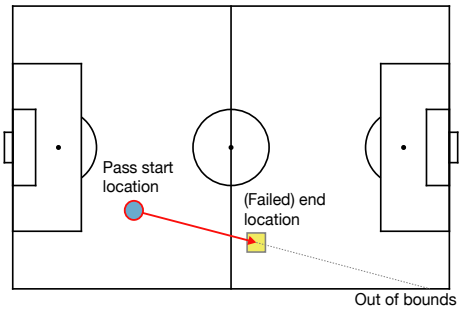

Figure 5.1: Illustration of an unsuccessful pass, starting at the circle and unsuccessfully ending at the square. The intended end location most likely lies along the line that starts from the pass' start location, passes through its failed end location, and ends when it goes out of bounds. If the pass was intercepted, the intended end location most likely lies on the dashed part of the line.

on this 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 5.1 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.[3] 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.[4] 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.

[3]See `https://www.statsperform.com/opta-event-definitions/`

[4]In reality, 86.5% of all crosses in the data set are "chipped into the air".

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.

5.2.2 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 $O_{t,s,a}$. A Bernoulli distribution naturally models the binary outcome of an action (i.e., success or fail):

$$O_{t,s,a} \sim \text{Bernoulli}(p_{t,s,a}) \tag{5.1}$$

$$p_{t,s,a} = \text{invlogit}(\gamma_{t,s,a}) \tag{5.2}$$

$$\gamma_{t,s,a} \sim \mathcal{N}(\mu_{s,a}, \sigma^2_{t,s,a}) \tag{5.3}$$

$$\sigma^2_{t,s,a} \sim \text{Half-Normal}(5.0) \tag{5.4}$$

Here, *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 $\sigma^2_{t,s,a}$. The normal distribution around the overall prior mean ensures that teams can be more or less successful that the "typical team" at action a in state s. The team-dependent variance is half-normally distributed, to

ensure only positive values for the variance, 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). A categorical distribution naturally models the different possible choices for actions in a state:

$$A_{t,s} \sim \text{Categorical}(\vec{p}_{t,s}) \tag{5.5}$$

$$\vec{p}_{t,s} = \text{softmax}(\vec{\lambda}_{t,s}) \tag{5.6}$$

$$\lambda_{t,s,a} \sim \mathcal{N}(\alpha_{s,a}, \nu^2_{t,s,a}) \tag{5.7}$$

$$\nu^2_{t,s,a} \sim \text{Half-Normal}(5.0) \tag{5.8}$$

Here, *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 $\nu^2_{t,s,a}$. The normal distribution around the overall prior mean ensures that teams can choose action a in state s more or less often that the "typical team" does. The team-dependent variance is half-normally distributed, to ensure only positive values for the variance, 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.

5.3 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 effectiveness of tactical 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 tactics 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 tactics in soccer.

5.3.1 Reasoning About the Effect of Offensive Tactics

Ultimately, a team's offensive tactics revolve 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 $(s_1, \ldots, s_{n-1})$ where $s_1, \ldots, s_{n-1} \in \mathcal{L}$, the probability of yielding a goal immediately after this sequence can be formulated as:

$$\mathtt{P}_{s=?}[\mathtt{X}\,(s_1 \wedge (\ldots \wedge (\mathtt{X}\,(s_{n-1} \wedge (\mathtt{X}\,\mathit{goal})))))] \tag{5.9}$$

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.

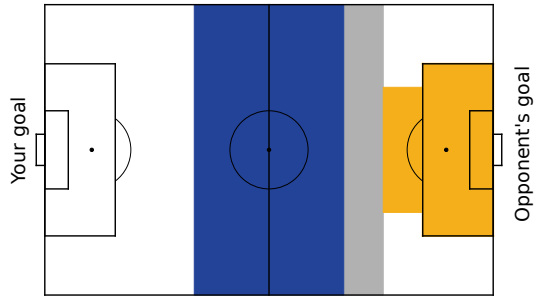

Figure 5.2: Illustration of three interesting areas used to optimize the defensive game plan. The region where most shots are taken from is shaded in yellow, the final third entry region is shaded in gray, and the middle third of the pitch is shaded in blue.

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

$$\mathtt{P}_{s=?}[\mathtt{X}\,(L_1 \wedge (\mathtt{F}\,(L_2 \wedge (\mathtt{X}\,goal))))] \tag{5.10}$$

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.

5.3.2 Reasoning About the Effect of Defensive Tactics

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 the 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.

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 5.2, 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:

$$\mathtt{P}_{s=?}[\mathtt{F}\ (\mathit{shot_locations}\ \mathtt{U}\ (\mathit{goal} \vee \mathit{no_goal}))] \tag{5.11}$$

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 a defending team could prevent an opponent from ever entering $s' \in \mathit{non_shot}$, where $\mathit{non_shot} = \mathcal{L} \setminus \mathit{shot_locations}$. 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:

$$\mathtt{P}_{s=?}[\neg s'\ \mathtt{U}\ (\mathit{shot_locations}\ \mathtt{U}\ (\mathit{goal} \vee \mathit{no_goal}))] \tag{5.12}$$

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:

$$1 - \frac{\sum_{s \in \mathit{non_shot}} \mathtt{P}_{s=?}[\neg s'\ \mathtt{U}\ (\mathit{shot_locations}\ \mathtt{U}\ (\mathit{goal} \vee \mathit{no_goal}))]}{\sum_{s \in \mathit{non_shot}} \mathtt{P}_{s=?}[\mathtt{F}\ (\mathit{shot_locations}\ \mathtt{U}\ (\mathit{goal} \vee \mathit{no_goal}))]} \tag{5.13}$$

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:

$$\mathtt{P}_{s=?}[\mathtt{F}\ \mathit{shot_locations}] \tag{5.14}$$

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:

$$\mathtt{P}_{s=?}[\neg s'\ \mathtt{U}\ \mathit{shot_locations}] \tag{5.15}$$

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':

$$1 - \frac{\sum_{s \in \mathit{non_shot}} \mathtt{P}_{s=?}[\neg s'\ \mathtt{U}\ \mathit{shot_locations}]}{\sum_{s \in \mathit{non_shot}} \mathtt{P}_{s=?}[\mathtt{F}\ \mathit{shot_locations}]} \tag{5.16}$$

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.

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 tactics, the defending team may be able to make it more challenging to get the ball into dangerous positions.

Figure 5.2 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 $s \in \mathit{region}$ by at least x percentage points. Formally, this can be framed as the following query:

$$\forall s \in \mathit{region} : (\mathtt{P}_{s=?}[\mathtt{F}\ \mathit{goal}] - \mathtt{P}_{s=?}[\neg \mathit{area}\ \mathtt{U}\ \mathit{goal}]) \geq x \tag{5.17}$$

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:

$$area = \{s'' \in \mathcal{L} \setminus region \mid \mathtt{P}_{s''=?}[\mathtt{F}\ s'] \geq b\} \tag{5.18}$$

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.

5.4 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 tactics (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.

5.4.1 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.

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 A.

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:

$$\pi'(move_to(s') \mid s) = (1 + x) * \pi(move_to(s') \mid s) \tag{5.19}$$

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 [129, 130], 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:

$$\pi'(move_to(s'') \mid s) = \frac{\pi(move_to(s'') \mid s)}{total_other} * (total_other - x * total_area) \tag{5.20}$$

with *total_other* and *total_area* defined as:

$$total_other = \sum_{s'' \in \mathcal{L} \setminus area} \pi(move_to(s'') \mid s) \tag{5.21}$$

$$total_area = \sum_{s' \in area} \pi(move_to(s') \mid s) \tag{5.22}$$

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:

$$\pi'(move_to(s'') \mid s) = \frac{\pi(move_to(s'') \mid s)}{total} * (total - x * \pi(move_to(s') \mid s)) \tag{5.23}$$

with *total* defined as:

$$total = \sum_{s'' \in s^h} \pi(move_to(s'') \mid s) \tag{5.24}$$

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 $s'' \in s^l \subset \mathcal{L} \setminus area$ 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.

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 π' compared to their original policy π:

$$1 - \frac{\sum_{s \in \mathcal{L}} V_{\pi'}(s)}{\sum_{s \in \mathcal{L}} V_{\pi}(s)} \tag{5.25}$$

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.

5.4.2 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.

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:

$$\pi'(shoot \mid s) = (1 + x) * \pi(shoot \mid s) \tag{5.26}$$

$$\forall s' \in \mathcal{L} : \pi'(move_to(s') \mid s) = \frac{\pi(move_to(s') \mid s)}{total} * (total - x * \pi(shoot \mid s)) \tag{5.27}$$

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

$$\pi'(shoot \mid s) = (1 - x) * \pi(shoot \mid s) \tag{5.28}$$

$$\forall s' \in \mathcal{L} : \pi'(move_to(s') \mid s) = \frac{\pi(move_to(s') \mid s)}{total} * (total + x * \pi(shoot \mid s)) \tag{5.29}$$

with in both cases

$$total = \sum_{s' \in \mathcal{L}} \pi(move_to(s') \mid s) \tag{5.30}$$

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:

$$E[goals \mid \pi] = \sum_{s \in \mathcal{L}} E[shots\, in\, s \mid \pi] * P(s, shoot, goal) \tag{5.31}$$

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

$$E[shots\, in\, s \mid \pi] = \pi(shoot \mid s) * E[visits\, to\, s \mid \pi] \tag{5.32}$$

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:

$$N = (I - Q)^{-1} \tag{5.33}$$

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 $Q_{ij} = P(i, move_to(j), j) * \pi(move_to(j) \mid i)$ and π can either be the observed (to estimate the current expected number of goals) or altered policy. Appendix A provides an analysis on the accuracy of the outlined approach when using the observed policy.

Estimating the success probability of shots. Finally, estimating $P(s, shoot, goal)$ (i.e., the location-based xG value for any shot in state s) for the observed policy π 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 [56, 129, 130].

It seems reasonable to assume that such a relationship would also exist for soccer. Therefore, we adjust for this trade-off when computing $P(s, shoot, goal)$ when an altered policy π' is in place. In our model, $P(s, shoot, goal)$ is the same for *any* shot taken in state s. 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.[5] For a detailed description of the model, we refer to Robberechts and Davis [119]. We then use the distribution of these xG values to derive an adjustment to the estimate of $P'(s, shoot, goal)$ for an altered policy π'. 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, $E[shots\,in\,s \mid \pi'] - E[shots\,in\,s \mid \pi]$ more shots will be taken than in the original policy. We only modify the xG value for the additional shots taken in π' by awarding them an xG of:

$$P(s, shoot, goal) - (\mu_s - \mu_s^{low}) \tag{5.34}$$

where μ_s is the average computed xG of all shots occurring in state s and μ_s^{low} is the average computed xG of the below-average shots occurring in state s. We illustrate this for one state in Figure 5.3.

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:

$$P(s, shoot, goal) + (\mu_s^{high} - \mu_s) \tag{5.35}$$

[5] `https://github.com/ML-KULeuven/soccer_xg`

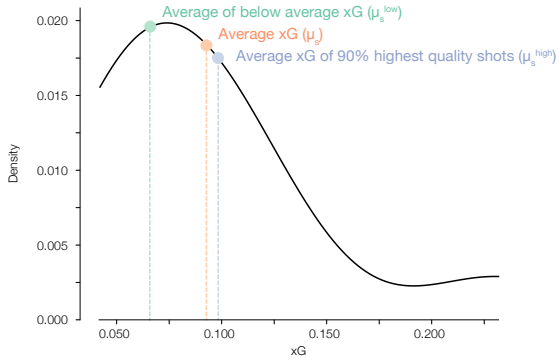

Figure 5.3: Possible expected goals (xG) distribution in one state. Orange indicates the average xG value (μ_s), green indicates the average of the below-average xG values (μ_s^{low}), and purple indicates the average xG value of the shots with an xG value in the top 90% of all shots in the considered zone (μ_s^{high}).

where μ_s is again the average computed xG of all shots occurring in state s and μ_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 5.3.

5.5 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 5.1: 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

5.5.1 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 5.1 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 a subset of 20 teams, ensuring that we both include teams that are at the top, middle, and towards the bottom of the league table, for each of the three leagues. We include 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).

5.5.2 Transient State Spaces

We consider two different sets of transient states (Figure 5.4), which are tailored to the use cases.

An offensive grid. This grid divides the offensive half of the field into zones of 3m $\times$ 3m and captures the defensive half in one state (Figure 5.4a). 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 one expects the chances of scoring to be higher and more coarse-grained where one expects the chances of

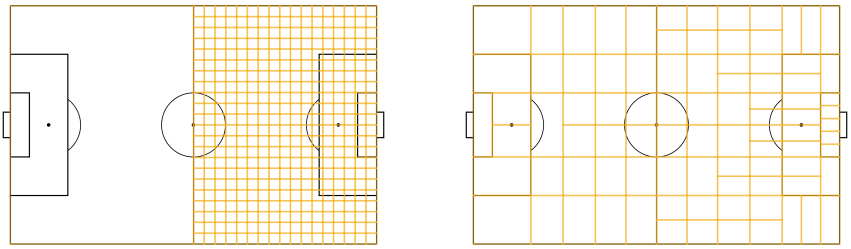

(a) Offensive grid (b) Custom grid

Figure 5.4: 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 3m × 3m. This state space is more suited for analyzing shot policies 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 policies.

scoring to be lower. These states also capture the differences between locations in the defensive half (Figure 5.4b). The size of the states ensures that there is sufficient data available in each state to analyze the different movement tactics and their effects. Additionally, the states also correspond to zones that are frequently used by practitioners when discussing movement tactics (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.

5.5.3 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 5.2.1) are trained using XGBoost [31].[6] 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 contains more information on the experimental setup. After training, the intended end

[6] See `https://xgboost.ai/`. We used XGBoost v1.5.

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 [128] is used to model each of the Bayesian models (Section 5.2.2).[7] 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 4,000 times and the average of these samples is used to compute the final probabilities. Appendix A contains more information on the experimental setup.

5.5.4 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 5.2 reports the area under the ROC curve (AUROC) and the original Brier score [23] 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] See `https://docs.pymc.io/`. We used PyMC3 v3.11.

Table 5.2: AUROC and Brier scores for the end location prediction models averaged (± 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 (see Section 5.2.1). 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 (± 0.00)	0.73 (± 0.04)
		XGBoost	0.73 (± 0.04)	0.57 (± 0.04)
		DK	**0.99 (± 0.00)**	**0.51 (± 0.03)**
	Cross	Baseline	0.50 (± 0.00)	0.94 (± 0.01)
		XGBoost	0.52 (± 0.03)	0.97 (± 0.04)
		DK	**0.92 (± 0.02)**	**0.78 (± 0.07)**
CUSTOM	Pass	Baseline	0.50 (± 0.00)	0.98 (± 0.00)
		XGBoost	0.93 (± 0.01)	0.80 (± 0.02)
		DK	**0.99 (± 0.00)**	**0.52 (± 0.02)**
	Cross	Baseline	0.50 (± 0.00)	0.87 (± 0.02)
		XGBoost	0.53 (± 0.03)	0.89 (± 0.03)
		DK	**0.96 (± 0.01)**	**0.55 (± 0.05)**

5.5.5 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).[8] As training the models can be quite time-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 5.3 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

[8] See `https://arviz-devs.github.io/arviz/`. We used the PSIS-LOO-CV implementation from Arviz v2.0.

Table 5.3: 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 (± 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 (± 0.73)	-27.1 (± 0.13)
	Hierarchical	**-161.4 (± 0.68)**	**-26.2 (± 0.28)**
CUSTOM	Unpooled	-159.1 (± 0.29)	-19.4 (± 0.14)
	Hierarchical	**-157.5 (± 0.29)**	**-18.9 (± 0.18)**

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.

5.5.6 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 (± 1 std) between the results from the model and the empirical results is 0.021 (± 0.002) using the offensive grid (Figure 5.4a) and 0.008 (±

0.002) using the custom grid (Figure 5.4b). 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.

5.6 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 5.5 for the English Premier League. This has possibly 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 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 5.4) as transient states.

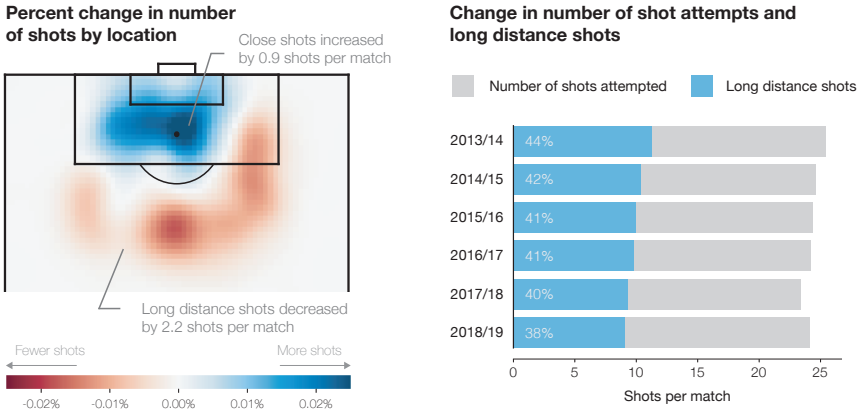

Figure 5.5: Evolution of shooting in the English Premier League between 2013/14 and 2018/19. A long-distance shot is defined as a shot taken from outside the penalty box. Over the course of this time period, the number of these shots has declined by around 20%.

5.6.1 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 5.3.1 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

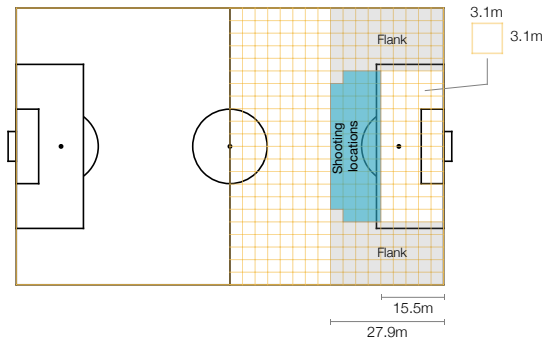

Figure 5.6: Illustration of the state space and interesting areas used to analyze shot policies. The long-distance zone from which shots can occur is shown in blue. The flanks to which the ball is often played to open up space are shaded in gray.

exactly two move actions to any field location after which a shot is attempted. We use the approach of Section 5.3.1 to calculate the exact probability of scoring in both scenarios for each state inside the blue-colored box in Figure 5.6. Figure 5.7 visualizes for each movement scenario the difference in xG 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 5.7). Hence, in some zones shooting is now preferable whereas moving was better in the one-move scenario. This is most

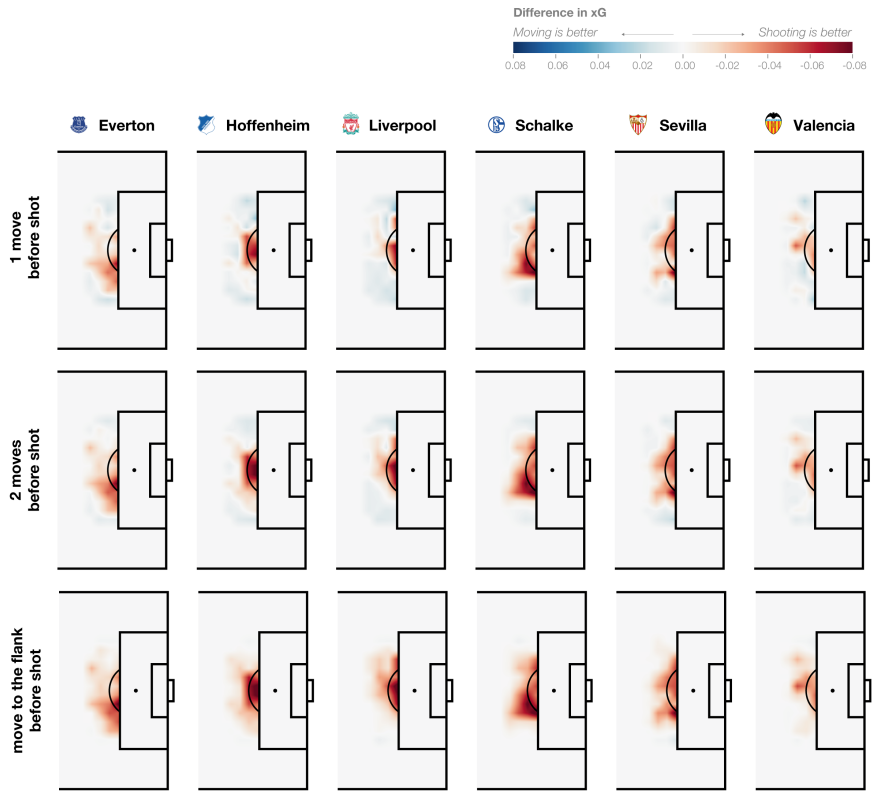

Figure 5.7: The difference in xG between moving once prior to shooting (top), moving exactly twice to any location prior to shooting (middle), or moving twice prior to shooting with the constraint that the first move action must move the ball to one of the flanks (bottom) versus directly shooting. Red 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.

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 5.6. 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 5.7 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.

5.6.2 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 5.3.1 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 5.8 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 Everton 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 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

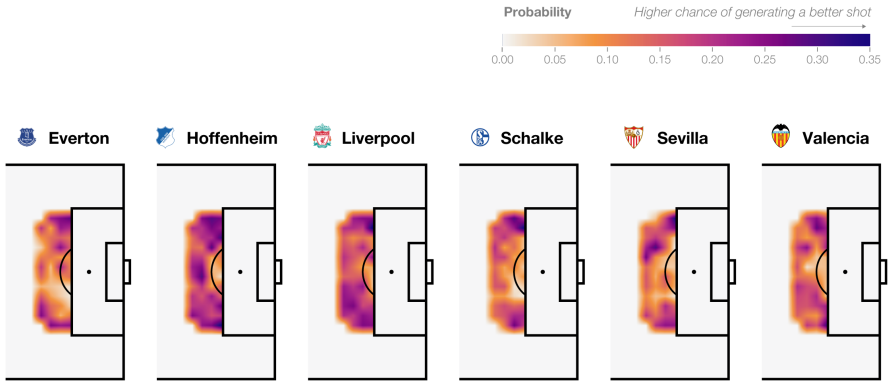

Figure 5.8: The probability of ever generating a better shot in the same possession sequence for Everton, Hoffenheim, Liverpool, Schalke, Sevilla, and Valencia. The regions with the lowest probabilities correspond to the zones where shooting was preferable to moving in the previous analyses. The magnitudes of these probabilities vary by team.

generate a better shot, whereas Everton clearly has a large region from which immediately shooting would be preferred.

5.6.3 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 5.4.2. Specifically, we explore the effect of increasing and decreasing the frequency of shooting from long-distance by 5%, 10%, and 20%.

Figure 5.9 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

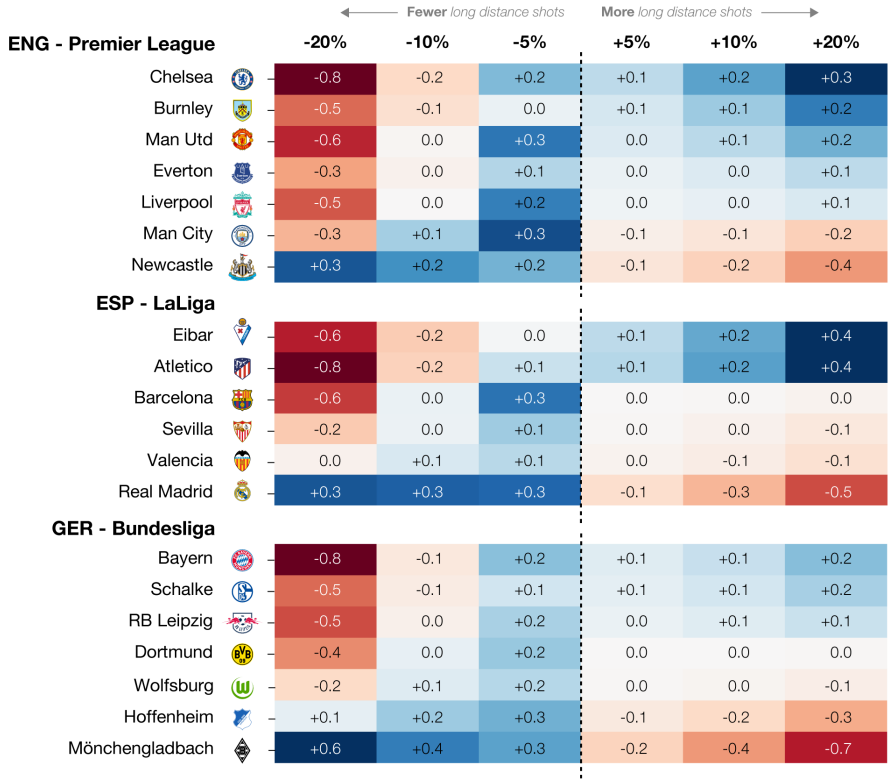

Figure 5.9: Effect on the expected number of goals a team would score when uniformly increasing or decreasing the frequency of shooting from the long-distance region shown in Figure 5.6 by 5%, 10%, and 20% for all 20 considered teams in the English Premier League, Spanish LaLiga, and German Bundesliga. 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.

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).

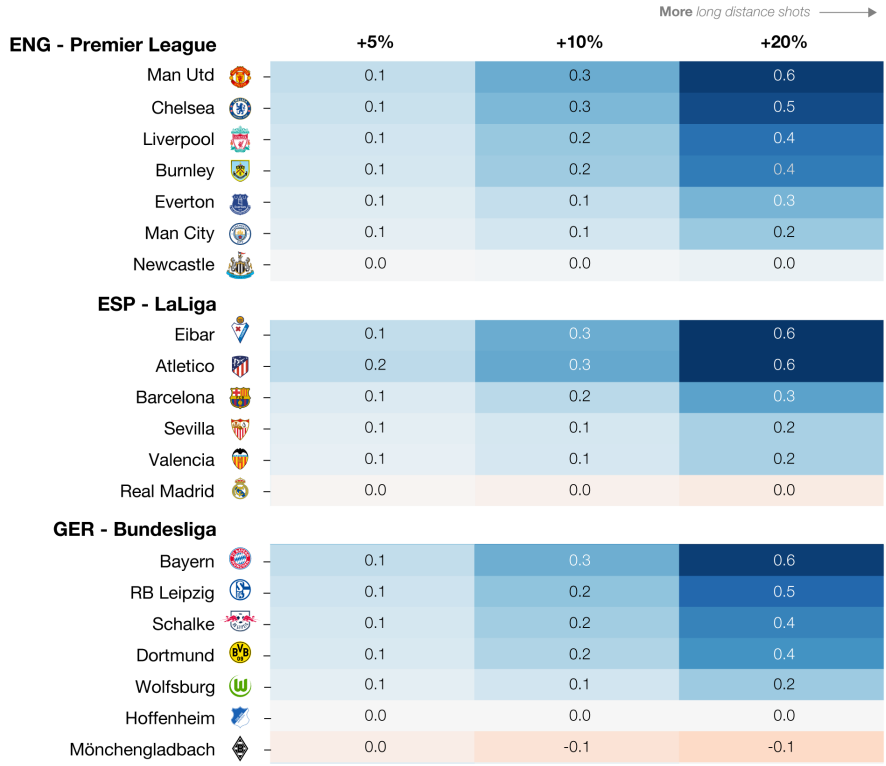

Figure 5.10: Effect on the expected number of goals a team would score when performing a targeted increase in the frequency of shooting from the long-distance region shown in Figure 5.6 by 5%, 10%, and 20% for all 20 considered teams in the English Premier League, Spanish LaLiga, and German Bundesliga. Half of the teams would score almost an extra goal. Interestingly, Mönchengladbach would still see a decrease.

5.6.4 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.

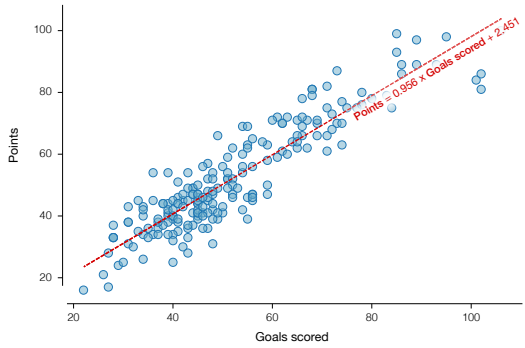

Figure 5.11: Points in the final table as a function of the goals scored. The image visualizes the data of ten seasons of the English Premier League (i.e., 2010/11 until 2019/20). Every goal scored equates to one point in the final table.

Figure 5.10 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 5.11 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 with a cost that can be estimated at around $250 million in the English Premier League.[9] For the top teams, an extra goal could mean the difference between qualifying for the Champions League or not.

5.7 Use Case: Reasoning About Defensive Tactics

In this use case, we reason about the efficacy of various defensive tactics 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?

[9] `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`

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 5.4) as transient states.

5.7.1 Shot Suppression

In this first analysis, we will use the methodology outlined in Section 5.3.2 to identify which states a defending team should pay attention to in order to *directly* and *indirectly* decrease their opponents' shots.

Figure 5.12 visualizes the percent decrease in the probability of taking a shot from and reaching the favored *shot_locations* (see Figure 5.2) 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 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

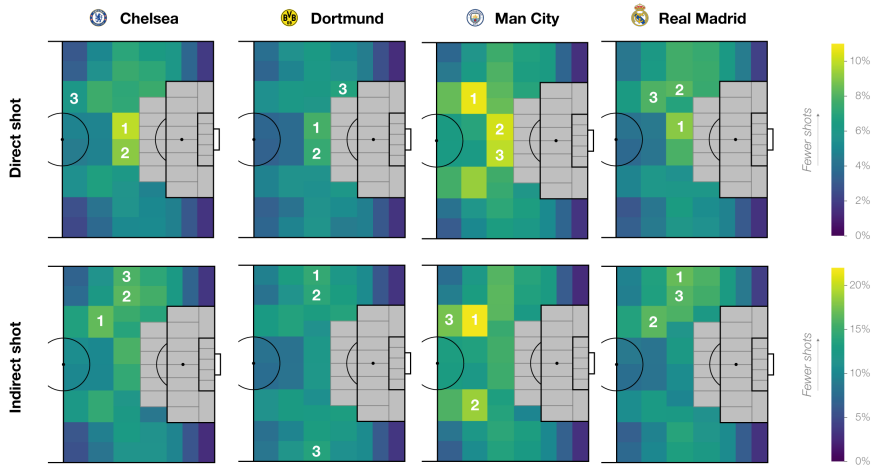

Figure 5.12: The percent decrease in shooting from (top row) and reaching (bottom row) the common shot locations for Chelsea, Dortmund, Manchester City, and Real Madrid. Yellow shading indicates states that result in a large decrease when the given team is forced to avoid them. Dark blue shading indicates states that result in a smaller decrease. The top three states with the largest impact are labeled in each figure.

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).

5.7.2 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 with the ball and will decrease the danger created by the attacking team. Here, we will use the methodology outlined in

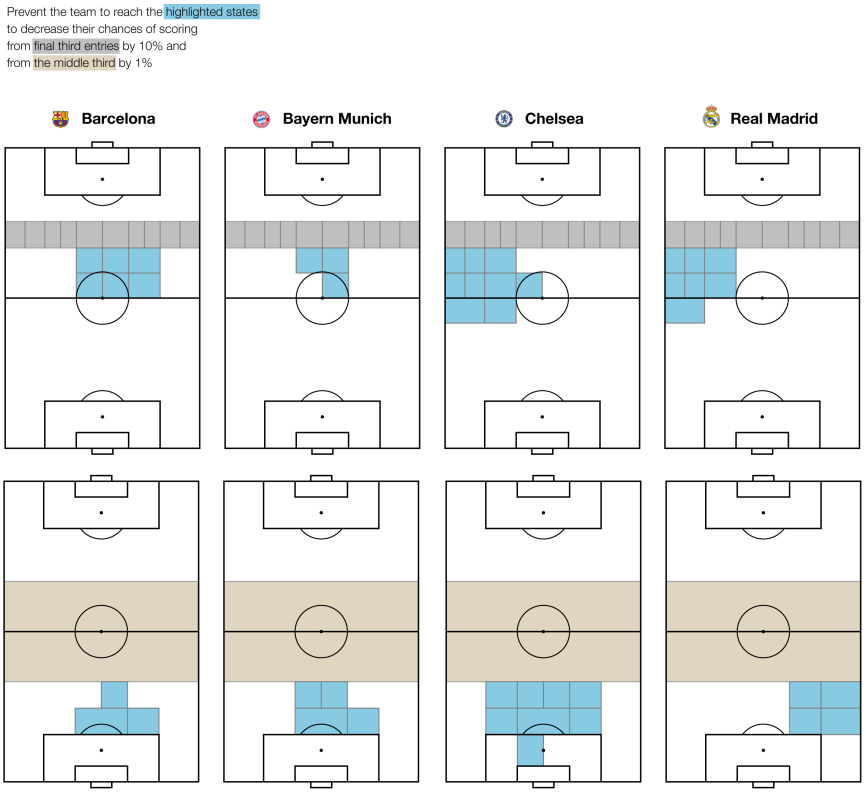

Figure 5.13: Illustrates for Barcelona, Bayern Munich, Chelsea, and Real Madrid, two areas (blue) to prevent them from reaching in order to decrease their chances of scoring in each final third entry state by at least 10% and in each middle third state by at least 1%.

Section 5.3.2 and focus on decreasing the attacking team's danger in the final third entry and middle third regions of the pitch.

Figure 5.13 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 playmakers Ivan Rakitić and Frenkie de Jong, and winger-striker 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

Table 5.4: Percent decrease in each team's probability of scoring when it is forced to avoid its respective region in the top row of Figure 5.13. The results are shown both before and after adapting to the defensive tactics, 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

cases. In the defensive part of the pitch, this is the playing field of Jérôme Boateng and 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.

5.7.3 Effect of Defensive Tactics 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 tactic, a team will eventually react and adapt their old policy π towards a new one π'. The new policy π' 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 5.4.1). 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 5.13.

Table 5.4 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 approach, 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 tactics is hard.

5.8 Related Work

Markov models and reinforcement learning paradigms have been widely used for analyzing sports such as soccer [51, 69, 70, 90, 125, 133, 148, 162], American football [56], basketball [30, 129, 130, 155], table tennis [115, 157], and ice hockey [124, 131]. 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[10], Barcelona [51]) 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 tactics 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 [125], identifying where teams create value from by using the values of players' actions [51, 133], analyzing the performance relevance of certain actions [90, 115, 157], and the analysis of the optimal or an alternative policy [129, 130, 155].

Our proposed methods also correspond to this setting, with the most closely related work being that of Sandholtz and Bornn [129, 130] for basketball. There are several differences with our work, beyond looking at different sports. First,

[10] `https://freakonomics.com/podcast/london-live/`

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 analyzing 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 [162]. With respect to the idea of analysing the effect of changing the underlying transition probabilities, recently Hirotsu et al. [69] have studied tactical changes by altering the transition rates (e.g., increasing transition rates according to home advantage) in a soccer Markov Game model with nine location-based zones. 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 tactics. Applying model checking techniques to sports models has not been extensively explored. Dong et al. [47, 48, 91] have applied probabilistic model checking to a tennis MDP to predict the win probability and identify a player's best action to improve. 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 tactics 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 [118] and uses Q-learning to reason about a learned MDP. Liu et al. [89] 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 [1] and Bai et al. [5] 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 policies 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. [138] 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 [57, 95, 124, 125, 131, 133, 150, 162]. 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 [30, 129, 130]. 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 [70]. 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 [35]. In the case of valuing player actions, some approaches validate their ratings against other metrics such as salaries, goals, and assists [87, 88, 95, 124, 131]. 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 validating 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 [131, 133]. 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 [30, 51, 129, 130].

5.9 Conclusion

This chapter proposed a novel framework to learn and reason about the policies 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 tactics 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 to reason about a soccer team's in-game behavior.

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 tactics 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 tactics 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 tactics, 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 required data is available, the MDP's state space could be expanded to more contextual and realistic states (e.g., including time and/or player information) which would allow for a more in-depth team or player analysis. This could e.g., facilitate attributing certain aspects of the used tactics to specific players. However, the more context is included in the state space, the sparser the data becomes. Therefore, an important future work is to investigate the context-sparsity trade-off in more detail. Finally, evaluating and validating the effects of alternative policies is hard. There are no ground truth labels to compare with and setting up controlled experiments in the real-world for validation is not possible. At best, one can compare an alternative simulation with a similar alternative scenario that has already happened [47]. However, this does not give any guarantees about counterfactual scenarios that were never seen before. In the future, the approaches to evaluate the effects of alternative policies could be expanded, and possibly made more realistic, 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.

Chapter 6

A Markov Framework for Analyzing Build-Up Effectiveness and Defensive Setups

In soccer, devising an effective tactical plan is an important part of match preparation. While artificial intelligence (AI) methods are increasingly being used as valuable tools in this endeavor, current methods predominantly focus on identifying offensive patterns and do not provide insight into the efficacy of defensive tactics that can be used to counteract them [25, 40, 54, 156, 161]. From an AI perspective, considering both the defensive and offensive tactics presents a challenging problem requiring the ability to learn from and generalize over sparse spatio-temporal data, and to reason about the interplay between both types of tactics.

To address this gap in data-driven tactical analysis, we quantify the impact of defensive tactics on the opponent's offensive patterns using the learning and reasoning approach developed in the previous chapter. Moreover, this chapter advances the framework in two important ways. First, it goes beyond simply tracking the ball to consider the positions of the teammates in the state space. Second, instead of only modeling the offensive team, it also reasons about how the defensive team is positioned. A final contribution of this chapter is the novel approach to quantitatively measure this impact.

To keep this chapter concise, we specifically focus on formulating a tactical plan for the defending team during the opponent's *build-up* phase, which is one of soccer's four attacking phases [84]. In this phase, the attacking team starts in their own half and attempts to break the opposition's first line of pressure and progress play into midfield, either centrally or out wide. Typically, this phase involves a structured arrangement of the team in possession against an organized defense. Furthermore, due to a lack of direct pressure, the attacking team can typically take their time passing the ball around until an opportunity presents itself to execute rehearsed patterns. This structured nature of build-ups contrasts with more chaotic phases of the game, such as transitions and counterattacks, which may require quick decision-making and involve more spontaneity. Consequently, build-ups emerge as a key phase for tactical planning that is well-suited for data-based analysis.

The content of this chapter is based on the following published and submitted publications [145, 146].

Van Roy, M., Robberechts, P., and Davis, J. Optimally Disrupting Opponent Build-ups. In *StatsBomb Conference* (2021), pp. 1–16

Van Roy, M., Robberechts, P., and Davis, J. A Markov Model Approach to Analyzing Build-Up Efficiency and Defensive Setups in Soccer. *Submitted to Data Mining and Knowledge Discovery* (2024)

6.1 Shortcomings of Raw Data Querying Techniques

Quantifying the impact of defensive tactics on the opponent's offensive patterns involves three different tasks. First, it is necessary to discover the opponent's offensive patterns. Second, it is necessary to represent the defensive tactics and model their interaction with the discovered offensive patterns. This requires integrating information about how the ball moves in response to the various arrangements of the defensive players on the pitch. Finally, a metric to measure the efficacy of defensive tactics must be defined. Arguably, a defensive tactic is effective when it prevents the opponent from scoring, but goals are rare. Even shots account for $\sim 1\%$ of all actions during a match. Hence, it is crucial to take into account near misses (i.e., one step away from success).

One approach to address these challenges would be querying the data to find frequently occurring patterns and count how often they lead to some kind of success criteria (e.g., final third entries), possibly conditioned on defensive information. However, this approach has several significant drawbacks. First, it involves significant hand-crafting (e.g., the features used to represent the sequences). Second, soccer is highly dynamic and teams rarely perform the exact same sequence of actions repeatedly. Thus, there is a need to capture small variations both in patterns and what constitutes success (i.e., near misses) which introduces a mismatch with the discrete approach often taken in pattern mining. Finally, the data is hand annotated (i.e., there are errors and disagreements) and incomplete (i.e., when actions fail, it is unknown where the player was aiming to move the ball to).

6.2 A Markov Model of Build-Ups

To tackle the above tasks and challenges, we take an alternative approach of combining learning and reasoning. The learning component entails learning a behavioral model of the teams that can be used to reason about the impact of various defensive tactics. Similar to the approach in Chapter 5 and related approaches [30, 51, 131, 138], we use a Markov decision process (MDP) to model the behavior of teams. However, these approaches typically only model the *general* behavior of the attacking team. Our proposed MDP differs from existing ones in two key ways. First, we specifically focus on *build-up play* as opposed to the general attacking play. Second, we analyze the *efficacy of defensive setups* and provide two ways to incorporate defensive information: (1) for each team, we learn one MDP per coarse-grained defensive setup of the opponent, and (2) we introduce a novel state space for the MDPs that captures defensive information by modeling the available passing options.

6.2.1 Defining Build-Ups and Defensive Setups

Our approach uses contextual event stream data, which is extracted from broadcast video. It consists of (1) event stream data that describes all on-the-ball actions in a game (e.g., type, start and end location, result, etc.) and (2) information about off-the-ball players that are visible on camera such as their location and relationship to the ball carrier (i.e., teammate or opponent). Given this data, we define a *build-up* and the considered *defensive setups* as follows.

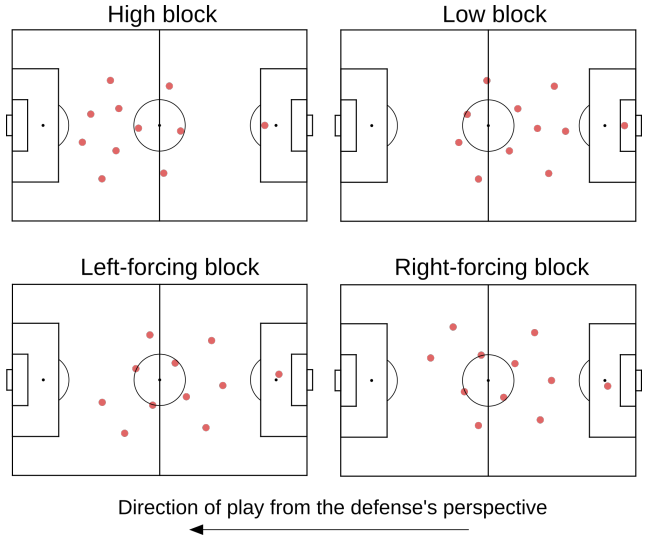

Figure 6.1: Illustration of the defensive setups with dots depicting the defending team's players.

Build-up

We define a build-up as a possession sequence starting with a center back in ball possession and no opponent pressure on him. It ends when the team loses possession or shoots.

Defensive setups

We consider defensive setups that correspond with typical setups used by practitioners and in literature [8, 60, 92, 116]. Thus, the insights can be formulated in terms practitioners know, increasing the explainability of our approach. We consider the following two scenarios, as illustrated in Figure 6.1:

1. **A high vs. low block.** In a high block, the defending team's players are positioned closer to the opponent's goal. This is done to exert pressure on the ball-possessing team, often with the objective of winning the ball back near their goal. The disadvantage is that it leaves space behind the defenders if the possessing team can break the pressure. A low block is more conservative with the defending players closer to their own goal. We define a high (low) block as a situation in which the majority of visible

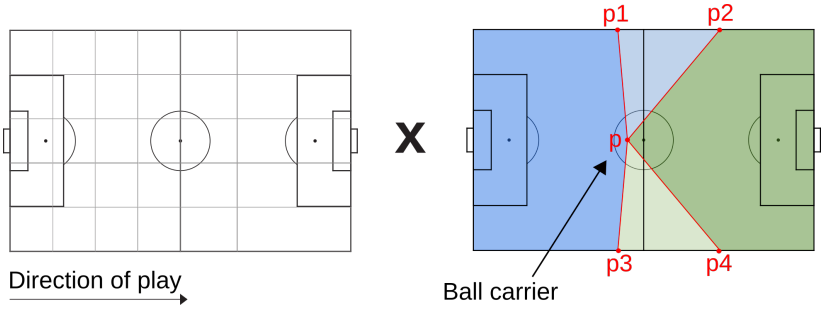

Figure 6.2: The location-based states (left) and extended states (left × right) with patches created relative to the ball carrier.

players of the defending team is higher (lower) up the field than five meters past the halfway line on their opponent's (own) half.

2. **A left- vs. right-forcing block.** For various reasons, e.g., weaker opponent players on one side, the defending team might force the opponent toward one side of the field. We define a left- (right-) forcing block as a situation in which the defending team's players are higher up the field on the opponent's right (left) side and lower up the field on the opponent's left (right) side. The opponent should also move the ball to its own left (right) side in the first two actions.

6.2.2 Defining the Markov Model

Next, we define the state space ($\mathcal{S}$), action space ($\mathcal{A}$), transition function (P), and policy (π) of the MDP.

State space

The state space includes the states that the team can enter and, possibly, exit during a build-up sequence. Absorbing states ($\mathcal{E}$) can never be exited and denote the end of a possession sequence. We define three absorbing states: $s_{loss_possession}$ (possession lost after a failed move action), s_{goal} (a goal is scored), and s_{no_goal} (a missed shot) [151]. Transient states ($\mathcal{L}$) can be exited. We define two transient state spaces, allowing us to analyze the defensive setups both at a coarse- and more fine-grained level:

1. **Location-based state space.** These 30 states denote a particular location on the field (Figure 6.2). The partitioning is chosen such that there are more fine-grained states in a team's own half (e.g., a distinction can be made between the sides, half-spaces, center, high or low up the pitch) and more coarse-grained states in the opponent's own half (e.g., a distinction between final third of the pitch or not) which facilitates the analysis of build-ups. Learning an MDP per defensive setup with these states allows us to analyze how a team builds up against that specific defensive setup.

2. **Location + reachability state space.** Each location-based state is extended by adding four binary features indicating whether there is a reachable teammate in each of the colored patches in Figure 6.2. This yields 480 transient states. A teammate is reachable in a patch if no opponent is within 1.5 meter of him and there is a 1 meter wide passing lane between the ball carrier and him. The patches that define the four additional binary features are formally defined as follows. Given the ball carrier's location $p = (p_x, p_y)$, the four intersection points with the borders of the pitch as depicted in Figure 6.2 are computed as:

$$p_1 = intr([(\sin(-\alpha)+p_x, \cos(-\alpha)+p_y), p]) \tag{6.1}$$

$$p_2 = intr([(\sin(\beta)+p_x, \cos(\beta)+p_y), p]) \tag{6.2}$$

$$p_3 = intr([(\sin(180°+\alpha)+p_x, \cos(180°+\alpha)+p_y), p]) \tag{6.3}$$

$$p_4 = intr([(\sin(180°-\beta)+p_x, \cos(180°-\beta)+p_y), p]) \tag{6.4}$$

   with $\alpha = 5°$, $\beta = 40°$, and $intr([a, b])$ a function computing the intersection point of the line between points a and b with the borders of the pitch. This state space allows for analyzing how a team reacts when players in certain areas are blocked by opponents.

A fundamental difference between both representations is how information about the defensive setup is integrated. Whether it is best to integrate this information in the MDP or learn a separate model per defensive setup depends on the available data and specific use cases. An integrated model has the advantage of allowing multiple fine-grained analyses (e.g., different blocked passing options). In contrast, when using a separate model per defensive setup, a new model should be learned for each scenario. This only works well with high-level defensive setups (e.g., high block). A disadvantage of an integrated model is the blow-up of the state space, which could cause sparsity issues when learning. In contrast, separate models have the advantage that data sparsity is

less of an issue. Additionally, these models are simple and, in practice, suffice for an analysis of high-level defensive setups.

Action space

Transitioning from a state $s \in \mathcal{L}$ to another state is done using the actions $move_to(s')$ and $shoot$. Here, $move_to(s')$ denotes the deliberate ball-moving action to state $s' \in \mathcal{L}$ and encompasses all ball-moving actions between two locations on the field (i.e., pass, dribble, cross), and $shoot$ denotes a shot on goal.

Transition function

The probability of transitioning from one state to another using a specific action is defined by the transition function. For $s \in \mathcal{E}$, this is a self-loop with probability one. For $s \in \mathcal{L}$, this is defined as:

- $P(s, move_to(s'), s')$ is the probability of successfully moving the ball from s to s', and $s' \in \mathcal{L}$.
- $P(s, move_to(s'), s_{loss_possession})$ is the probability of unsuccessfully moving the ball from s to s'. The possession sequence ends in $s_{loss_possession}$.
- $P(s, shoot, s_{goal})$ is the probability of scoring from s.
- $P(s, shoot, s_{no_goal})$ is the probability of missing the shot from s. The possession sequence ends in s_{no_goal}.
- $P(s, a, s') = 0$ in all other cases.

Policy

The policy defines the probability distribution over the actions (i.e., how a team chooses to perform the different actions) in each state. For an action a and a state $s \in \mathcal{L}$:

$$\pi(a \mid s) = Pr[A = a \mid S = s]. \tag{6.5}$$

6.2.3 Learning the Markov Model

Learning the transition function and policy involves two main challenges. First, the intended end location of the ball-moving actions is not always known (e.g., for intercepted passes, it is unclear where the player was aiming). Second, the data is extremely sparse. We address these two challenges below.

Predicting the intended end state

To infer the missing intended end state, we apply the method outlined in Chapter 5. This method learns to predict a probability distribution over the intended end states from a team's successful executions of each action type. The distributions are post-processed using domain knowledge to exclude unlikely states.

We adapt the post-processing step to also consider the positions of teammates and opponents, which are available in the contextual event stream data. Including this additional context should improve the predictions. Our hypothesis is that a state is more likely to be the intended end state when there is a teammate in it, but slightly less likely when that teammate is close to an opponent. We incorporate this risk-reward trade-off and boost the states with teammates based on the proximity of opponents. That is, we rank the teammates from closest to furthest from opponents and multiply the probabilities of their corresponding states by their rank. The distribution is renormalized afterward.

Learning the transition function and policy

When using the location-based states, learning a team's transition function and policy follows a standard counting approach. Given a state $s \in \mathcal{L}$:

$$\pi(shoot \mid s) = \frac{\#shots(s)}{\#actions(s)} \tag{6.6}$$

$$\pi(move_to(s') \mid s) = \frac{\#attempts(s, s')}{\#actions(s)} \tag{6.7}$$

$$P(s,\ move_to(s'),\ s') = \frac{\#success(s, s')}{\#attempts(s, s')} \tag{6.8}$$

$$P(s,\ shoot,\ s_{goal}) = \frac{\#goals(s)}{\#shots(s)}. \tag{6.9}$$

Afterward, Gaussian smoothing is applied to ensure the policy and transition function are spatially coherent. Given that these states cover large areas of the pitch, we ensure that while smoothing the policy and transition function in a state, we only take adjacent states into account. Therefore, the standard deviation of the kernel is set to 0.25, yielding a kernel size in both dimensions of $2 * \lceil 3 * 0.25 \rceil + 1 = 3$ [132]. Thus, to smooth the value of the state in the center of the kernel, only the values of its immediate adjacent states are considered.

When using the extended states, a similar smoothed prior for the policy and transition function is learned using the data of all other available teams. A Bayesian average with this prior is used to compute the policy and transition function of a specific team. Given a state $s \in \mathcal{L}$, these are calculated as:

$$\pi(shoot \mid s) = \frac{C_1 m_1 + \#shots(s)}{C_1 + \#actions(s)} \tag{6.10}$$

$$\pi(move_to(s') \mid s) = \frac{C_2 m_2 + \#attempts(s, s')}{C_2 + \#actions(s)} \tag{6.11}$$

$$P(s,\ move_to(s'),\ s') = \frac{C_3 m_3 + \#success(s, s')}{C_3 + \#attempts(s, s')} \tag{6.12}$$

$$P(s,\ shoot,\ s_{goal}) = \frac{C_4 m_4 + \#goals(s)}{C_4 + \#shots(s)}. \tag{6.13}$$

Here, m_1, m_2, m_3, and m_4 are the prior probabilities of shooting in state s, moving to state s' from state s, successfully moving to state s' from state s, and scoring from state s, respectively. All C_i are constants denoting the number of data points of the corresponding value m_i that are added. In case there are less than 30 actions in a state, all C_i for that state are set in function of the number of actions in the state to ensure a tight 99% confidence interval around the prior. Otherwise, all $C_i = 0$ and no prior information is used. The minimum number of actions needed to use the prior was determined empirically by analyzing the average number of visits per state. This approach mitigates data sparsity issues in states where limited data is available.

6.3 Reasoning About Build-Up Effectiveness

Using the constructed models, we want to identify which defensive setup(s) can disrupt the attacking team's build-up sequences. To this end, we analyze the *build-up effectiveness* of the attacking team when playing against various defensive setups. We define build-up effectiveness as the *probability of reaching*

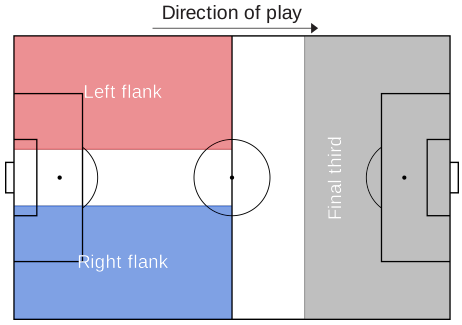

Figure 6.3: Illustration of the final third of the pitch (gray), and the left (red) and right (blue) flank during build-up.

the final third of the pitch (Figure 6.3). This can be seen as having reached a desired and threatening situation from which the next game phase can begin. In comparison to goals or shots, final third entries are much more frequent and therefore provide a stronger signal.

Computing the effectiveness using the observed build-ups leading to a final third entry is not meaningful due to the small sample sizes and variability in outcomes. For example, only half of the build-up sequences reach the final third and a quarter of the build-up sequences are near misses.[1] Therefore, we compute a team's build-up effectiveness by using the probabilistic verification tool PRISM [83] on the MDPs. This allows us to generalize over the various build-up patterns. We formulate the property "reaching the final third" as a Probabilistic Computational Tree Logic (PCTL) [65] formula and use PRISM to verify this property against the team's model. We do this for each state $s \in \mathcal{L}$ in which a build-up can start:

$$p_s = \mathtt{P}_{s=?}[\mathtt{F}\ final_third]. \tag{6.14}$$

Here, $\mathtt{P}_{s=?}$ indicates that the query computes the probability of the given property in state s, $\mathtt{F}$ denotes "future/eventually", and $final_third$ is the set of states denoting the final third of the pitch. For each state s, this is weighed by the probability of the build-ups starting in that state ($init_p_s$), resulting in the total probability of the team reaching the final third during a build-up sequence:

$$p_{eff} = \sum_{s \in \mathcal{L}} p_s * init_p_s. \tag{6.15}$$

We analyze p_{eff} under three opposing circumstances:

[1] See Appendix B for detailed percentages of build-ups ending in each state.

1. **Starting the build-up from the left or right flank.** This allows us to identify if a team is less effective when they start their build-up from a particular side of the pitch. For this, Equation 6.15 is adapted to only aggregate over states s that are part of one flank (Figure 6.3). We further contextualize this analysis by investigating it when playing against both a high and a low defensive block.

2. **Playing against a left- or right-forcing block.** This allows us to identify to which side the defending team should force their opponent to play the ball in order to diminish their effectiveness. All states s are used in Equation 6.15.

3. **Playing against a team that actively blocks players to prevent the first pass going to the left or right side.** This illustrates the added value of the reachability information and allows us to analyze the effect of certain unreachable players. Equation 6.14 is constrained to only allow the first action ($\texttt{X}$ = "next") to go to one of the "free" players, which are denoted by their states $free_states$:

$$p_s = \texttt{P}_{s=?}[\texttt{X}\ free_states \wedge \texttt{X}\ (\texttt{F}\ final_third)]. \tag{6.16}$$

6.4 Experiments

As no ground truth for (the effectiveness of) tactical patterns exists, evaluating our approach is challenging. Therefore, we apply two evaluation strategies. First, we illustrate our approach on a number of use cases and qualitatively evaluate them by comparing the results with the analysis of a domain expert.[2] Second, as an objective measure of our model, we quantitatively evaluate the stability of the results and evaluate several of our design choices. In summary, we address the following components and research questions.

The Markov model Each team's build-up tactics should be faithfully captured by its MDP in order to be useful for practitioners.

Q1 When comparing with the expert's intuition, do the most likely build-up sequences according to each team's model align with it?

[2]We compare our results for RB Leipzig with the analysis of a domain expert of the Belgian First Division team Club Brugge, which was their opponent during the 2021/22 Champions League. Our analysis is based on data of 2020/21 when Leipzig was coached by Julian Nagelsmann, while Club Brugge's analysis was mainly based on the 2021/22 season with Jesse Marsch as coach.

Q2 When comparing with the raw data, does the proposed approach learn a faithful representation of each team's build-up tactics? Is this affected by the different state spaces, and hence, different learning approaches?

Q3 Knowing the intended end state is crucial to learn the model. Does the additional information about the positions of teammates and opponents improve the predictions?

Build-up effectiveness The proposed probabilistic model checking approach should be able to extract a team's build-up effectiveness and the effect of defensive setups to aid the practitioners.

Q4 Is there a difference in build-up effectiveness between flanks for certain teams, and does this differ between a high and a low defensive block? Do these insights match with the expert's intuition?

Q5 Do certain teams become less efficient when the defensive setup contains an active component (e.g., actively forcing to a side, actively blocking opponents)?

Evaluation and applicability The proposed approach should be applicable to multiple data sets with different characteristics and the extracted insights should be stable for each team.

Q6 Is the proposed approach applicable to data sets of different leagues and competitions with different characteristics (e.g., availability of data)?

Q7 How stable are the results when training the models on slightly different subsets of data from the same season?

6.4.1 Data

Our data set consists of StatsBomb 360 contextual event stream data for the 2020/21 German Bundesliga, the 2020/21 Spanish LaLiga and the 2020 UEFA Euro Championship. The 2020/21 German Bundesliga and Spanish LaLiga data sets each include all games played by the top 10 teams in the league. This results in a total of 250 and 290 games, yielding 34 and 38 games per team in the top 10, for the respective leagues. The 2020 UEFA Euro Championship data set includes all 51 games of 24 teams, starting at the Group Stage up to the Final. Each team plays at least 3 and at most 7 games. We considered the five teams with the most played games: Belgium (5), Denmark (6), England

Table 6.1: The average number of actions (± 1 std) per team during the extracted build-up situations in each data set.

Data	All Actions	Passes	Crosses	Dribbles	Shots
Bundesliga	15603 (± 2269)	8143 (± 1113)	110 (± 19)	7273 (± 1158)	79 (± 15)
LaLiga	16718 (± 4751)	8670 (± 2412)	122 (± 31)	7849 (± 2289)	79 (± 31)
Euro	3832 (± 960)	1967 (± 491)	28 (± 13)	1820 (± 454)	18 (± 8)

(7), Italy (7), and Spain (6).[3] The data sets are converted into the SPADL[4] representation before applying the proposed approaches. SPADL is a unified representation for on-the-ball soccer actions and facilitates the analysis.

We retrieve about 1700 (±160) build-up sequences for each team in the Bundesliga and LaLiga, and about 300 (±10) build-up sequences per team in Euro 2020. Against each defensive setup, we obtained about 850 (±94) build-up sequences per team in the Bundesliga and LaLiga, and 150 (±7) build-up sequences per team in Euro 2020. Table 6.1 provides an overview of the average number of actions of each type per team during the build-up situations.

6.4.2 Experimental Settings

To predict the intended end state, we apply the proposed approach per team to each set of their build-up sequences separately. We use a random forest [22] as opposed to gradient-boosted decision trees as the underlying model for computational reasons. We apply it to passes only and split the successful pass executions in a train and test set using an 80-20 split. We optimize the max tree depth (within the interval $[2, 10]$) using 3-fold cross-validation on the train set. The test set is used to evaluate the models. As crosses are limited, we use the distribution of observed end states of successful crosses.

The Bayesian approach for the models with the extended state space requires learning a prior for the policy and transition function over all other teams. We only consider the teams that played within the same league or competition to learn the prior.

Finally, for each team, we construct four MDPs using the location-based states (i.e., one against each defensive setup), and two MDPs using the extended states (i.e., against a high and a low block). The corresponding build-up sequences and intended end state predictions are used for learning.

[3]Part of the data is publicly available at `https://github.com/statsbomb/open-data`.
[4]`https://github.com/ML-KULeuven/socceraction`

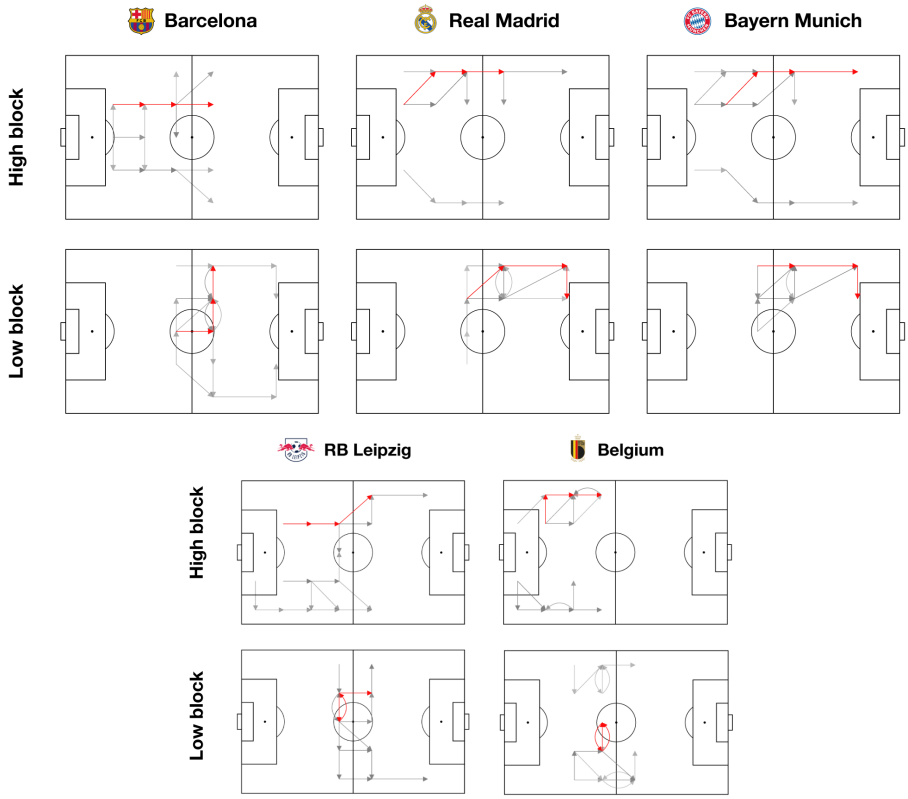

Figure 6.4: Top 10 most likely build-up sequences against a high/low block. The most likely sequence for each team is in red.

6.4.3 Identifying Offensive Build-Up Patterns

Identifying the most likely build-up patterns can provide teams with a general idea of how each opponent behaves during their build-up and help formulate plans to disrupt it. One typical approach is to apply a sequential pattern miner on the observed build-up sequences. However, simply mining all frequent patterns does not immediately yield interesting insights.[5] One would need to heavily constrain the mined patterns (i.e., specific start state, specific time steps between successive itemsets or states, etc.) to yield relevant insights. However, this is challenging and constraining the Markov model is more straightforward.

[5]See Appendix B for the results of applying a sequential pattern miner.

Therefore, as a first use case, we generate the most likely build-up sequences according to each team's model. We identify these sequences by exploiting the fact that the Markov model is a generative model. For each initial state where a build-up can start, we generate all sequences of length three.[6] We use a standard approach to compute the probability of a sequence, by multiplying the probability of starting the sequence in the initial state with the transition probabilities between the states of the sequence.

Figure 6.4 shows the ten most likely build-up sequences of Barcelona, Real Madrid, Bayern Munich, Leipzig and Belgium. Barcelona clearly had a tendency to focus on central combinations, while Real Madrid and Bayern Munich tended to use the left flank. For Leipzig, Club Brugge's analysts identified two distinct styles of build-up play. The first style is characterized by a vertical build-up through the center, often neglecting the sides until the halfway line. This style is clearly visible on Leipzig's left side when playing against a high block. The second style is characterized by a patient build-up in midfield with fast horizontal ball circulation, which is indeed confirmed by our analysis against a low block. During Euro 2020, Belgium's two defensive midfielders or wing-backs usually dropped deep. In general, we see that Belgium had trouble progressing the ball.

6.4.4 Optimizing Defensive Tactics

The ultimate goal of our method is to optimize the defensive team's tactics. To this end, we demonstrate our approach on three use cases. We analyze the build-up effectiveness when the defending team (1) employs a high or low block while the opponent starts their build-ups from their left or right side, (2) actively forces the opponent to play the ball to a certain side using the left- and right-forcing blocks, and (3) actively blocks all players on a center back's left or right side during the first action of the build-up.

Figure 6.5 shows the results of the first use case. Regardless of the block they played against, teams like Real Madrid, Stuttgart and Belgium were more efficient from their left side. For Real Madrid, this corresponds to the previously identified most likely build-up sequences. In contrast, Belgium used both flanks during the build-up, but there is a clear difference in effectiveness between them. The Belgians created more danger from their left side, where offensive wingers like Carrasco filled the wing-back position. On the right, this position was filled by Castagne and Meunier who play as more defensive full-backs at their domestic club. For some teams like Granada, we see that their most efficient

[6] Considering all possible lengths is prohibitively expensive. Three actions suffice to provide an indication of how each team's build-up sequences progress.

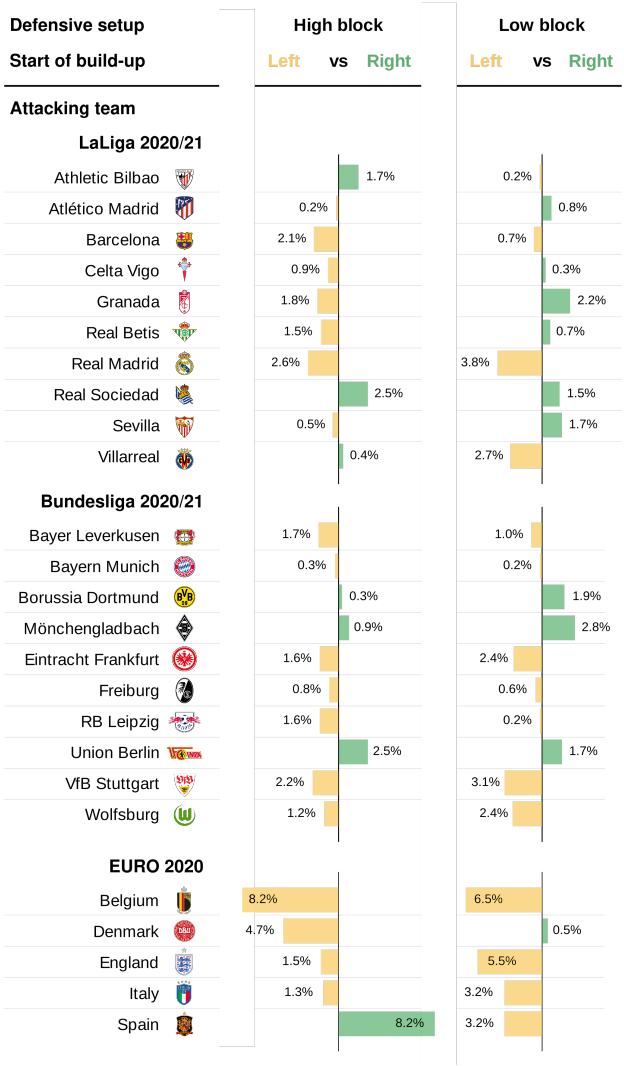

Figure 6.5: Percentage increase in build-up effectiveness when starting build-ups from the left/right side, both against a high/low block. Bars show a team's most efficient side.

side switches when playing against a high vs. low block. However, such a switch is not common. Finally, Leipzig is more efficient on their preferred left side when playing against a high block. Given they prefer fast-paced midfield play

against a low block, as confirmed by the expert's analysis, we did not see a significant difference in effectiveness between both sides.

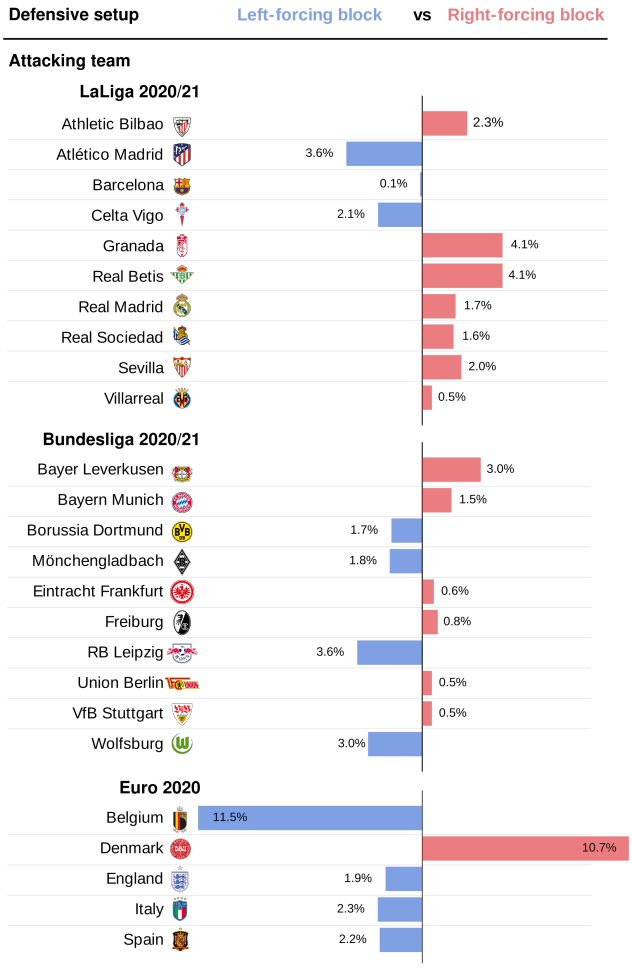

Figure 6.6: Percentage decrease in build-up effectiveness when forced to play the ball towards a certain side.

Figure 6.6 shows the results of the second use case. For most of the teams in LaLiga, it is best to force them to their right side, whereas the opposite is found for most of the teams in Euro 2020. For the Bundesliga, there is about a 50/50 split between sides. It is advantageous to force Atlético Madrid to their left side, away from the creative Trippier who finished the season with six assists from

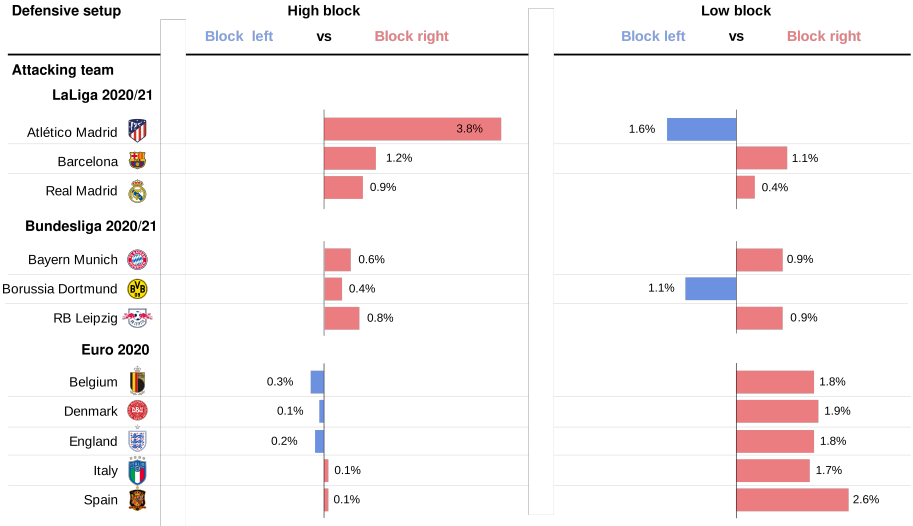

Figure 6.7: Percentage decrease in build-up effectiveness when all players are blocked on the center back's left/right side, both against a high/low block.

the right-back position. For Belgium, we again see a difference in effectiveness between sides, with the left side being less efficient. Notably, Barcelona shows no real difference, which corresponds to them predominantly using the center of the field.

Figure 6.7 shows the results of the third use case for the five considered teams in Euro 2020 and the top 3 teams in the Bundesliga and LaLiga. Figure 6.8 illustrates the definition of blocking all players on a certain side of the center back. When using a high block, it is generally best to block the players on the center back's right side for all six teams in the Bundesliga and LaLiga, with Atlético incurring a clear decrease in effectiveness. For these six teams, we see more differences when using a low block. For Dortmund and Atlético, it is best to block the players on the center back's left side, whereas for Barcelona, Bayern Munich, and Leipzig, the opposite is found. For the teams in Euro 2020, there is no clear difference visible when using a high block. However, with a low block, we see it is generally best to block the players on the center back's right side.

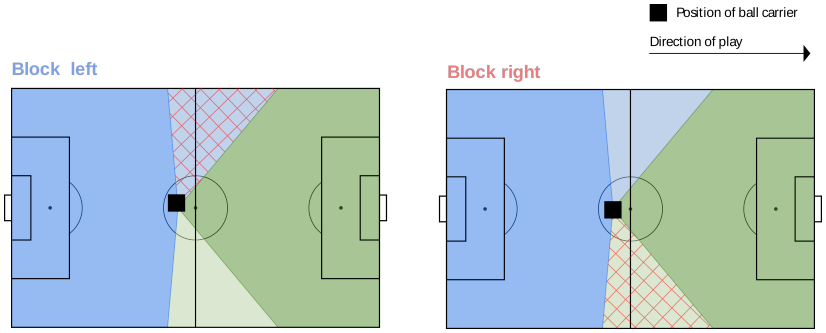

Figure 6.8: Illustrates the definition of "blocking all players on a certain side of the center back". The cross-hatched patch indicates the patch in which no reachable teammate is allowed.

6.4.5 Evaluating Stability

Within one season, there should not be much variation in the results and we would expect our approach to give quantitatively similar results when trained on different subsets of the data from the same season. To this end, we randomly split the build-up sequences of a team against each block into two distinct sets, and learn separate MDPs for each setup. We compute the build-up effectiveness for each split and compare the results per team. We repeat this ten times.

Figure 6.9 shows the stability of the build-up effectiveness against a left- and right-forcing block for all LaLiga teams. The probabilities in the first vs. the second split lie along the diagonal, indicating that the results are stable across slightly different data sets. For each team, a negative correlation is visible over the ten runs. This variance is due to the small number of successful build-up sequences and the resampling not being stratified to ensure equal numbers in each split. Additionally, this experiment gives an insight into which teams are similar in terms of build-up effectiveness against certain defensive setups. The results of the other defensive setups and leagues can be found in Appendix B.

6.4.6 Evaluating Design Choices

Finally, we evaluate our approach for predicting the intended end state and evaluate the MDPs to assess whether they faithfully capture each team's build-up tactics.

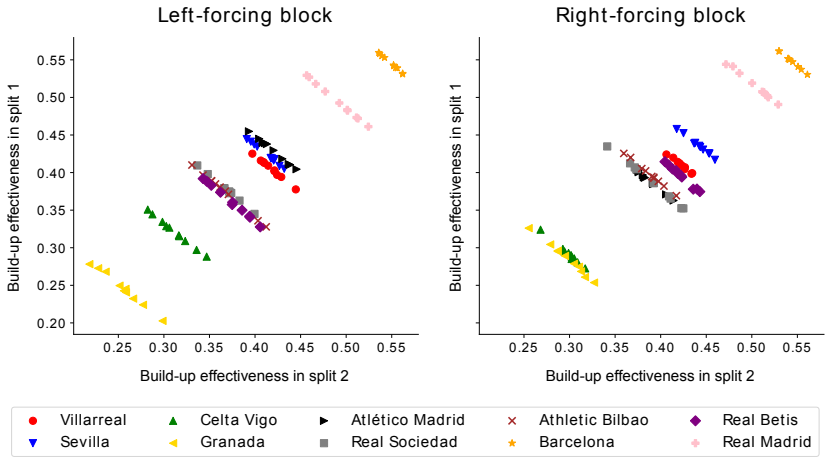

Figure 6.9: Stability of the build-up effectiveness against a left- and right-forcing block, compared split-wise over 10 runs.

Intended end state predictions

To evaluate our adaptation of the intended end state prediction models for passes, we perform an ablation study on the held-out test set. We use the original Brier score [23], averaged over all teams (±1 std), to assess whether the predicted distributions are well-calibrated. The additional post-processing step based on teammate information improves upon the results obtained with the original approach for both state spaces (Table 6.2). Predictions are less accurate for the extended states as the number of possible intended end states for which a probability can be predicted is increased by a factor of 16.

Table 6.2: Original Brier scores (± 1 std) for the intended state prediction with the original (O) and teammate information (TM) approach, averaged over all teams. The baseline predicts the distribution of observed end locations in the train set.

State space	Baseline	O	TM
Location-based	0.953 (±0.003)	0.435 (±0.031)	**0.415** (±**0.029**)
Extended	0.992 (±0.002)	0.547 (±0.024)	**0.511** (±**0.026**)

The Markov model

Evaluating whether the learned model faithfully captures the build-up tactics of each team is not straightforward as there is no ground-truth model to compare with. Instead, we use the raw data of each team and assess whether the model captures the same information by comparing the results on a soccer-relevant metric. More precisely, we calculate the probability of scoring in the future for each state in the model and compare this with the values obtained from the raw data [151].

The probability of scoring is computed for all $s \in \mathcal{L}$ using the value function V_π [12] with a reward function $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{N}$ that rewards a goal scored from $s \in \mathcal{L}$ with 1, and 0 otherwise; goals are not discounted (γ=1):

$$V_\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}} P(s, a, s')(R(s, a, s') + \gamma V_\pi(s')). \tag{6.17}$$

Since this considers all the model's learned probabilities, it enables to evaluate all aspects of the model.

Using the raw data, the probabilities can be computed by (1) identifying all regular build-up sequences and all build-up sequences that yield a goal, (2) identifying their start state, and (3) calculating for each state the probability of a build-up sequence yielding a goal given that it started in that state. The mean absolute error ($\pm$1 std) between the results from the model and the raw data, averaged over all teams, is 0.005 ($\pm$0.0009) using the location-based state space and 0.006 ($\pm$0.003) using the extended state space. The small difference indicates the model faithfully models the build-up tactics. The difference may arise because the Markov model generalizes over all *possible* build-up sequences, regardless of whether they ever appeared in the data. The slight difference between the two state spaces is due to the intended end state predictions being less accurate for the extended state space.

6.5 Related Work

Markov models have been used extensively to model the general attacking game-play in team sports [30, 51, 130, 131, 138, 151]. Their most prominent uses are to objectively quantify a player's contributions during a game and to analyze in-game decision-making. In contrast, our MDP is tailored towards analyzing the build-up as opposed to general game-play, and encodes the defensive setups which are not explicitly modeled, or even considered, in other approaches.

Analyzing the patterns that sports teams use has mainly been done using three approaches: graph-based, where typically the passes between players are inspected [25, 32, 54, 100, 110, 111]; clustering, applied to single actions or action sequences [36, 61, 62, 105, 106]; and pattern mining, possibly in combination with clustering techniques [11, 26, 36]. These works mainly analyze the frequencies of the patterns and not their effectiveness.

Probabilistic model checking techniques have only recently been applied to sports models [48, 91, 151]. They have mainly been used for aiding in-game offensive decision making, and not for evaluating the efficacy of defensive tactics.

6.6 Conclusion

This chapter proposed a learning and reasoning approach for (1) modeling the build-up tactics and defensive setups of professional soccer teams using a Markov model, and (2) reasoning about the effectiveness of build-up tactics in relation to the defensive setups. The obtained insights were presented to a professional soccer analyst, who confirmed their relevance and stated they could be applied during match preparations. This work provides a valuable real-world use case of combining learning and reasoning techniques to analyze the impact of defensive tactics and is a relevant contribution to improving soccer practitioners' interpretation of the complex patterns involved.

Currently, the approach contains similar limitations related to the included context as in Chapter 5. However, in this work, we have attempted to explore how additional defensive information could be included in the Markov model. While this is only a first step, future research should explore how additional contextual information such as game time (e.g., to model the speed of a build-up) or player identities and locations (e.g., to model the offensive and defensive positioning in more detail) could be incorporated, possibly by investigating other modeling approaches as well (e.g., an implicit versus an explicit state space, as has been done with deep learning approaches [87]).

Chapter 7

Conclusion

This chapter summarizes the current research gaps and related data challenges, outlines the contributions of this dissertation, and provides possible directions for future research.

7.1 Research Gaps and Data Challenges

This dissertation focused on the novel research direction of combining learning and reasoning to provide insights about policies used by agents in dynamic environments. In particular, the dissertation drew motivation from the task of analyzing in-game behavior and decision-making of soccer players and teams to drive new technical advances.

Deriving insights into policies requires both the ability to "learn" the current policies and their effects from the available data, and "reason" about the learned, and possibly counterfactual, policies and their effects. Current approaches focus mostly on learning-based techniques. Typically, they are used to (1) learn the current policies from data, and (2) derive insights about individual actions by combining the learning-based techniques with the most basic form of reasoning, i.e., reasoning about how an action changes the environment in one time step. This is often used to estimate the impact of a single action on the likelihood of scoring a goal in the future. While the combination of learning and the simplest form of reasoning is already useful, we identified four main research gaps.

First, a one-step reasoning approach does not explicitly take the subsequent actions and their effects into account. However, this is required to be able to

reason about the sequential behavior of actions and their effects. Thus, current approaches do not allow for trading off the different merits and pitfalls of the various sequences of actions, and hence cannot give insights into the learned policies.

Second, it is not possible to reason about counterfactual situations and the implications of alternative decisions, as learning-based approaches focus only on the current trends in the data. Yet, tackling counterfactuals is at the heart of many of the stakeholders' questions.

Third, learning-based models struggle to extrapolate beyond what they have seen in the training data. As sports continuously evolve, the examples the models were trained on might deviate from the examples for which predictions need to be made. Thus, the predictions for these examples might not be trustworthy. This is problematic because the most common use case of these learning-based frameworks, i.e., predicting the probability of a shot resulting in a goal (xG), is often shown on live television, and scouting decisions are often made based on the values of such action-value frameworks. In general, it would be bad to distribute untrustworthy values or base costly decisions on them. Thus, it would be beneficial to flag such cases to the end user.

Fourth, the most prominent use case of one-step reasoning in soccer is to estimate how much an action changes the likelihood of scoring a goal. While there exist numerous such frameworks, they are never directly compared, either conceptually or empirically. Thus, it is uncertain what the advantages and disadvantages of the different approaches are and a deeper understanding of them is lacking.

To provide insights into the policies of players and teams, one must tackle the above limitations. This is complicated further by the nature of the in-game data in soccer which one must learn from and reason about. Concretely, the data poses five main challenges from an AI perspective.

C1: Spatial and Temporal Information. There are many heterogeneous components present in the data (e.g., what, where, when, who). Ideally, to capture the in-game behavior and its full context, automated techniques must be able to combine multiple components.

C2: Data Sparsity. The most important reward in soccer, i.e., goals, is extremely rare. Additionally, teams and players rarely perform the exact same sequences of actions in the same order at the same locations. The more context is included in an analysis (e.g., specific action type, location, time, setup of players), the fewer actions the data will contain that match the specific context. Thus, the data is very sparse, which hinders AI techniques to identify general patterns in the data.

C3: Observational Data. The data is collected from what can be observed from a match, either by humans or sensors. Hence, the data can be incomplete and noisy. Therefore, one must be able to remove noise or handle the noisy nature of the data, and possibly impute missing information.

C4: Uncertainty. Soccer is a dynamic sport, involving human participants, and therefore contains lots of uncertainty and randomness (e.g., which actions are chosen, what their effect is). Appropriate techniques should be chosen that can capture and handle this uncertainty.

C5: Counterfactuals. Many of the stakeholders' questions contain a notion of counterfactuals, i.e., situations that were not captured in the data. Addressing counterfactual questions is immediately related to the above limitations of the current approaches. In general, it requires an ability to reason about the consequences of various sequences of actions.

7.2 Contributions

In this dissertation, we made three major contributions with respect to learning, reasoning, and their combination in the setting of soccer. This is not only interesting from a sports perspective, but also from an artificial intelligence perspective as soccer provides a dynamic environment with real-world data that is particularly suited for the development and evaluation of AI techniques. The insights and methodologies are also transferable to many areas that involve agents that interact with a (dynamic) environment and where one attempts to learn, analyze, and reason about the agents' behavior.

Concretely, we have made one contribution with respect to the understanding of action-value frameworks and two contributions with respect to the combination of learning and reasoning. In particular, we combined learning and reasoning to contextualize the predictions made by action-value frameworks, and to understand and reason about team behavior. Next, we summarize each of these contributions and how the various data challenges (**C1 - C5**) are tackled in each of them.

7.2.1 Contribution 1: A Critical Comparison of Action-Value Frameworks

To better understand the advantages and disadvantages of action-value frameworks, which allow for reasoning about how an action changes the environment in one step, we performed the first in-depth comparative study on

two state-of-the-art learning-based action-value frameworks, Expected Threat (xT) [133] and VAEP [36]. The key differences arise in how both approaches represent the game state and what actions are valued. These design choices influence how the spatial and temporal information is captured (**C1**) and consequently how data sparsity is dealt with (**C2**). On the one hand, xT describes each game state by only considering the location of the ball and uses a Markov model to describe the movements of the ball within a possession sequence. As the location of the ball is the only feature that is considered, xT can only value ball-progressing actions that move the ball from one zone to another (i.e., pass, dribble, cross). On the other hand, the VAEP framework uses a detailed feature-based representation that captures both the action and the game context. Therefore, it can value any on-the-ball action. However, the detailed representation comes at a cost of interpretability. In addition, VAEP uses fixed length sequences as opposed to possessions and therefore looks beyond turnovers. Our analyses showed that VAEP is better at capturing the risk-reward trade-off of actions, whereas xT is more robust.

7.2.2 Contribution 2: Contextualizing Action-Values by Reasoning

The predicted action values might not always be trustworthy. First, the collected data might include erroneous measurements (e.g., the start and/or end location may be wrong, **C3**), for which the predicted values are inherently wrong. Second, soccer naturally contains randomness and players sometimes perform unique actions that are dissimilar to what is mostly observed (**C2**). As learning techniques struggle to extrapolate beyond what they have seen during training, the predicted values for these actions should be treated with caution.

To better contextualize these predicted values, we applied a tree ensemble-based reasoning approach OC-Score [44] on top of a learned VAEP [36] action-value framework. This approach encodes each example as the ordered set of leaf nodes that are reached when executing the ensemble (i.e., output configurations). Even though examples might be similar based on their feature representations, small changes in these features (e.g., a lucky goal) can yield very different output configurations and thus final predictions. We use this approach to compare the output configurations of these examples to those of the examples on which the ensemble was trained. Our analyses show that we are able to detect annotation errors, which could help data providers automatically verify their data before release, as well as infrequent examples that are dissimilar to the training set and for which the values should be treated with caution.

7.2.3 Contribution 3: A Framework for Learning and Reasoning about Policies

We introduced a novel learning+reasoning framework for tactical analyses in soccer. The learning component entails learning the model underlying a Markov decision process (MDP) that models the offensive behavior of a team. The Markov model-based approach allows us to incorporate the uncertainty and randomness that is present in the data (**C4**). The state space of the Markov model consists of the locations on the pitch and the actions involve moving between these zones or shooting on goal (**C1**). To learn the underlying model, we (1) combine a predictive modeling approach with domain knowledge to resolve missing information (**C3**), and (2) employ a hierarchical Bayesian approach that leverages information from a "typical team" to combat the sparsity of the data (**C2**). The reasoning component operates on two levels. On the one hand, we demonstrated the application of probabilistic model checking techniques to reason about a team's tactics (**C4**). This allowed us to identify good shooting areas as well as areas that are crucial for a team to create danger, both of which are individual to each team. On the other hand, we developed analytical approaches to (slightly) modify and evaluate counterfactual tactics (**C5**). This allowed us to reason about the effects of, e.g., shooting more or less often from certain areas of the pitch.

In addition, we extended this framework to allow for analyzing defensive tactics of the opponent and the impact of these tactics on the offensive output of the attacking team. In particular, we provided two approaches for incorporating the opponent's defensive setups in the MDP (**C1, C2**): (1) by learning an MDP per defensive setup, and (2) by capturing the defensive setups in the state space. Both options are equally valuable and whether to choose one over the other depends on the available data (e.g., sparsity can cause problems with the second option) and the specific use case (e.g., a high level defensive setup can easily be evaluated using the first option). The reasoning component using probabilistic model checking allowed us to reason about the efficacy of various defensive setups for different teams (**C4**).

7.3 Future Work

This dissertation has made three main contributions regarding the combination of learning and reasoning in the context of soccer. Nevertheless, there remain many open problems, both from the sports as well as the artificial intelligence side. This section outlines four interesting directions for future work.

Context - Sparsity - Interpretability Trade-Off

One of the main challenges in this dissertation has been the trade-off of context, sparsity, and interpretability. Ideally, when modeling the in-game behavior of teams, one can include as much important and interesting context as is necessary. However, the more context that is included, the sparser the data becomes. In addition, with more context, the models themselves often become less interpretable. The approach taken in this dissertation has been to use a discretized location-based state space, possibly augmented with defensive setup information of the opponent. This allows us to investigate parts of both the offensive and defensive play of a team. Offensively, it allows for analyzing ball positioning, which is a major aspect of how teams play offensively. Defensively, it allows for analyzing a (limited) version of the defensive setup of teams. Additionally, the discretization approach allows for domain knowledge to be included, yielding practically relevant and interpretable pitch locations. However, the discretization also results in crude boundaries between areas of the pitch that are in fact very similar. Furthermore, while interesting, including additional contextual information (e.g., game time, scoreline, or player identities) in an explicit state space as used in this dissertation, makes it impossible to derive a statistically sound transition model from the data due to its sparsity. An important future work is therefore to investigate this context-sparsity-interpretability trade-off in more detail, and to investigate modeling directions such as Markov model variants with continuous time and space, using an implicit state space such as has been done with deep learning approaches [87], or possibly learning a data-driven yet interpretable state space representation (e.g., with SoccerMix [42]).

Modeling the Opponent

Currently, our in-game models depict the behavior of single teams, possibly against a fixed defensive setup. However, when evaluating which decision would be optimal, it would be better to take the behavior of each specific opponent into account. The optimal decision will most likely not be the same when playing against a high-pressing opponent versus an opponent that sits back and waits for the counter. Within the Markov model-based approaches, there has been work on using Markov Game models for this purpose, where each opponent is modeled as an actor in the game. However, most of these approaches do not focus on analyzing tactics, but only focus on a specific aspect of the game strategy such as formations and substitutions [9], or use it to value player actions [87].

Counterfactual Reasoning using Probabilistic Model Checking

In this dissertation, we have touched upon various counterfactual questions. For example, in Chapter 5, we have attempted to answer questions such as "What if players shot more often from outside the penalty box?" and "What if a defending team forced an attacking team away from a certain area on the pitch?". Our approach involved both analytical techniques and probabilistic model checking to make the required changes to the model and estimate the impact. On the one hand, analytical techniques are necessary when specific dynamics of the underlying transition model need changes. On the other hand, probabilistic model checking techniques are used to pose counterfactual questions about the states that are visited. However, probabilistic model checking cannot be used in itself to model and evaluate counterfactual changes to the underlying dynamics of the transition model. That is, the standard temporal logics in which the formal properties need to be specified (e.g. PCTL [65]) only allow for reasoning about the current *fixed* behavior that the modeled system must have (e.g. never reaching a certain state or eventually reaching a target). It is not possible to formulate a "counterfactual commitment" or "soft constraint" on the underlying dynamics (e.g., transitioning from A to B is decreased by 10% with respect to the original model), and thus reason about this alternative behavior. An important future direction for AI research is to investigate the possibility of such a counterfactual commitments or soft constraints in temporal logics. This would allow one to investigate any counterfactual question without the need for sophisticated analytical adjustments to the model, which often include various assumptions. Such counterfactual logics would not only facilitate many applications in sports, but in various other areas as well (e.g., robotics, finance).

Learning and Reasoning for Managing Physical Performance

In this dissertation, we have focused on the areas of player recruitment and match preparation to apply our learning and reasoning techniques to. However, in addition to these two areas, lots of data is being collected to monitor and manage the physical performance of players as well. To the best of our knowledge, the combination of learning and reasoning techniques has not yet been applied to provide insights into the physical performance of players, to manage it, and possibly improve it. Given that tree ensemble-based models are often used within this area due to their interpretability, one interesting future direction would be to investigate whether tree-based reasoning techniques, as those used in this dissertation, can be used to further contextualize the predictions. In addition, such reasoning techniques can also be used to ask domain-specific questions about the model. A concrete example for future work is to investigate

whether the reasoning tool `veritas` [43] can be used to identify which external load features should be kept within certain ranges to reduce a runner's injury risk [96]. By combining a learned injury prediction model and domain-specific reasoning techniques, it would be possible to give clear and actionable advice to runners about their next training sessions and the associated injury risk.

Appendix A

A Markov Framework for Learning and Reasoning About Tactics

This appendix contains information on the experimental setup, policy modifications, and ablation study of Chapter 5.

A.1 Experimental Setup

In this section, we provide all details of our experimental setup. All experiments in Chapter 5 were performed using Python v3.9.

A.1.1 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]
```

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.1.2 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 Chapter 5.

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 A.1 shows the training time needed for each of the considered state spaces and models.

Table A.1: 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

A.2 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.

A.2.1 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 $s' \in area$ by x percent can be done as follows:

$$\pi'(move_to(s') \mid s) = (1 - x) * \pi(move_to(s') \mid s) \tag{A.1}$$

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

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

$$\pi'(move_to(s'') \mid s) = \frac{\pi(move_to(s'') \mid s)}{total_other} * (total_other + x * total_area) \tag{A.2}$$

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:

$$\pi'(move_to(s'') \mid s) = \frac{\pi(move_to(s'') \mid s)}{total} * (total + x * \pi(move_to(s') \mid s)) \tag{A.3}$$

with $total$ defined as:

$$total = \sum_{s'' \in s^h} \pi(move_to(s'') \mid s) \tag{A.4}$$

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.

A.2.2 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 5.4 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.

A.3 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.1.1. 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 A.2 reports the average AUROC 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 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

Table A.2: AUROC and Brier scores for the ablation study on the end location prediction models. Results are averaged (± 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 (± 0.033)	0.535 (± 0.036)
		NO_HISTORY	0.683 (± 0.069)	0.594 (± 0.157)
		NO_BODYPART	0.711 (± 0.030)	0.534 (± 0.037)
		NO_DIRECTION	0.537 (± 0.019)	0.891 (± 0.119)
	Cross	ALL	0.511 (± 0.025)	0.961 (± 0.040)
		NO_HISTORY	0.510 (± 0.043)	0.965 (± 0.034)
		NO_BODYPART	0.521 (± 0.020)	0.969 (± 0.039)
		NO_DIRECTION	0.515 (± 0.026)	0.966 (± 0.037)
CUSTOM	Pass	ALL	0.925 (± 0.011)	0.799 (± 0.019)
		NO_HISTORY	0.922 (± 0.010)	0.785 (± 0.017)
		NO_BODYPART	0.926 (± 0.011)	0.800 (± 0.019)
		NO_DIRECTION	0.840 (± 0.019)	0.945 (± 0.006)
	Cross	ALL	0.511 (± 0.039)	0.882 (± 0.027)
		NO_HISTORY	0.516 (± 0.025)	0.867 (± 0.023)
		NO_BODYPART	0.514 (± 0.037)	0.878 (± 0.023)
		NO_DIRECTION	0.502 (± 0.022)	0.889 (± 0.038)

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.

Appendix B

A Markov Framework for Analyzing Build-Up Effectiveness and Defensive Setups

This appendix contains additional information on build-ups, the experimental setup, and the sequential pattern miner and stability experiments of Chapter 6.

B.1 Where Do Build-Ups End?

Figure B.1 shows for each state the percentage of build-up sequences that end in it. Overall, we can see that roughly half of all the build-up sequences end in the final third and about a quarter of the build-up sequences can be classified as near misses (i.e., ending in one of the states immediately to the left of the final third).

B.2 Compute Infrastructure and Timing

All models were trained on a computing server running Ubuntu 20.04 with 32GB RAM and an Intel(R) Xeon(TM) CPU E3-1225 v3 @ 3.20GHz.

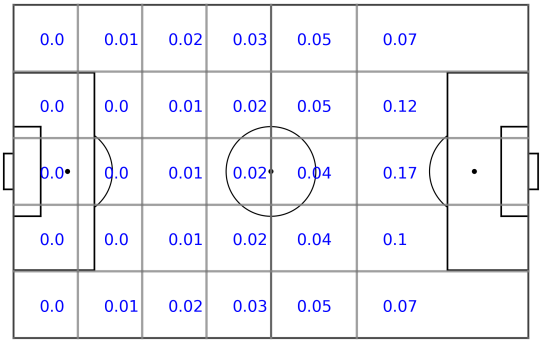

Figure B.1: For each state, the percentage of build-up sequences that ends in it.

Training the random forest for the intended end location prediction of passes takes on average 0.5 seconds (using the location-based states) and 1 second (using the extended states) for one team. Learning the MDP takes on average 5 seconds (using the location-based states) and 2 minutes (using the extended states) for one team.

B.3 Frequent Patterns Using a Sequential Pattern Miner

We illustrate the results of simply applying a sequential pattern mining algorithm to identify the most frequent patterns without constraining which patterns are found. Therefore, we apply the standard sequential pattern mining algorithm PrefixSpan on the extracted build-up sequences. Figure B.2 shows the top 10 most frequent patterns involving three actions for Barcelona and RB Leipzig found with this approach. As can be seen, the sequential pattern mining algorithm tends to find patterns that are mostly situated in midfield, often involving self-loops within the same state. As midfield is where most of the action takes place, the most frequent patterns are also situated there. However, this does not yield any interesting insight into how a team builds up. In order to find interesting patterns, one would further need to constrain the patterns e.g., to start in a certain state, to ensure a specific time step between successive itemsets, etc. This is however not straightforward.

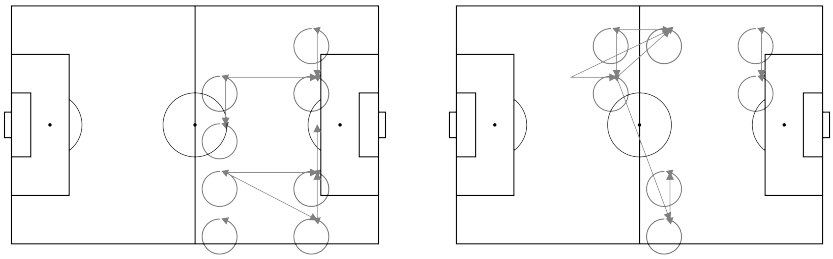

(a) Barcelona (b) RB Leipzig

Figure B.2: The top 10 most frequent patterns involving three actions for (a) Barcelona and (b) RB Leipzig according to PrefixSpan on the build-up sequences against a high block for the respective teams. Barcelona and RB Leipzig are playing left to right.

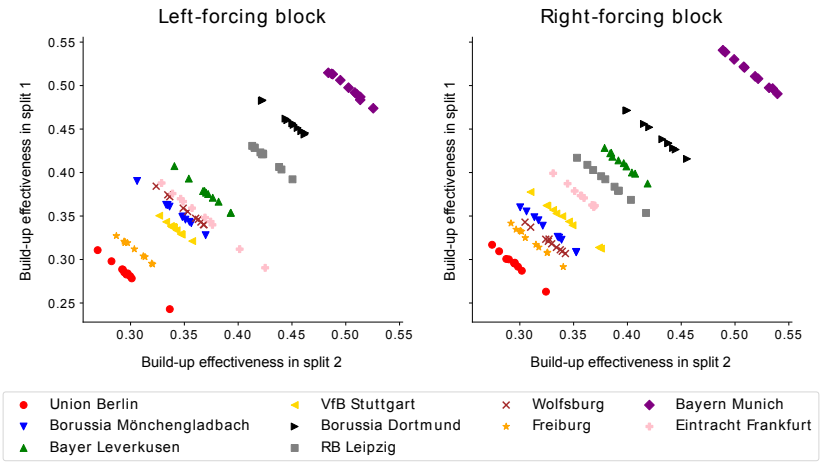

Figure B.3: Stability of the build-up effectiveness against a left- and right-forcing block for all Bundesliga teams, compared split-wise over 10 runs.

B.4 Stability Experiment

We include the additional results of the stability experiment using the data of the 2020/21 Spanish LaLiga and German Bundesliga seasons. We exclude the data of Euro 2020 as it is too limited for performing this experiment.

Figure B.3 shows the stability of the build-up effectiveness against a left- and

right-forcing block for all Bundesliga teams. The results for all LaLiga teams for these queries can be found in the main paper. Figures B.4 and B.5 show the stability of the build-up effectiveness against a high and low block, both when starting the build-up from the left and right side of the pitch. Figures B.6 and B.7 show the stability of the build-up effectiveness against a high and low block, both when the players on the center back's left and right side are blocked. For all queries, the results of our approach are stable across multiple runs with a slightly different subset of the data within one season, similar to the results presented in the main paper. The variance in the results for each team when using the extended state space is slightly different from when using the location-based state space. This is due to the use of the priors for the transition function and policy when learning the models with the extended state space.

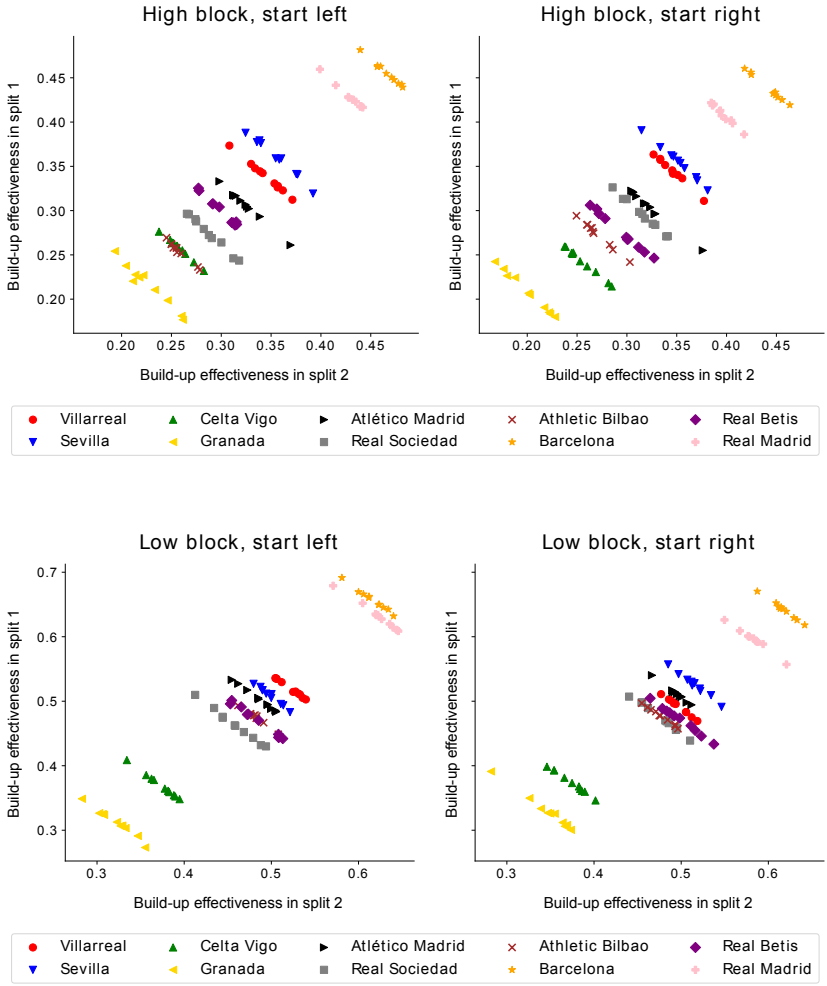

Figure B.4: Stability of the build-up effectiveness against a high and low block when starting the build-up from the left and right side for all LaLiga teams, compared split-wise over 10 runs.

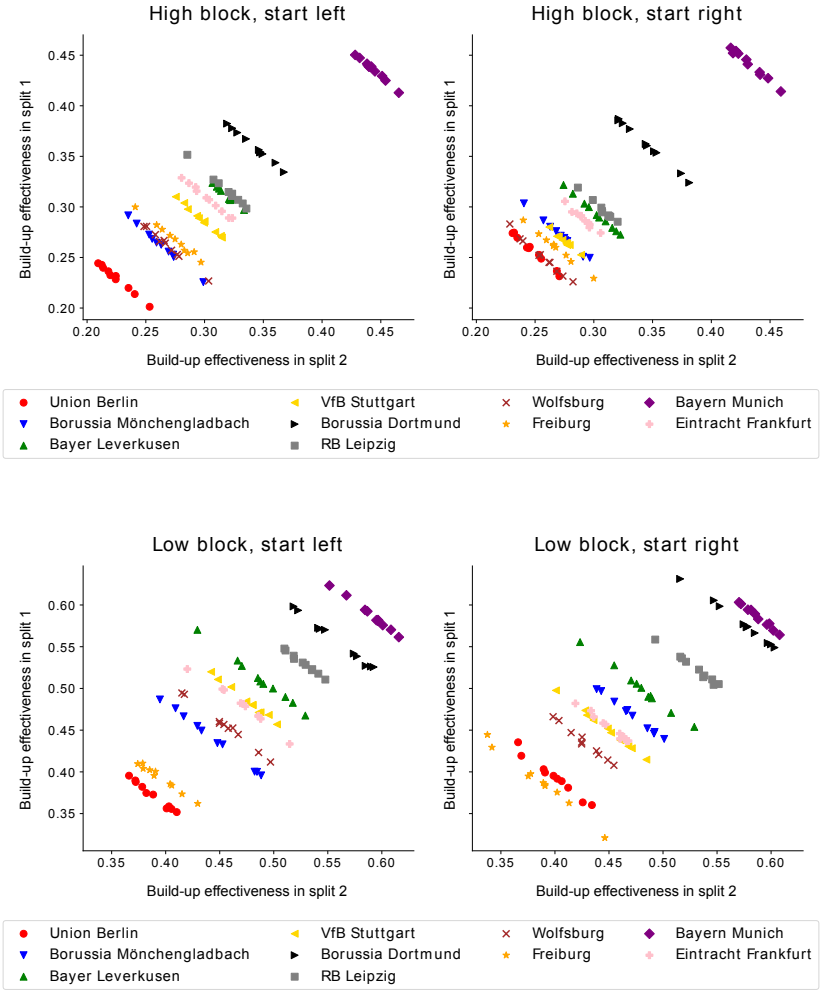

Figure B.5: Stability of the build-up effectiveness against a high and low block when starting the build-up from the left and right side for all Bundesliga teams, compared split-wise over 10 runs.

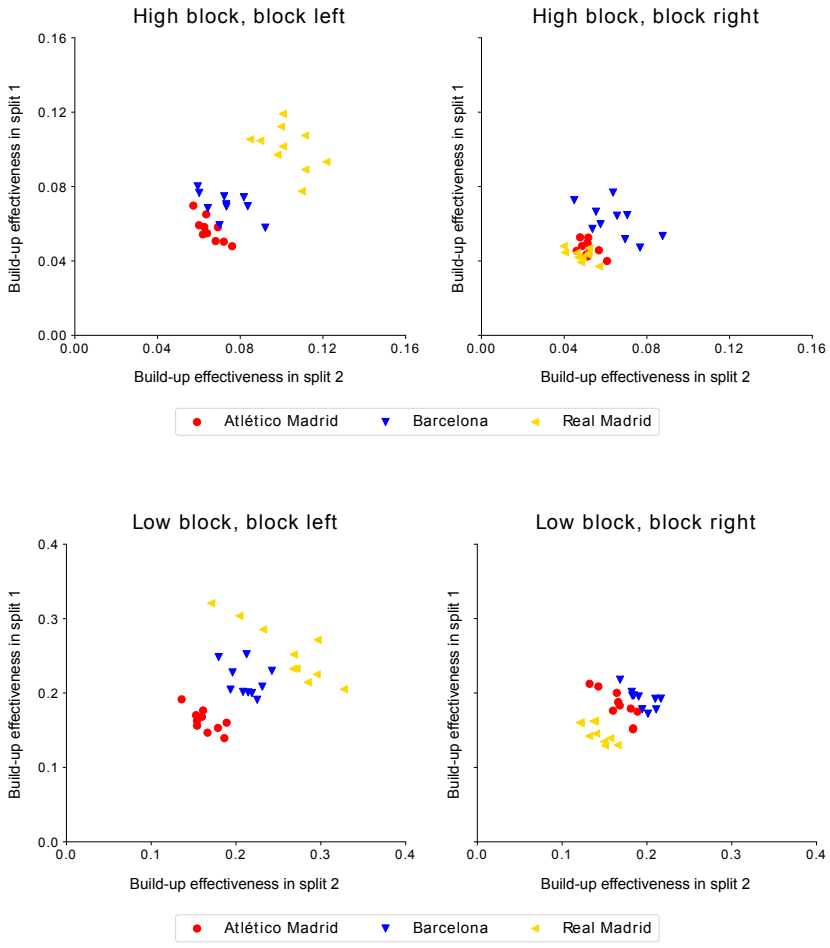

Figure B.6: Stability of the build-up effectiveness against a high and low block when blocking the players on the center back's left and right side for all considered LaLiga teams, compared split-wise over 10 runs.

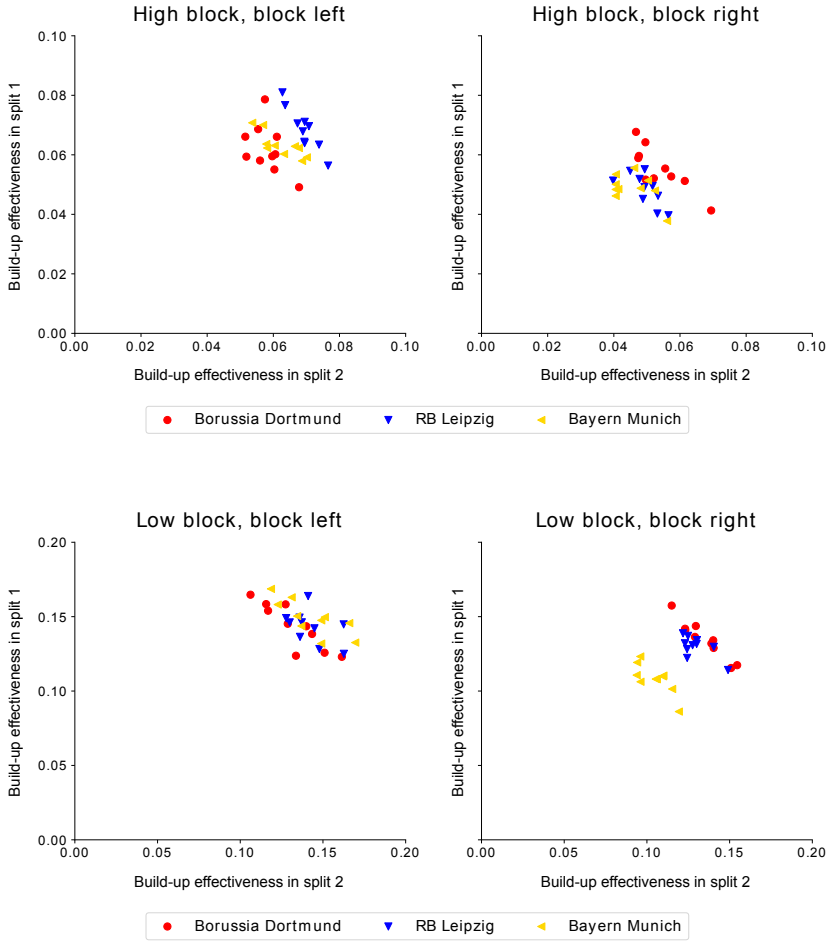

Figure B.7: Stability of the build-up effectiveness against a high and low block when blocking the players on the center back's left and right side for all considered Bundesliga teams, compared split-wise over 10 runs.

Bibliography

[1] Ahmadi, M., and Stone, P. Instance-Based Action Models for Fast Action Planning. In *RoboCup 2007: Robot Soccer World Cup XI* (2008), vol. 5001 of *Lecture Notes in Computer Science*, pp. 1–16.

[2] Amendolara, A., Pfister, D., Settelmayer, M., Shah, M., Wu, V., Donnelly, S., Johnston, B., Peterson, R., Sant, D., Kriak, J., and Bills, K. An Overview of Machine Learning Applications in Sports Injury Prediction. *Cureus 15*, 9 (2023), e46170–e46170.

[3] Anzer, G., Bauer, P., Brefeld, U., and Fassmeyer, D. Detection of tactical patterns using semi-supervised graph neural networks. In *16th MIT Sloan Sports Analytics Conference, 2022* (2022), pp. 1–15.

[4] Arbués Sangüesa, A. *A journey of computer vision in sports: from tracking to orientation-base metrics.* PhD thesis, Universitat Pompeu Fabra, 2021.

[5] Bai, A., Wu, F., and Chen, X. Towards a Principled Solution to Simulated Robot Soccer. In *RoboCup 2012: Robot Soccer World Cup XVI* (2012), vol. 7500 of *Lecture Notes in Computer Science*, pp. 141–153.

[6] Baier, C. *On algorithmic verification methods for probabilistic systems.* Habilitation thesis, Faculty of Mathematics and Informatics, University of Mannheim, 1998.

[7] Baier, C., and Katoen, J.-P. *Principles of Model Checking.* The MIT Press, 2008.

[8] Bauer, P., Anzer, G., and Shaw, L. Putting team formations in association football into context. *Journal of Sports Analytics 9* (2023), 39–59.

[9] Beal, R., Chalkiadakis, G., Norman, T. J., and Ramchurn, S. D. Optimising Game Tactics for Football. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)* (2020), International Foundation for Autonomous Agents and Multiagent Systems, p. 141–149.

[10] Beal, R., Changder, N., Norman, T., and Ramchurn, S. Learning the Value of Teamwork to Form Efficient Teams. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2020), pp. 7063–7070.

[11] Bekkers, J., and Dabadghao, S. Flow motifs in soccer: What can passing behavior tell us? *Journal of Sports Analytics 5*, 4 (2019), 299–311.

[12] Bellman, R. Dynamic Programming. *Science 153*, 3731 (1966), 34–37.

[13] Bialik, C. The People Tracking Every Touch, Pass And Tackle in the World Cup. Available at: `https://fivethirtyeight.com/features/the-people-tracking-every-touch-pass-and-tackle-in-the-world-cup/`, 2014.

[14] Bialkowski, A., Lucey, P., Carr, P., Yue, Y., Sridharan, S., and Matthews, I. Large-Scale Analysis of Soccer Matches Using Spatiotemporal Tracking Data. In *2014 IEEE International Conference on Data Mining* (2014), pp. 725–730.

[15] Borg, G. Psychophysical bases of perceived exertion. *Medicine & Science in Sports Exercise 14*, 5 (1982), 377–381.

[16] Bourdon, P. C., Cardinale, M., Murray, A., Gastin, P., Kellmann, M., Varley, M. C., Gabbett, T. J., Coutts, A. J., Burgess, D. J., Gregson, W., and Cable, N. T. Monitoring Athlete Training Loads: Consensus Statement. *International Journal of Sports Physiology and Performance 12*, Suppl 2 (2017), 161–170.

[17] Bransen, L., Robberechts, P., Van Haaren, J., and Davis, J. Choke or Shine? Quantifying Soccer Players' Abilities to Perform Under Mental Pressure. In *13th MIT Sloan Sports Analytics Conference* (2019), pp. 1–25.

[18] Bransen, L., and Van Haaren, J. Measuring Football Players' On-the-Ball Contributions from Passes During Games. In *Proceedings of the Machine Learning and Data Mining for Sports Analytics Workshop at ECML/PKDD* (2018), pp. 3–15.

[19] Bransen, L., and Van Haaren, J. Player Chemistry: Striving for a Perfectly Balanced Soccer Team. In *MIT Sloan Sports Analytics Conference* (2020), pp. 1–25.

[20] Bransen, L., Van Haaren, J., and van de Velden, M. Measuring soccer players' contributions to chance creation by valuing their passes. *Journal of Quantitative Analysis in Sports 15*, 2 (2019), 97–116.

[21] Brefeld, U., Davis, J., de Jong, L., and Kovalchik, S. Computational Approaches to Strategy and Tactics in Sports (Dagstuhl Seminar 24081). *Dagstuhl Reports 14*, 2 (2024), 164–181.

[22] Breiman, L. Random Forests. *Machine learning 45*, 1 (2001), 5–32.

[23] Brier, G. W. Verification of forecasts expressed in terms of probability. *Monthly Weather Review 78*, 1 (1950), 1–3.

[24] Buchheit, M., Cholley, Y., and Lambert, P. Psychometric and physiological responses to a preseason competitive camp in the heat with a 6-hour time difference in elite soccer players. *International Journal of Sports Physiology and Performance 11*, 2 (2016), 176–181.

[25] Buldú, J. M., Busquets, J., Martínez, J. H., Herrera-Diestra, J. L., Echegoyen, I., Galeano, J., and Luque, J. Using Network Science to Analyse Football Passing Networks: Dynamics, Space, Time, and the Multilayer Nature of the Game. *Frontiers in Psychology 9* (2018), 1900–1900.

[26] Bunker, R., Fujii, K., Hanada, H., and Takeuchi, I. Supervised sequential pattern mining of event sequences in sport to identify important patterns of play: An application to rugby union. *PLOS ONE 16*, 9 (2021), 1–19.

[27] Caley, M. Premier League Projections and New Expected Goals. Available at `https://cartilagefreecaptain.sbnation.com/2015/10/19/9295905/premier-league-projections-and-new-expected-goals`, 2015.

[28] Carling, C., Williams, A., and Reilly, T. *Handbook of Soccer Match Analysis: A Systematic Approach to Improving Performance*, First ed. Routledge, 2005.

[29] Cascioli, L., Goldsmith, M., Stradiotti, L., Van Roy, M., Robberechts, P., Wouters, M., and Davis, J. Boot It: A Pragmatic Alternative To Build Up Play. In *StatsBomb Conference* (2024), pp. 1–17.

[30] Cervone, D., D'Amour, A., Bornn, L., and Goldsberry, K. A Multiresolution Stochastic Process Model for Predicting Basketball Possession Outcomes. *Journal of the American Statistical Association 111*, 514 (2016), 585–599.

[31] Chen, T., and Guestrin, C. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016), pp. 785–794.

[32] Cintia, P., Giannotti, F., Pappalardo, L., Pedreschi, D., and Malvaldi, M. The harsh rule of the goals: Data-driven performance indicators for football teams. In *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)* (2015), pp. 1–10.

[33] Clijmans, J., Van Roy, M., and Davis, J. Looking Beyond The Past: Analyzing The Intrinsic Playing Style Of Soccer Teams. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)* (2022), vol. 2035 of *Communications in Computer and Information Science*, Springer, pp. 370–385.

[34] Davis, J., Bransen, L., Devos, L., Jaspers, A., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Methodology and Evaluation in Sports Analytics: Challenges, Approaches, and Lessons Learned. *Machine Learning 113* (2024), 6977–7010.

[35] Davis, J., Bransen, L., Devos, L., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Evaluating Sports Analytics Models: Challenges, Approaches, and Lessons Learned. In *AI Evaluation beyond Metrics Workshop at IJCAI* (2022), vol. 3169, CEUR-WS, pp. 1–11.

[36] Decroos, T., Bransen, L., Van Haaren, J., and Davis, J. Actions Speak Louder Than Goals: Valuing Player Actions in Soccer. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2019), pp. 1851–1861.

[37] Decroos, T., and Davis, J. Player Vectors: Characterizing Soccer Players' Playing Style from Match Event Streams. In *Machine Learning and Knowledge Discovery in Databases: European Conference (ECML PKDD)* (2019), Springer-Verlag, p. 569–584.

[38] Decroos, T., and Davis, J. Interpretable Prediction of Goals in Soccer. In *Artificial Intelligence in Team Sports Workshop at AAAI* (2020), pp. 1–8.

[39] Decroos, T., Dzyuba, V., Van Haaren, J., and Davis, J. Predicting Soccer Highlights from Spatio-Temporal Match Event Streams. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence* (2017), pp. 1302–1308.

[40] Decroos, T., Van Haaren, J., and Davis, J. Automatic Discovery of Tactics in Spatio-Temporal Soccer Match Data. In *Proceedings of the*

24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2018), pp. 223–232.

[41] Decroos, T., Van Haaren, J., Dzyuba, V., and Davis, J. STARSS: A spatio-temporal action rating system for soccer. In *Machine Learning and Data Mining for Sports Analytics Workshop at ECML/PKDD* (2017), vol. 1971, pp. 11–20.

[42] Decroos, T., Van Roy, M., and Davis, J. SoccerMix: Representing Soccer Actions with Mixture Models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)* (2020), vol. 12461 of *Lecture Notes in Computer Science*, Springer, pp. 459–474.

[43] Devos, L., Meert, W., and Davis, J. Versatile Verification of Tree Ensembles. In *Proceedings of the 38th International Conference on Machine Learning* (2021), vol. 139 of *Proceedings of Machine Learning Research*, PMLR, pp. 2654–2664.

[44] Devos, L., Perini, L., Meert, W., and Davis, J. Detecting Evasion Attacks in Deployed Tree Ensembles. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)* (2023), vol. 14173 of *Lecture Notes in Computer Science*, Springer Nature Switzerland, pp. 120–136.

[45] Dick, U., and Brefeld, U. Learning to Rate Player Positioning in Soccer. *Big Data 7*, 1 (2019), 71–82.

[46] Dick, U., Tavakol, M., and Brefeld, U. Rating Player Actions in Soccer. *Frontiers in Sports and Active Living 3*, 682986 (2021).

[47] Dong, J. S., Jiang, K., Liu, Z., Dong, C., Hou, Z., Hundal, R. S., Guo, J., and Lin, Y. Sports Analytics Using Probabilistic Model Checking and Deep Learning. In *Proceedings of the 27th International Conference on Engineering of Complex Computer Systems (ICECCS)* (2023), pp. 7–11.

[48] Dong, J. S., Shi, L., Chuong, L. V. N., Jiang, K., and Sun, J. Sports Strategy Analytics Using Probabilistic Reasoning. In *Proceedings of the 20th International Conference on Engineering of Complex Computer Systems (ICECCS)* (2015), pp. 182–185.

[49] Fernández, J., and Bornn, L. SoccerMap: A Deep Learning Architecture for Visually-Interpretable Analysis in Soccer. In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track (ECML/PKDD)* (2021), vol. 12461 of *Lecture Notes in Computer Science*, Springer, pp. 491–506.

[50] Fernández, J., Cervone, D., and Bornn, L. Decomposing the Immeasurable Sport: A deep learning expected possession value framework for soccer. In *13th MIT Sloan Sports Analytics Conference* (2019), pp. 1–20.

[51] Fernández, J., Bornn, L., and Cervone, D. A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. *Machine Learning 110*, 6 (2021), 1389–1427.

[52] Forejt, V., Kwiatkowska, M., Norman, G., and Parker, D. 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* (2011), pp. 53–113.

[53] Gabbett, T. J., and Ullah, S. Relationship Between Running Loads and Soft-Tissue Injury in Elite Team Sport Athletes. *Journal of Strength and Conditioning Research 26*, 4 (2012), 953–960.

[54] Garrido, D., Antequera, D., Busquets, J., López del Campo, R., Resta Serra, R., Jos Vielcazat, S., and Buldú, J. Consistency and identifiability of football teams: a network science perspective. *Scientific Reports 10*, 1 (2020), 19735–19735.

[55] Goes, F. R., Brink, M. S., Elferink-Gemser, M. T., Kempe, M., and Lemmink, K. A. The tactics of successful attacks in professional association football: large-scale spatiotemporal analysis of dynamic subgroups using position tracking data. *Journal of Sports Sciences 39*, 5 (2021), 523–532.

[56] Goldman, M., and Rao, J. M. Misperception of Risk and Incentives by Experienced Agents. Available at: `https://ssrn.com/abstract=2435551`, 2014.

[57] Goldner, K. A Markov Model of Football: Using Stochastic Processes to Model a Football Drive. *Journal of Quantitative Analysis in Sports 8* (2012), 1–18.

[58] Goodfellow, I., Bengio, Y., and Courville, A. *Deep Learning*. MIT Press, 2016.

[59] Green, S. Assessing The Performance of Premier League Goalscorers. Available at `https://www.statsperform.com/resource/assessing-the-performance-of-premier-league-goalscorers/`, 2012.

[60] Grehaigne, J.-F., Godbout, P., and Zerai, Z. How the "rapport de forces" evolves in a soccer match: The dynamics of collective decisions in a complex system. *Revista de Psicología del Deporte 20* (2011), 747–765.

[61] Gurpinar-Morgan, W. Passing Out at the Back. Available at: `https://www.statsperform.com/resource/passing-out-at-the-back/`, 2019.

[62] Gyarmati, L., and Anguera, X. Automatic Extraction of the Passing Strategies of Soccer Teams. arXiv preprint arXiv:1508.02171, 2015.

[63] Gyarmati, L., Kwak, H., and Rodriguez, P. Searching for a Unique Style in Soccer. arXiv preprint arXiv:1409.0308, 2014.

[64] Gyarmati, L., and Stanojevic, R. QPass: a Merit-based Evaluation of Soccer Passes. In *Large-Scale Sports Analytics Workshop at KDD* (2016), pp. 1–4.

[65] Hansson, H., and Jonsson, B. A logic for reasoning about time and reliability. *Formal Aspects of Computing 6* (1995), 512–535.

[66] Hei Ho, C., and McKinley, E. Using k-means to learn what soccer passing tells us about playing styles. Available at: `https://www.americansocceranalysis.com/home/2019/3/11/using-k-means-to-learn-what-soccer-passing-tells-us-about-playing-styles`, 2019.

[67] Heiter, E., Martens, L., Seurinck, R., Guilliams, M., Bie, T. D., Saeys, Y., and Lijffijt, J. Pattern or Artifact? Interactively Exploring Embedding Quality with TRACE. In *Machine Learning and Knowledge Discovery in Databases. Research Track and Demo Track (ECML/PKDD)* (2024), Lecture Notes in Computer Science, Springer, pp. 379–382.

[68] Hensel, C., Junges, S., Katoen, J.-P., Quatmann, T., and Volk, M. The Probabilistic Model Checker Storm. *International Journal on Software Tools for Technology Transfer 24*, 4 (2022), 589—610.

[69] Hirotsu, N., Masui, Y., Shimasaki, Y., and Yoshimura, M. Modelling tactical changes in association football using a Markov game. *IMA Journal of Management Mathematics 35*, 3 (2024), 503–535.

[70] Hirotsu, N., and Wright, M. 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 (2002), 88–96.

[71] Hooker, G. Diagnosing extrapolation: tree-based density estimation. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2004), Association for Computing Machinery, p. 569–574.

[72] Hvattum, L. M. Offensive and Defensive Plus–Minus Player Ratings for Soccer. *Applied Sciences 10*, 20 (2020).

[73] Ijtsma, S. 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, 2015.

[74] J. Theiner and W. Gritz and E. Muller-Budack and R. Rein and D. Memmert and R. Ewerth. Extraction of Positional Player Data from Broadcast Soccer Videos. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (2022), IEEE Computer Society, pp. 1463–1473.

[75] Jaspers, A., Brink, M., Probst, S., Frencken, W., and Helsen, W. Relationships between training load indicators and training outcomes in professional soccer. *Sports Medicine 47*, 3 (2017), 533–544.

[76] Johnson, N. Extracting Player Rracking Data from Video Using Non-Stationary Cameras and a Combination of Computer Vision Techniques. In *MIT Sloan Sports Analytics Conference* (2020), pp. 1–14.

[77] Kemeny, J. G., and Snell, J. L. *Finite Markov Chains*. Undergraduate Texts in Mathematics. Springer, 1976.

[78] Kharrat, T., McHale, I. G., and Peña, J. L. Plus–minus player ratings for soccer. *European Journal of Operational Research 283*, 2 (2020), 726–736.

[79] Knutson, T. StatsBomb Release Expected Goals with Shot Impact Height. Available at `https://statsbomb.com/2020/07/statsbomb-release-expected-goals-with-shot-impact-height/`, 2020.

[80] Kovalchik, S. A. Player Tracking Data in Sports. *Annual Review of Statistics and Its Application 10*, 1 (2023), 677–697.

[81] Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., and Blei, D. M. Automatic Differentiation Variational Inference. *Journal of Machine Learning Research 18*, 1 (2017), 430–474.

[82] Kullowatz, M. Goals Added: Deep Dive Methodology. Available at: `https://www.americansocceranalysis.com/home/2020/5/4/goals-added-deep-dive-methodology`, 2020.

[83] Kwiatkowska, M., Norman, G., and Parker, D. PRISM 4.0: Verification of Probabilistic Real-time Systems. In *Proceedings of the 23rd International Conference on Computer Aided Verification (CAV)* (2011), vol. 6806 of *LNCS*, pp. 585–591.

[84] Laurie, E., and Fenner, P. How Football Actually Works. Available at: `https://www.youtube.com/watch?v=1W4bRAKfMeM`, 2021.

[85] Lawrence, T. Introducing xGChain and xGBuildup. Available at `https://statsbomb.com/articles/soccer/introducing-xgchain-and-xgbuildup/`, 2018.

[86] Link, D., Lang, S., and Seidenschwarz, P. Real Time Quantification of Dangerousity in Football Using Spatiotemporal Tracking Data. *PLOS ONE 11*, 12 (2016).

[87] Liu, G., Luo, Y., Schulte, O., and Kharrat, T. Deep soccer analytics: Learning an action-value function for evaluating soccer players. *Data Mining and Knowledge Discovery 34*, 5 (2020), 1531–1559.

[88] Liu, G., and Schulte, O. Deep Reinforcement Learning in Ice Hockey for Context-Aware Player Evaluation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence* (2018), pp. 3442–3448.

[89] 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., and Heess, N. From motor control to team play in simulated humanoid football. *Science Robotics 7*, 69 (2022), 1–16.

[90] Liu, T., and Hohmann, A. 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 (2013), 79–86.

[91] Liu, Z., Jiang, K., Hou, Z., Lin, Y., and Dong, J. S. Insight Analysis for Tennis Strategy and Tactics. In *IEEE International Conference on Data Mining (ICDM)* (2023), pp. 1169–1174.

[92] Llana, S., Burriel, B., Madrero, P., and Fernández, J. Is it worth the effort? Understanding and contextualizing physical metrics in soccer. arXiv preprint arXiv:2204.02313, 2022.

[93] Lucey, P., Bialkowski, A., Mofort, M., Carr, P., and Matthews, I. "Quality vs Quantity": Improved Shot Prediction In Soccer Using Strategic Features From Spatiotemporal Data. In *9th MIT Sloan Sports Analytics Conference* (2015), pp. 1–9.

[94] Lucey, P., Oliver, D., Carr, P., Roth, J., and Matthews, I. Assessing team strategy using spatiotemporal data. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2013), Association for Computing Machinery, p. 1366–1374.

[95] Luo, Y., Schulte, O., and Poupart, P. Inverse Reinforcement Learning for Team Sports: Valuing Actions and Players. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence* (7 2020), pp. 3356–3363.

[96] Lövdal, S. S., Den Hartigh, R. J., and Azzopardi, G. Injury Prediction in Competitive Runners With Machine Learning. *International Journal of Sports Physiology and Performance 16*, 10 (2021), 1522–1531.

[97] Mackay, N. Improving my 'xG added' model. Available at: `https://mackayanalytics.nl/2017/07/28/improving-my-xg-added-model/`, 2017.

[98] McHale, I., and Scarf, P. Modelling soccer matches using bivariate discrete distributions with general dependence structure. *Statistica Neerlandica 61*, 4 (2007), 432–445.

[99] McHale, I., Scarf, P., and Folker, D. On the development of a soccer player performance rating system for the English Premier League. *Interfaces 42*, 4 (2012), 339–351.

[100] McHale, I. G., and Relton, S. D. Identifying key players in soccer teams using network analysis and pass difficulty. *European Journal of Operational Research 268*, 1 (2018), 339–347.

[101] McInnes, L., Healy, J., Saul, N., and Grossberger, L. UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software 3*, 29 (2018), 861.

[102] Merhej, C., Beal, R. J., Matthews, T., and Ramchurn, S. What Happened Next? Using Deep Learning to Value Defensive Actions in Football Event-Data. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (2021), Association for Computing Machinery, p. 3394–3403.

[103] Merlin, M., Cunha, S. A., Moura, F. A., Torres, R. d. S., Gonçalves, B., and Sampaio, J. Exploring the determinants of success in different clusters of ball possession sequences in soccer. *Research in Sports Medicine 28*, 3 (2020), 339–350.

[104] Michalczyk, K. An attempt to extend xG gain. Available at: `https://kubamichalczyk.github.io/2018/11/18/An-attempt-to-extend-xG-gain.html`, 2018.

[105] Michalczyk, K. A glance at building out from the back. In *Opta Pro Forum, London* (2019).

[106] Moore, J. Pass Chain Analysis: Should Teams Build From The Back? Available at: `https://www.americansocceranalysis.com/home/2018/7/3/gdomsoqkfxu27tatucddmu2a55u2pw`, 2018.

[107] Oruç, D. C., Cascioli, L., Stradiotti, L., Van Roy, M., Robberechts, P., and Davis, J. An Analysis of the Influence of Game Context on Team Playing Style. In *StatsBomb Conference* (2023), pp. 1–17.

[108] Oruç, D. C., Cascioli, L., Stradiotti, L., Van Roy, M., Robberechts, P., and Davis, J. An Analysis of the Influence of Game Context on Team Playing Style. In *Machine Learning and Data Mining for Sports Analytics Workshop at ECML/PKDD* (2024), Communications in Computer and Information Science, Springer Nature Switzerland, pp. 1–17.

[109] Pappalardo, L., Cintia, P., Ferragina, P., Massucco, E., Pedreschi, D., and Giannotti, F. PlayeRank: Data-driven Performance Evaluation and Player Ranking in Soccer via a Machine Learning Approach. *ACM Transactions on Intelligent Systems and Technology 10*, 5 (2019), 59:1–59:27.

[110] Pappalardo, L., Cintia, P., Rossi, A., Massucco, E., Ferragina, P., Pedreschi, D., and Giannotti, F. A public data set of spatio-temporal match events in soccer competitions. *Scientific Data 6*, 1 (2019), 1–15.

[111] Passos, P., Davids, K., Araújo, D., Paz, N., Minguéns, J., and Mendes, J. Networks as a novel tool for studying team ball sports as complex social systems. *Journal of Science and Medicine in Sport 14*, 2 (2011), 170–176.

[112] Pearl, J., and Mackenzie, D. *The Book of Why*. Basic Books, 2018.

[113] Peña, J. L., and Touchette, H. A network theory analysis of football strategies. arXiv preprint arXiv:1206.6904, 2012.

[114] Pfaff, C., Hunter, E., Hong, H., Forestell, D., Fialkov, A., Drassinower, Z., and Chan, T. Controlling ball progression in soccer. arXiv preprint arXiv:2210.16474, 2022.

[115] Pfeiffer, M., Zhang, H., and Hohmann, A. A Markov Chain Model of Elite Table Tennis Competition. *International Journal of Sports Science & Coaching 5*, 2 (2010), 205–222.

[116] Power, P., Ruiz, H., Wei, X., and Lucey, P. Not All Passes Are Created Equal: Objectively Measuring the Risk and Reward of Passes in Soccer from Tracking Data. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2017), pp. 1605—1613.

[117] Rein, R., and Memmert, D. Big data and tactical analysis in elite soccer: future challenges and opportunities for sports science. *SpringerPlus 5*, 1 (2016), 1410–1410.

[118] Riley, P., and Veloso, M. Advice Generation from Observed Execution: Abstract Markov Decision Process Learning. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI)* (2004), pp. 631—636.

[119] Robberechts, P., and Davis, J. How data availability affects the ability to learn good xG models. In *Machine Learning and Data Mining for Sports Analytics Workshop at ECML/PKDD* (2020), pp. 17–27.

[120] Robberechts, P., Van Haaren, J., and Davis, J. A Bayesian Approach to In-Game Win Probability in Soccer. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2021), Association for Computing Machinery, p. 3512–3521.

[121] Robberechts, P., Van Roy, M., and Davis, J. un-xPass: Measuring Soccer Player's Creativity. In *StatsBomb Conference* (2022), pp. 1–21.

[122] Robberechts, P., Van Roy, M., and Davis, J. un-xPass: Measuring Soccer Player's Creativity. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2023), Association for Computing Machinery, p. 4768–4777.

[123] Rossi, A., Pappalardo, L., Cintia, P., Iaia, F. M., Fernàndez, J., and Medina, D. Effective injury forecasting in soccer with GPS training data and machine learning. *PLOS ONE 13*, 7 (2018), 1–15.

[124] Routley, K., and Schulte, O. A Markov Game Model for Valuing Player Actions in Ice Hockey. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence* (2015), pp. 782–791.

[125] Rudd, S. A Framework for Tactical Analysis and Individual Offensive Production Assessment in Soccer Using Markov Chains. In *New England Symposium on Statistics in Sports* (2011).

[126] Russell, S., and Norvig, P. *Artificial Intelligence: A Modern Approach*, Fourth ed. Pearson series in Artificial Intelligence. Pearson, 2022.

[127] Saebo, O. D., and Hvattum, L. M. Modelling the financial contribution of soccer players to their clubs. *Journal of Sports Analytics 5*, 1 (2018), 23–34.

[128] Salvatier, J., Wiecki, T. V., and Fonnesbeck, C. Probabilistic programming in Python using PyMC3. *PeerJ Computer Science 2* (2016), e55.

[129] Sandholtz, N., and Bornn, L. Replaying the NBA. In *12th MIT Sloan Sports Analytics Conference* (2018), pp. 1–13.

[130] Sandholtz, N., and Bornn, L. Markov decision processes with dynamic transition probabilities: An analysis of shooting strategies in basketball. *The Annals of Applied Statistics 14*, 3 (2020), 1122–1145.

[131] Schulte, O., Khademi, M., Gholami, S., Zhao, Z., Javan, M., and Desaulniers, P. A Markov Game model for valuing actions, locations, and team performance in ice hockey. *Data Mining and Knowledge Discovery 31*, 6 (2017), 1735–1757.

[132] Shapiro, L. G., and Stockman, G. C. *Computer Vision*. Prentice Hall, 2001.

[133] Singh, K. Introducing Expected Threat. Available at `https://karun.in/blog/expected-threat.html`, 2019.

[134] Spearman, W. Beyond Expected Goals. In *12th MIT Sloan Sports Analytics Conference* (2018), pp. 1–17.

[135] Stan Development Team. Stan Modeling Language Users Guide and Reference Manual. Available at: `https://mc-stan.org`, 2024.

[136] StatsBomb. Introducing On-Ball Value. Available at `https://statsbomb.com/articles/soccer/introducing-on-ball-value-obv/`, 2021.

[137] Tureen, T., and Olthof, S. Estimated Player Impact (EPI): Quantifying the Effects of Individual Players on Football (Soccer) Actions Using Hierarchical Statistical Models. In *StatsBomb Conference* (2022), pp. 1–29.

[138] Tuyls, 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., and Hassabis, D. Game Plan: What AI Can Do for Football, and What Football Can Do for AI. *Journal of Artificial Intelligence Research 71* (2021), 41–88.

[139] van der Maaten, L., and Hinton, G. Visualizing Data using t-SNE. *Journal of Machine Learning Research 9*, 86 (2008), 2579–2605.

[140] Van Haaren, J. "Why Would I Trust Your Numbers?" On the Explainability of Expected Values in Soccer. In *AI for Sports Analytics Workshop at IJCAI* (2021), pp. 1–7.

[141] Van Haaren, J., Dzyuba, V., Hannosset, S., and Davis, J. Automatically Discovering Offensive Patterns in Soccer Match Data. In *International Symposium on Intelligent Data Analysis (Lecture Notes in Computer Science)* (2015), vol. 9385, Springer, pp. 286–297.

[142] Van Haaren, J., Robberechts, P., Decroos, T., Bransen, L., and Davis, J. Analysing Performance and Playing Style Using Ball Event Data. In *Football Analytics: Now and Beyond. A Deep Dive into the Current State of Advanced Data Analytics*. Barça Innovation Hub, 2019, pp. 36–47.

[143] Van Roy, M., Cascioli, L., and Davis, J. ETSY: A Rule-Based Approach to Event and Tracking Data SYnchronization. In *Machine Learning and Data Mining for Sports Analytics Workshop at ECML/PKDD* (2023), vol. 2035 of *Communications in Computer and Information Science*, Springer Nature Switzerland, pp. 11–23.

[144] Van Roy, M., and Davis, J. DataDebugging: Enhancing Trust in Soccer Action-Value Models by Contextualization. In *Proceedings of 13th World Congress of Performance Analysis of Sport and 13th International Symposium on Computer Science in Sport* (2023), vol. 1448 of *Advances in Intelligent Systems and Computing*, Springer Nature Switzerland, pp. 193–196.

[145] Van Roy, M., Robberechts, P., and Davis, J. Optimally Disrupting Opponent Build-ups. In *StatsBomb Conference* (2021), pp. 1–16.

[146] Van Roy, M., Robberechts, P., and Davis, J. A Markov Model Approach to Analyzing Build-Up Efficiency and Defensive Setups in Soccer. *Submitted to Data Mining and Knowledge Discovery* (2024).

[147] Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. DTAI Sports Analytics Lab Blog. Available at `https://dtai.cs.kuleuven.be/sports/blog/valuing-on-the-ball-actions-in-soccer:-a-critical-comparison-of-xt-and-vaep`, 2019.

[148] Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. In *Artificial Intelligence in Team Sports Workshop at AAAI* (2020), pp. 1–8.

[149] Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. Learning a Markov Model for Evaluating Soccer Decision Making. In *Reinforcement Learning for Real Life Workshop at ICML* (2021), pp. 1–10.

[150] Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. Leaving Goals on the Pitch: Evaluating Decision Making in Soccer. In *15th MIT Sloan Sports Analytics Conference* (2021), pp. 1–25.

[151] Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. A Markov Framework for Learning and Reasoning About Strategies in Professional Soccer. *Journal of Artificial Intelligence Research 77* (2023), 517–562.

[152] Van Roy, M., Yang, W.-C., De Raedt, L., and Davis, J. Analyzing Learned Markov Decision Processes Using Model Checking for Providing Tactical Advice in Professional Soccer. In *AI for Sports Analytics Workshop at IJCAI* (2021), pp. 1–7.

[153] Vanrenterghem, J and Nedergaard, NJ and Robinson, MA and Drust, B. Training Load Monitoring in Team Sports: A Novel Framework Separating Physiological and Biomechanical Load-Adaptation Pathways. *Sports Medicine 47*, 11 (2017), 2135–2142.

[154] Walid, A. How are football stats collected? Visiting StatsBomb's data centre. Available at: `https://www.nytimes.com/athletic/4263797/2023/03/24/xg-statsbomb-inspired-by-you/`, 2023.

[155] Wang, J., Fox, I., Skaza, J., Linck, N., Singh, S., and Wiens, J. The Advantage of Doubling: A Deep Reinforcement Learning Approach to Studying the Double Team in the NBA. In *12th MIT Sloan Sports Analytics Conference* (2018), pp. 1–12.

[156] Wang, Q., Zhu, H., Hu, W., Shen, Z., and Yao, Y. Discerning Tactical Patterns for Professional Soccer Teams: An Enhanced Topic Model with Applications. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2015), pp. 2197–2206.

[157] Wenninger, S., and Lames, M. Performance Analysis in Table Tennis - Stochastic Simulation by Numerical Derivation. *International Journal of Computer Science in Sport 15*, 1 (2016), 22–36.

[158] Whitmore, J. Evolving our possession value framework. Available at: `https://www.statsperform.com/resource/evolving-our-possession-value-framework/`, 2020.

[159] Whitmore, J. Slice It Up: Introducing Opta Player Radars. Available at: `https://theanalyst.com/eu/2023/06/introducing-opta-radars-compare-players/`, 2023.

[160] Wildman, L. Probabilistic Model Checking of Handball Action Sequences. In *2023 IEEE 28th Pacific Rim International Symposium on Dependable Computing (PRDC)* (2023), pp. 332–336.

[161] Xie, X., Wang, J., Liang, H., Deng, D., Cheng, S., Zhang, H., Chen, W., and Wu, Y. PassVizor: Toward Better Understanding of the Dynamics of Soccer Passes. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (2021), 1322–1331.

[162] Yam, D. Attacking Contributions: Markov Models for Football. Available at `https://statsbomb.com/2019/02/attacking-contributions-markov-models-for-football/`, 2019.

Statement on the Use of Generative AI

I did not use generative AI assistance tools during the research/writing process of my thesis, except for mere language assistance.

The text, code, and images in this thesis are my own (unless otherwise specified). Generative AI has only been used in accordance with the KU Leuven guidelines and appropriate references have been added. I have reviewed and edited the content as needed and I take full responsibility for the content of the thesis.

Curriculum Vitae

Maaike Van Roy was born in Leuven on August 19, 1996. She obtained a Bachelor's degree in Informatics (2017) and a Masters degree in Engineering in Computer Science (2019) at KU Leuven. In October 2019, Maaike started her doctoral studies at KU Leuven in the DTAI research group under the supervision of Prof. Dr. Jesse Davis.

Maaike has published papers at both technical (ECML/PKDD, KDD, JAIR, Machine Learning) and sports-oriented conferences and journals (IACSS&ISPAS, MIT Sloan Sports Analytics Conference, StatsBomb Conference), and received 3rd place for the Routledge Young Researcher Award at the 2022 IACSS&ISPAS Joint Conference. Additionally, her work has been picked up by various media. It was covered in an interview with John Muller for FiveThirtyEight and in articles by ESPN and the New York Times.

Publications

Journal Papers

Van Roy, M., Robberechts, P., and Davis, J. A Markov Model Approach to Analyzing Build-Up Efficiency and Defensive Setups in Soccer. *Submitted to Data Mining and Knowledge Discovery* (2024)

Davis, J., Bransen, L., Devos, L., Jaspers, A., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Methodology and Evaluation in Sports Analytics: Challenges, Approaches, and Lessons Learned. *Machine Learning 113* (2024), 6977–7010

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. A Markov Framework for Learning and Reasoning About Strategies in Professional Soccer. *Journal of Artificial Intelligence Research 77* (2023), 517–562

Conference Papers

Robberechts, P., Van Roy, M., and Davis, J. un-xPass: Measuring Soccer Player's Creativity. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2023), Association for Computing Machinery, p. 4768–4777

Van Roy, M., and Davis, J. DataDebugging: Enhancing Trust in Soccer Action-Value Models by Contextualization. In *Proceedings of 13th World Congress of Performance Analysis of Sport and 13th International Symposium on Computer Science in Sport* (2023), vol. 1448 of *Advances in Intelligent Systems and Computing*, Springer Nature Switzerland, pp. 193–196
★ *This paper received 3rd place for the Routledge Young Researcher Award.*

Clijmans, J., Van Roy, M., and Davis, J. Looking Beyond The Past: Analyzing The Intrinsic Playing Style Of Soccer Teams. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)* (2022), vol. 2035 of *Communications in Computer and Information Science*, Springer, pp. 370–385

Decroos, T., Van Roy, M., and Davis, J. SoccerMix: Representing Soccer Actions with Mixture Models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)* (2020), vol. 12461 of *Lecture Notes in Computer Science*, Springer, pp. 459–474

Conference Papers without Proceedings

Cascioli, L., Goldsmith, M., Stradiotti, L., Van Roy, M., Robberechts, P., Wouters, M., and Davis, J. Boot It: A Pragmatic Alternative To Build Up Play. In *StatsBomb Conference* (2024), pp. 1–17

Oruç, D. C., Cascioli, L., Stradiotti, L., Van Roy, M., Robberechts, P., and Davis, J. An Analysis of the Influence of Game Context on Team Playing Style. In *StatsBomb Conference* (2023), pp. 1–17

Robberechts, P., Van Roy, M., and Davis, J. un-xPass: Measuring Soccer Player's Creativity. In *StatsBomb Conference* (2022), pp. 1–21

Van Roy, M., Robberechts, P., and Davis, J. Optimally Disrupting Opponent Build-ups. In *StatsBomb Conference* (2021), pp. 1–16

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. Leaving Goals on the Pitch: Evaluating Decision Making in Soccer. In *15th MIT Sloan Sports Analytics Conference* (2021), pp. 1–25

Workshop Papers

Oruç, D. C., Cascioli, L., Stradiotti, L., Van Roy, M., Robberechts, P., and Davis, J. An Analysis of the Influence of Game Context on Team Playing Style. In *Machine Learning and Data Mining for Sports Analytics Workshop at ECML/PKDD* (2024), Communications in Computer and Information Science, Springer Nature Switzerland, pp. 1–17

Van Roy, M., Cascioli, L., and Davis, J. ETSY: A Rule-Based Approach to Event and Tracking Data SYnchronization. In *Machine Learning and Data*

Mining for Sports Analytics Workshop at ECML/PKDD (2023), vol. 2035 of *Communications in Computer and Information Science*, Springer Nature Switzerland, pp. 11–23

Davis, J., Bransen, L., Devos, L., Meert, W., Robberechts, P., Van Haaren, J., and Van Roy, M. Evaluating Sports Analytics Models: Challenges, Approaches, and Lessons Learned. In *AI Evaluation beyond Metrics Workshop at IJCAI* (2022), vol. 3169, CEUR-WS, pp. 1–11

Van Roy, M., Yang, W.-C., De Raedt, L., and Davis, J. Analyzing Learned Markov Decision Processes Using Model Checking for Providing Tactical Advice in Professional Soccer. In *AI for Sports Analytics Workshop at IJCAI* (2021), pp. 1–7

Van Roy, M., Robberechts, P., Yang, W.-C., De Raedt, L., and Davis, J. Learning a Markov Model for Evaluating Soccer Decision Making. In *Reinforcement Learning for Real Life Workshop at ICML* (2021), pp. 1–10

Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP. In *Artificial Intelligence in Team Sports Workshop at AAAI* (2020), pp. 1–8

FACULTY OF ENGINEERING SCIENCE
DEPARTMENT OF COMPUTER SCIENCE
Celestijnenlaan 200A box 2402
3001 Leuven
maaike.vanroy@kuleuven.be
https://maaikevr.github.io/

