UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

Models and Algorithms to Optimise Team Performance in Football

by

Gregory Alan Everett

Co-Supervisors:
Prof. Sarvapali Ramchurn
Prof. Timothy Norman
ORCiD: 0000-0003-0949-1689

A thesis submitted for the degree of
Doctor of Philosophy

March 2026

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

Doctor of Philosophy

Models and Algorithms to Optimise Team Performance in Football

by Gregory Alan Everett

With the growing use of Artificial Intelligence (AI) and machine learning, there is increasing potential to model and optimise team behaviour across important domains such as disaster response, security, and team sports. These domains are inherently spatiotemporal, requiring models that capture both spatial and temporal dynamics. This thesis focuses on football, a complex, dynamic sport with rich spatiotemporal data and clear objectives, making it an ideal testbed for developing and validating team-based AI models. Moreover, football analytics is a rapidly expanding industry, with European clubs generating €38 billion in revenue during the 2023/24 season, driving the demand for models that provide competitive advantages.

This thesis proposes a number of novel methods that utilise spatiotemporal data to advance team prediction, analysis, and decision-making in football and other team-based domains. In particular, we introduce a spatiotemporal agent behaviour imputation model that reduces predictive error by 62% compared to baselines in limited observability settings, significantly improving the accessibility of off-ball football analytics through imputed tracking data. We also present a novel spatial teamwork model, combining Monte Carlo tree search (MCTS) and linear programming, that optimises agent decision-making in real-world football defence, reducing opponent threat by 24%. In addition, we develop new metrics, derived from a graph attention network (GAT), to assign credit to indirect agent contributions in team-based defence. The GAT model predicts football passes with a 6% reduction in loss compared to baselines, and we show how these new metrics can greatly improve off-ball football player evaluation. Finally, we propose a dynamic team formation and agent replacement model that accounts for agent fatigue and unavailability and optimises decision-making using a multi-step MCTS algorithm. Applied to football team selection and substitutions, this model improves long-term team performance by 1% and reduces first-team injuries by 15%. This thesis also highlights key remaining challenges for AI in football, including the use of richer player data (e.g., body pose) and greater use of explainable models to build trust in clubs.

Contents

List of Figures

List of Tables

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:
   (a) Gregory Everett, Ryan J Beal, Tim Matthews, Joseph Early, Timothy J Norman, and Sarvapali D Ramchurn. Inferring player location in sports matches: Multi-agent spatial imputation from limited observations. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 1643–1651, 2023b
   (b) Gregory Everett, Ryan Beal, Tim Matthews, Timothy J Norman, and Sarvapali D Ramchurn. The strain of success: A predictive model for injury risk mitigation and team success in soccer. In *Proceedings of the 18th MIT Sloan Sports Analytics Conference*, 2024

(c) Gregory Everett, Ryan J Beal, Tim Matthews, Timothy J Norman, and Sarvapali D Ramchurn. Optimising spatial teamwork under uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23168–23176, 2025b

(d) Gregory Everett, Ryan J Beal, Tim Matthews, Timothy J Norman, and Sarvapali D Ramchurn. Evaluating defensive influence in multi-agent systems using graph attention networks. In *IEEE 12th International Conference on Data Science and Advanced Analytics*, 2025a

Signed:.. Date:..................

Acknowledgements

Firstly, I would like to thank my PhD supervisors, Professor Gopal Ramchurn and Professor Timothy Norman, for their invaluable help, advice and support throughout my studies. Their input during our regular meetings provided tremendous insight and expertise at all stages of my PhD. I would also like to thank Dr. Ryan Beal and Dr. Tim Matthews (of Sentient Sports) for their support throughout my PhD. I am deeply appreciative of the consistent attention you gave my work, and of your willingness to share advice, expertise, and insight, especially on how the research applies to football. In addition, thank you to everyone in the AIC group and the IRIDIS support team.

Sincere thanks are extended to Bepro Group Ltd, Gradient Sports, and Hudl Statsbomb for generously providing the data resources that supported the projects presented in this thesis. I am also grateful to my fellow PhD student, Dr. Joseph Early, for collaborating with me on a paper during the early stages of my PhD and for your support and guidance throughout the PhD journey.

I have been fortunate to meet many close friends during my years in Southampton, and I would like to thank them for their encouragement and wonderful friendship. From my undergraduate days, my thanks goes to Ewan, Lauren, Ellen, Lucy, Mat, Michael, Alex and Remi. During my PhD, I would like to thank my football teammates and housemates, whose friendship made the PhD challenges much more enjoyable. In particular, I would like to thank Lena and Hannah, for always offering a kind ear and for our countless PhD chats over games of cards in the past two years.

I am hugely appreciative of my oldest and dearest friends, including all those in my hometown of Bordon, who help remind me of life outside of Southampton. In particular, I'd like to thank Anna, Kieran M and Kieran S for their regular check-ins and unwavering friendship. Overall, I'd like to say I am incredibly proud of all my friends, both in and out of Southampton, for all that they have achieved.

Finally, thank you to my family for their unwavering love and encouragement throughout my PhD and indeed, my entire life. To Katrina, Jon, Steve, Mollie and Alfie, I am incredibly thankful. In particular, I'd like to express my appreciation to my Mum and Dad. You have always been there for me and pushed me to be the best that I can be. From counting pencils on the bed, to buying me my first Python coding book, all the way to my PhD research. Making you proud remains my greatest motivator.

I would like to acknowledge the use of AI-powered tools, specifically ChatGPT and Grammarly, for assistance with grammar, spelling, and stylistic improvements during the editing of this thesis. All content, arguments, and conclusions remain my own.

To Mum and Dad, for your love and support.

List of Acronyms

ACWR	Acute:Chronic Workload Ratio
AI	Artificial Intelligence
AUC	Area Under the ROC Curve
BCE	Binary Cross-Entropy
CNN	Convolutional Neural Network
CUDA	Compute Unified Device Architecture
CV	Cross-Validation
DI	Defender Influence
DP	Defender Performance
EPL	English Premier League
EPV	Expected Possession Value
GAPP	Graph Attention for Pass Probabilities
GAT	Graph Attention Network
GCN	Graph Convolutional Network
GNN	Graph Neural Network
GPU	Graphics Processing Unit
GRU	Gated Recurrent Unit
HPC	High Performance Computing
IL	Imitation Learning
IRL	Inverse Reinforcement Learning
LLM	Large Language Models
LSTM	Long Short-Term Memory
MCTS	Monte Carlo Tree Search
MDP	Markov Decision Process
ML	Machine Learning
MMDP	Multi-agent Markov Decision Process
ReLU	Rectified Linear Unit
RL	Reinforcement Learning
RNN	Recurrent Neural Network
RWP	Real-World Policy
SHAP	Shapley Additive Explanations

sRPE	Session Rate of Perceived Exertion
SVM	Support Vector Machine
TF	Team Formation
VAE	Variational Autoencoder
VAEP	Valuing Actions by Estimating Probabilities
VRAM	Video Random Access Memory
VRNN	Variational Recurrent Neural Network
xG	Expected Goals
xT	Expected Threat

List of Mathematical Notation

Thesis-wide Notation

The following notation is used throughout the thesis. Additional symbols introduced for specific chapters are listed in the corresponding chapter notation sections.

Core Notation

E	Sequence of events (e.g., on-ball football events)
e	An event (e.g., a pass in football)
T	Number of events in the time series
t	A timestep
x, y	Coordinate locations
B	The ball in football

Multi-Agent Systems

C	Set of agents (e.g., players in football)
N	Number of agents
$\mathbf{\Phi}^C$	Complete set of agent locations across all timesteps
Φ_t^C	Coordinate locations of all agents at timestep t
ϕ_t^c	Coordinate location of agent c at timestep t
ζ^c	Characteristics of an agent c
κ_c	Role of an agent c within the team (e.g., goalkeeper)

Machine Learning and MCTS

b	Batch size
h	Hidden layer size
η_s	MCTS tree node corresponding to a state s
$V(s)$	Value of a state s

Chapter 3: Multi-Agent Spatial Imputation

$\mathbf{M}$	Agent observation mask (binary matrix)
$\mathbf{M}_t^n$	Binary indicator: agent n observed at time t
$\hat{\boldsymbol{\Phi}}^C$	Predicted full set of agent locations across all timesteps
$\hat{\Phi}_t^C$	Predicted full set of agent location at time t
$\hat{\phi}_t^n$	Predicted location for agent n at timestep t
$e_t^{x,y}$	(x, y) location of event e at timestep t
l_{seq}	Length of input sequence
d_{imp}	Feature dimensionality for the *Agent Imputer* model

Chapter 4: Learning and Optimising Spatial Teamwork

Ω_E	Outcome of sequence E (e.g., possession loss, failed attack)
Λ	Two-dimensional Euclidean space (football pitch/plane)
$V_t^{x,y}$	Value of location (x, y) at timestep t
C_α	Defending team
C_β	Attacking team
a_t^c	Action of agent c at time t
$\vartheta_t^{x,y}$	Binary indicator: defenders control location (x, y) at timestep t
$\Pr(\vartheta_t^{x,y} \| \Phi_t^\alpha, \Phi_t^\beta)$	Probability that defenders control location (x, y) at timestep t
$U_t(C_\alpha)$	Utility of team C_α at timestep t
$\Phi_t^{\alpha,*}$	Optimal structure of team C_α at time t to maximise their utility U_t
$U_E(C_\alpha)$	Utility of team C_α across entire event sequence
Θ_t^α	Marginal contribution of agent α
$\boldsymbol{\Psi}$	Set of all agent subteams
Ψ	Subteam of agents
Θ_t^Ψ	Marginal contribution of subteam Ψ
$\mathcal{M}_{\text{MA}}$	MMDP
$s \in S$	State within the MMDP
P	MMDP transition function
R^c	Immediate reward for agent c
γ	MMDP discount factor
$\eta_{\mathbf{s}'}$	Set of child nodes expanded in parallel
$V(\mathbf{s}')$	Estimated values for the set of simulated child states expanded in parallel
ϖ	Number of parallel transitions
α_Ψ	Binary indicator: 1 if α acts individually, 0 if in a subteam
Z	Euclidean space (e.g., football pitch) divided into a grid of $(I \times J)$ zones
$Z_{i,j}$	Specific zone within the grid Z
$V_t^{Z_{i,j}}$	Value of a zone $Z_{i,j}$ at timestep t
B_t^Z	Zone occupied by the ball at timestep t
$\dot{x}_c, \dot{y}_c$	Velocity of an agent c in the x and y direction

Chapter 5: Evaluating Off-Ball Player Contributions

team : $C \to \{1, 2\}$	Function mapping players to teams
C_α	Defending team
C_β	Attacking team
c_t^n	Player performing on-ball action at event e_t
$\Pr(c_{t+1}^n \vert e_t)$	Probability that player c_n will be the next on-ball player at event e_{t+1}
$\mathcal{G}_t$	Graph for event e_t
$\mathcal{N}_t$	Set of graph nodes (players and ball)
ξ_t	Set of directed edges in the graph
$u, v \in \mathcal{N}_t$	Nodes in the graph
X_u	Node feature vector
$Y_{u,v}$	Edge feature vector
$d_\mathcal{N}$	Dimensions of *GAPP* node features
d_ξ	Dimensions of *GAPP* edge features
$f : \mathcal{G} \to [0, 1]$	Function (*GAPP* model) mapping graph to ball reception probability
$\mathbf{X_t}$	Node feature tensor
$\mathbf{Y_t}$	Edge feature tensor
δ_{uv}	Normalised attention weight on edge (u, v)
$\Pr(c_{t+1}^u \vert e_t, \delta_{uv} = 0)$	Attacker reception probability with defender attention removed
DI_{uv}	The influence of defender v on attacker u
DP_v	The defensive performance of defender v
xT_u	The expected threat of an attacker u

Chapter 6: Optimising Short- and Long-Term Team Selection

$\mathbf{G}$	Ordered set of games in a season
K	Number of games in a season
G_k	The k-th game in the season
C_α	Team of agents
C_β	Opposing team of agents
N_α	Number of players in the squad C_α
C_α^+	Starting players in a game
C_α^-	Reserve players in a game
τ	A game timestep, uniform 10-minute intervals
L	Number of timesteps τ in a game
Ω_k	Game outcome (final score, updated player states)
$\Omega_{k,\tau}$	Timestep outcome (score and player status at timestep τ of game G_k)
ψ_α	Historical performance of player α
$\psi_{k,\tau}^\alpha$	In-game performance during game G_k up to a time τ
$\mathcal{I}_{k,\tau}^\alpha$	Binary injury indicator for player α at timestep τ of game G_k
l_k^α	Number of future games missed due to injury for player α
$\mathcal{F}$	Injury duration distribution

ψ_α^*	Fatigue-adjusted player ability
$\Xi_{k,\tau}^{\alpha}$	Fatigue decay
λ	Fatigue scaling factor
$\Pr(\mathcal{I}_{k,\tau}^{\alpha})_{\max}$	Maximum injury probability for player α
D_α, D_β	Team dominance values
$D_{\alpha,k}^{\text{pre}}, D_{\beta,k}^{\text{pre}}$	Dominance of team starting lineups for game G_k
$D_{\alpha,k}^{\text{in}}(\tau), D_{\beta,k}^{\text{in}}(\tau)$	Cumulative in-game dominance up to timestep τ of game G_k
f_Ω^{pre}	Pre-game outcome model
f_Ω^{in}	In-game outcome model
$\mathcal{M}_{\text{pre}}$	Pre-game MDP
$s_k \in S_{\text{pre}}$	State at game G
$a_k \in A_{\text{pre}}$	Pre-game action (team selection at game G)
P_{pre}	Pre-game transition function
R_{pre}	Pre-game reward function (expected points)
γ_{pre}	Pre-game MDP discount factor
$\mathcal{M}_{\text{in}}$	In-game MDP
$s_\tau \in S_{\text{in}}$	State at timestep τ
$a_\tau \in A_{\text{in}}$	In-game action (Substitution)
$\text{score}_{k,\tau}$	Current score of game G_k at timestep τ
$\mathcal{O}_\alpha^+$	Set of players being replaced
$\mathcal{O}_\alpha^-$	Set of players being substituted on
P_{in}	In-game transition function
R_{in}	In-game reward function
γ_{in}	In-game MDP discount factor
$\mathcal{E}_{k,\leq\tau}$	Set of VAEP events in game k up to timestep τ
μ_α, μ_β	Expected team scoring rates
φ	Team dominance decay multiplier
$V_{\text{pre}}(s_k)$	Pre-game value function
$V_{\text{in}}(s_\tau)$	In-game value function
$\mathbb{I}$	Indicator function - equal to one if a condition is satisfied and zero otherwise
α_{sub}	Binary variable indicating if agent α has been substituted
π_α, π_β	Team policies
$\pi_\beta^{\text{pre}}(a_k \mid s_k)$	Pre-game policy (probability distribution over lineups)
$\pi_\beta^{\text{in}}(a_\tau \mid s_\tau)$	In-game substitution policy
$\pi_\beta^{\text{in}}(\lvert\mathcal{O}_\beta^+\rvert \mid s_\tau)$	Policy for number of substitutions
$\pi_\beta^{\text{in}}(\mathcal{O}_\beta^+ \mid s_\tau, \lvert\mathcal{O}^+\rvert)$	Policy for players substituted off
$\pi_\beta^{\text{in}}(\mathcal{O}_\beta^- \mid s_\tau, \mathcal{O}^+)$	Policy for players substituted on
a_s^α	Action selected for team C_α at state s
$Q(s,a)$	Action-value estimate
ϵ_{puct}	PUCT constant for MCTS exploration
W	Size of the explorable action space

Chapter 1

Introduction

Many real-world scenarios involve groups of humans or autonomous entities, referred to as agent teams, working together to achieve shared objectives in complex, and often adversarial, environments. For instance, emergency response teams distribute responsibilities to address incidents quickly (Ramchurn et al., 2010), while security teams protect assets by anticipating and countering adversary attacks (Shieh et al., 2012). In these situations, the actions of each agent are shaped not only by their surrounding environment but also by the actions of other agents, both team members and opponents. A security team's patrol strategy, for example, may be adapted to respond to anticipated attacks, while an emergency crew's response is influenced by the proximity of the incident and the availability of other crews. These environments are inherently dynamic and challenging to predict, with numerous interacting variables and rapidly changing circumstances, making optimal collective decision-making and performance analysis a significant challenge.

The multi-agent systems (MAS) paradigm provides a framework for modelling, evaluating, and optimising the behaviour of agent teams in complex, real-world environments (Dorri et al., 2018). In MAS, intelligent agents interact with each other and their environment to achieve individual and collective goals through actions such as sharing information, forming teams, and coordinating strategies against adversaries. MAS research has addressed many challenges by representing individuals as decision-makers with shared or competing objectives, with notable applications in air traffic control (Brittain and Wei, 2019) and disaster response (Ramchurn et al., 2016). Many MAS applications are inherently spatiotemporal, requiring agents to coordinate over both space and time to optimise team performance (Capezzuto et al., 2020). For example, security teams may adjust the assignment of officers to cover vulnerable areas as adversaries change attack locations (Shieh et al., 2012), while ambulance crews strategically position vehicles in high-demand areas to minimise response times (Schmid and Doerner, 2010). Effective planning also involves managing agent workload and availability, such as scheduling shifts to prevent fatigue and maintain long-term team effectiveness.

However, while many real-world teams must consider factors such as teamwork, agent fatigue, workload balance and human factors (e.g., social norms and nuanced human interactions) for effective team formation (TF) and long-term performance, many existing MAS applications, often evaluated on virtual or robot teams, do not fully address these considerations. Here, teamwork refers to how agents coordinate and communicate, such as sharing information and spatially organising their positions, to achieve collective objectives. Although some models address agent coordination and responding to environmental changes, such as forming teams based on positive past agent interactions (Beal et al., 2020b), allocating agents to tasks to minimise travel times (Capezzuto et al., 2020), or matching agent skills to tasks (Gaston and DesJardins, 2005), they often overlook spatial coordination between agents and the long-term effects of fatigue on agent availability. Addressing these open challenges, especially in optimising team planning under uncertainty and adversarial conditions, would significantly advance the field, greatly enhancing both the analysis and practical performance of agent teams across real-world domains.

We focus on Association Football[1] as a key application domain for MAS modelling, where team selection and tactical changes in games are shaped by proactive managerial decision-making. Football offers a suitable environment for developing and validating new artificial intelligence (AI) and MAS techniques due to the complexity and dynamic nature of the game, and the presence of well-defined team objectives over extended periods in competitive, adversarial settings. Additionally, there is a rich availability of real-world spatiotemporal datasets across these time periods, enabling rigorous testing and analysis of these models. This combination of factors makes football a strong domain for studying real-world team behaviour and performance. The alignment between AI and MAS modelling and football is well supported in existing research (Tuyls et al., 2021; Beal, 2022; Radke and Orchard, 2023).

Analysis in football involves many factors, such as opposition tactics, match schedules, team formations, and player form and fatigue. At the player level, individuals work collectively to create scoring opportunities and reduce the risk of conceding. For example, players position themselves strategically to receive passes or to limit the opposition's attacking options. Match outcomes often depend on which team can coordinate best to maximise scoring opportunities and minimise threats. These forms of player coordination are captured in rich spatiotemporal datasets that record ball and player locations over time, making them well-suited for MAS modelling. Over the course of a season, team success also depends on the consistent availability of key players, which requires managers to consider player welfare in team selection and recruitment. Many club processes, including tactical analysis, training, recruitment, and opposition analysis, aim to improve team performance through better decision-making, and MAS and AI models

[1] Referred to as football for the remainder of this thesis.

offer powerful insights to inform these processes. Figure 1.1 presents a decision-making framework for football, highlighting how data analysis informs each stage.

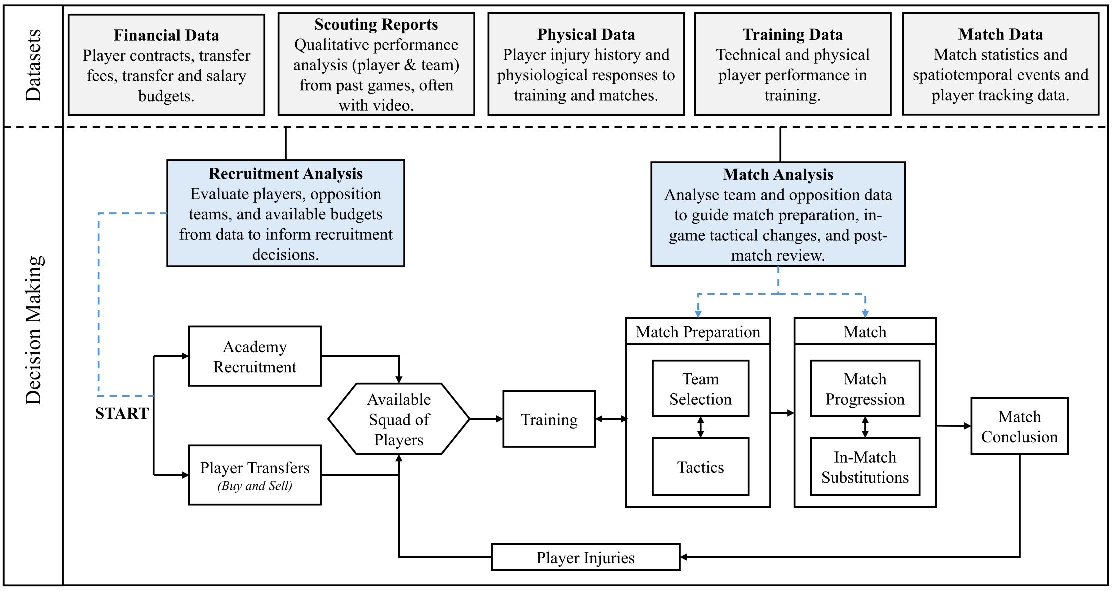

FIGURE 1.1: The decision-making process in football. The process begins with team recruitment and proceeds iteratively through training and matches, with data (grey) and analysis (blue) informing each stage. Datasets are continuously updated as the decision flow proceeds.

Football clubs have significant incentives to integrate new approaches to optimise team strategy due to the considerable social and economic benefits associated with team success, including substantial competition prizes and broadcast rights payments, with a total revenue of 38 billion Euros in 2023/24 across European leagues.[2] This has driven a fast growth in the application of various AI and machine learning (ML) techniques to gain advantages in club processes such as tactical planning (Beal et al., 2020a), player evaluation and recruitment (Decroos et al., 2019; Merhej et al., 2021), opposition analysis (Le et al., 2017b) and team training (Wang et al., 2024). A key enabler of the growth in AI modelling in football has been the emergence of detailed spatiotemporal datasets. While early research relied on basic match statistics, the introduction of event data provided information on the time and location of actions, such as shots or passes. More recently, advances in tracking technology have enabled the collection of high-frequency player and ball position data, providing more granular context, particularly for off-ball player positions, which are crucial given that players spend the majority of a match without possession of the ball. This has led to new objective metrics like Expected Goals (Spearman, 2018) and Expected Possession Value (Fernández et al., 2019), and tracking data has been used to model goal likelihood (Link et al., 2016), pass success (Spearman et al., 2017), and player physical statistics for injury prediction (Beal et al., 2019).

[2]Statistic found at: `https://www.deloitte.com/uk/en/services/consulting/research/annual-review-of-football-finance-europe-premier-league.html` (Last accessed: July 2025)

Despite recent advances in AI modelling using these datasets, football remains a challenging domain for AI and MAS research, with many open questions. Existing work has focused mostly on aspects such as evaluating on-ball actions (Decroos et al., 2019; Fernández et al., 2019) and optimising team selection based on player skill (Beal et al., 2020b). However, there is still limited research on evaluating and optimising off-ball positioning, player coordination, and long-term planning for challenges like player fatigue and injuries. Furthermore, the high cost and limited accessibility of tracking data, due to the requirement for advanced tracking systems, means that many researchers and analysts rely on less detailed, but cheaper, on-ball event data. This restricts the ability of these researchers to analyse and optimise team performance fully. As access to richer data improves and human decision-makers increasingly trust AI tools, AI-driven football modelling will likely become more sophisticated and widespread, enabling better tactical analysis, team selection, and evaluation of player performance and teamwork both on and off the ball.

In this thesis, we propose a set of research questions for applications of AI and MAS across key areas in football. In Chapter 3, we focus on the "Datasets" section of Figure 1.1, introducing a deep learning-based imputation method that reconstructs tracking data from event data, enabling off-ball analysis even when only event data is available. Chapters 4 and 5 contribute to player evaluation, addressing the open area of modelling and optimising off-ball positioning and coordination. These contributions can enhance the "Recruitment Analysis" and "Match Analysis" sections of Figure 1.1. These chapters demonstrate how these methods can improve player recruitment, match analysis, and team training. Finally, Chapter 6 focuses on the "Match Preparation" and "Match" sections of Figure 1.1, proposing a forward-looking model for team selection and substitutions that accounts for long-term player welfare and match outcome probabilities. Collectively, these contributions aim to support improved decision-making for football clubs across a wide range of strategic areas.

In the remainder of this introduction, the research questions explored in this thesis are outlined (Section 1.1), followed by the contributions of our research (Section 1.2). Finally, the structure of the remainder of the thesis is presented (Section 1.3).

1.1 Research Questions

The research presented in this thesis answers the following open research questions:

1. **Can we effectively predict spatiotemporal behaviour in team environments characterised by limited observability and irregular timesteps?**

   Background: Predicting agent trajectories in MAS is a well-studied problem (Sriram et al., 2020; Alahi et al., 2016), but challenges remain when observations are sparse, irregular, or incomplete. Existing models often assume continuous and uniform data, which is not always available in real-world settings such as drone footage with limited visibility or autonomous driving with obstructed views of pedestrians. In such situations, accurate estimations of agent locations are crucial to enhancing the effectiveness of response systems.

Football Application: Player tracking systems in football are expensive and not widely accessible, leaving many clubs limited to on-ball event data, where only one player location is observed at a time ($\sim$95% missing locations) and timesteps are non-uniform. Developing models to impute all player positions from on-ball event data would significantly expand access to off-ball analysis, which is vital for advancing tactical insights and performance evaluation.

2. **How can teamwork be modelled and optimised in adversarial MAS using spatiotemporal agent behaviour?**

   Background: Teamwork in MAS often involves complex spatial coordination between agents. While prior work has explored agent teamwork through pairwise interactions and task allocation (Beal et al., 2020b; Baker et al., 2016), an open question remains about how agents may coordinate to dominate adversaries spatially. Addressing this question will deepen our understanding of spatial teamwork dynamics and enable optimisation of agent coordination in dynamic, adversarial environments such as security settings and team sports.

   Football Application: Football is inherently a team sport where spatial coordination and off-ball movement are essential. Modelling spatial teamwork between players can reveal valuable insights into off-ball defensive organisation, attacking runs, and overall tactical effectiveness. Understanding and optimising these dynamics can significantly contribute to team analysis and coaching.

3. **What metrics can be used to quantify indirect agent contributions in team-based defensive scenarios?**

   Background: Evaluating individual agent performance within teams is crucial for optimising strategies in domains such as security and team sports. Traditional metrics typically assess performance through measurable actions, such as stopping a security breach or completing a high-value pass (Shieh et al., 2012; Decroos et al., 2019). However, indirect contributions, like an agent's positioning influencing an attacker's target choice in security or a defender's off-ball positioning affecting pass reception in team sports, are often overlooked. Developing metrics for evaluating these contributions will provide deeper insights into agent performance and enhance team strategies.

   Football Application: In football defence, players' off-ball positioning can significantly influence the opposition's attacking options, yet individual off-ball contributions remain underexplored in current research. Developing metrics to quantify these indirect contributions can improve player evaluation and tactical analysis, especially since players spend around 95% of their time off the ball.

4. **How can team performance and agent welfare be optimised in scenarios with uncertain agent availabilities?**

   Background: While prior TF models have focused on agent skills, spontaneous task emergency or robustness to agent failure (Gunn and Anderson, 2013; Okimoto et al., 2015; Demirović et al., 2018), limited research explicitly models how agent workload affects long-term availability and team performance. Effective TF in environments where agents are susceptible to fatigue or unavailability, such as emergency responders experiencing exhaustion or athletes at risk of injury, requires balancing immediate performance with managing fatigue and fault risk. Developing strategies that achieve this balance to enhance long-term team performance remains a key open area of research.

   Football Application: In football, teams face congested match schedules and frequent player injuries throughout a season. However, current team selection and substitution models mainly focus on player skill and chemistry (Beal et al., 2020b,a), often overlooking the long-term effects of player workload and injury risk. Developing strategies that balance immediate performance with effective fatigue and injury management can optimise squad rotation, reduce injuries, and help maintain consistent competitive success. These strategies have the potential to support more informed team selection decisions in football clubs.

1.2 Research Contributions

The key contributions of this thesis to the state-of-the-art are as follows:

- **Spatiotemporal Agent Imputation in Limited Observability:** In Chapter 3, we present a novel *Agent Imputer* model that addresses limited observations and non-uniform timesteps for agent movement prediction by capturing the temporal relationship between agent locations and observations, as well as the spatial relationship between agents. We test the model on real-world football events and tracking data and find that our model predicts player location to within $\sim$6.9 metres, outperforming multiple baseline models. We also demonstrate various applications of the model to improve the accessibility of football analysis for clubs and researchers.

> This work has been published and can be found in: "Inferring player location in sports matches: Multi-agent spatial imputation from limited observations". Gregory Everett, Ryan J Beal, Tim Matthews, Joseph Early, Timothy J Norman, and Sarvapali D Ramchurn. (2023). In: *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 1643–1651 (Everett et al., 2023b).

- **Learning and Optimising Spatial Teamwork in Multi-Agent Teams:** In Chapter 4, we model agent teamwork through spatial coordination. This approach models the impact of spatial proximity between agents on TF and long-term spatial dominance against adversaries using a Markov decision process (MDP), extended to handle multiple agents. Optimal agent teamwork is derived using Monte Carlo tree search (MCTS) and linear programming, where a set of agent subteams and movement-based actions is suggested. We apply this model to team defence in football using real-world data, finding that our approach reduces opponent threat by 24%, outperforming optimised individual behaviour by 7%. Additionally, our model enhances the predictive accuracy of future adversary attack locations and offers deeper insights into past teamwork models that do not explicitly account for the spatial impact of teamwork.

> This contribution was published in: "Optimising Spatial Teamwork Under Uncertainty". Gregory Everett, Ryan J. Beal, Tim Matthews, Timothy J. Norman, and Sarvapali D. Ramchurn (2025). In: *Proceedings of the AAAI Conference on Artificial Intelligence, vol. 39, no. 22, pp. 23168-23176. 2025* (Everett et al., 2025b).

> An initial, football-centric approach to this research contribution extends an existing football model that values goal probabilities using event data by focusing on how defensive and attacking coordination influences these probabilities. This contribution is beyond the scope of this thesis, but was published in: "Contextual Expected Threat using Spatial Event Data". Gregory Everett, Ryan J Beal, Tim Matthews, Timothy J Norman, and Sarvapali D Ramchurn. (2022). In: *StatsBomb Innovation in Football Conference* (Everett et al., 2023a).

- **Evaluating Indirect Agent Contributions in Multi-Agent Defence:** In Chapter 5, we present *GAPP*, a novel Graph Attention Network (GAT) model designed to predict football pass reception probabilities and offer interpretable insights into off-ball defending. Leveraging attention mechanisms, *GAPP* effectively captures interactions between off-ball players and introduces two new metrics, Defender

Influence (DI) and Defender Performance (DP), to quantify off-ball defender contributions. We evaluated *GAPP* on an English Premier League (EPL) dataset, demonstrating a 6.4% $\pm$ 1.5% reduction in binary cross-entropy (BCE) loss compared to multiple baseline models for pass reception prediction. Furthermore, we show that these defensive metrics correlate with future defensive actions and provide visualisations illustrating how they deliver actionable, data-driven insights for defender evaluation. The *GAPP* model and defensive metrics are designed to be generalisable across other team-based domains for measuring indirect agent contributions.

> This contribution was published in: "Evaluating Defensive Influence in Multi-Agent Systems Using Graph Attention Networks". Gregory Everett, Ryan J. Beal, Tim Matthews, Timothy J. Norman, and Sarvapali D. Ramchurn (2025). In: *2025 IEEE 12th International Conference on Data Science and Advanced Analytics* (Everett et al., 2025a).

- **Optimising Short- and Long-Term Team Selection:** In Chapter 6, we present a novel framework for dynamic TF and agent replacement. Specifically, the model optimises both pre-game team selection and in-game substitutions in football by explicitly accounting for player fatigue and injury risk. Our approach models these problems as stochastic MDPs and employs MCTS guided by expert policies to approximate optimal team selection and substitution decisions. The MDPs integrate a player-based match outcome prediction model alongside an injury risk model to simulate match transitions and evaluate player injury risk. Validated on real-world data from two EPL seasons, our framework demonstrates an average season points improvement of approximately 1% over baseline models, including a predictive model of human manager behaviour, while reducing squad injuries by around 3%, first-team injuries by 15%, and wages inefficiently spent on injured players by 3%. This model supports informed decision-making for team managers and is designed for extension to other dynamic, adversarial multi-agent domains.

> A preliminary version, focusing solely on pre-game team selection, was published as: "The Strain of Success: A Predictive Model for Injury Risk Mitigation and Team Success in Soccer". Gregory Everett, Ryan J. Beal, Tim Matthews, Timothy J. Norman, and Sarvapali D. Ramchurn (2024). In: *MIT Sloan Sports Analytics Conference 2024* (Everett et al., 2024).

1.3 Thesis Structure

The rest of this thesis is organised as follows:

- **Chapter 2:** This chapter provides a literature survey of current AI applications in football, with a particular focus on those leveraging spatiotemporal data. We also review relevant AI research in broader team-based domains. This chapter also identifies open research areas that will be used to expand on our research objectives.

- **Chapter 3:** In the first research chapter, we define a novel imputation model that estimates spatiotemporal agent behaviour in team environments of limited observability and non-uniform timesteps. We apply this model to football to estimate player tracking data and perform downstream off-ball analysis.

- **Chapter 4:** Next, we introduce a novel approach to teamwork, named spatial teamwork, that focuses on the spatial relationships that underpin inter-agent interactions. The model is designed to optimise coordination among agents in team-based defence scenarios and is applied specifically to football defence.

- **Chapter 5:** This chapter proposes new defensive metrics for quantifying indirect agent contributions in team-based defence by leveraging a GAT. The model is applied to football, where it predicts pass receptions and evaluates the contributions of off-ball defenders.

- **Chapter 6:** The final research chapter presents a novel TF and agent replacement framework that balances team performance and long-term agent availability. This framework is applied to football team selection and substitutions.

- **Chapter 7:** We reflect on the key findings of the thesis and their significance for the field. We also discuss the emerging challenges and explore potential future applications of AI in the football industry.

- **Chapter 8:** This chapter summaries the conclusions from our research.

- **Appendices A, B, C and D:** These appendices contain additional information organised by each research chapter in the thesis.

Chapter 2

Literature Review

In this chapter, we provide a background on football and critically evaluate key areas of research relevant to understanding and optimising behaviour and decision-making within football and other team-based domains. We begin by providing an overview of football (Section 2.1), followed by a review of the most recent data collection techniques and datasets available in football analytics (Section 2.2).

After providing a background to football, we examine state-of-the-art research in football and other team-based domains. This includes a detailed review of spatiotemporal prediction models for agent behaviour and trajectory prediction (Section 2.3), followed by an analysis of spatiotemporal performance models such as those used to evaluate player performance (Section 2.4). Finally, we explore current artificial intelligence (AI) models and algorithms designed to support team formation (TF) and strategic planning (Section 2.5).

Throughout the review, we identify key gaps and open challenges in the literature. These are summarised in a discussion (Section 2.6), where we highlight how addressing these gaps could advance decision-making and analysis in team-based settings. We also connect these open areas directly to the research questions in this thesis.

2.1 Overview of Football

In this section, we first outline the core rules and structure of football (Section 2.1.1), then summarise the key tactical processes used by teams in the sport (Section 2.1.2).

2.1.1 Fundamentals of Football

Football is a team sport played between two sides of 11 players, typically divided into four key roles: goalkeeper, defender, midfielder, and attacker. The match is played on a grass pitch and consists of two halves of 45 minutes each, plus additional stoppage time. Pitches can vary in size; however, the standard dimensions are 105 meters by 68 meters. The primary objective is to score more goals than the opposing team by getting the ball into the opponent's goal, which is being defended. At the end of a match, the possible outcomes are a win, a draw, or a loss.

Players perform a wide range of actions during a game, including shots, passes, tackles, interceptions, aerial duels, and saves. Set pieces, such as corners, penalties, and free kicks, are a subset of player actions that occur when the ball leaves the pitch or a foul is committed, and represent structured opportunities to influence the game. Every sequence of actions performed by a team ultimately results in either scoring a goal or losing possession.

Compared to many other sports, football is a low-scoring sport, with an average of just 2.74 goals per game over nine English Premier League (EPL) seasons from 2009/10 to 2018/19 (Zhao and Zhang, 2019). This low scoring frequency makes predicting match outcomes particularly challenging, as results are often determined by rare events that are difficult to predict and may not accurately reflect overall team dominance. Football's complexity is also due to its dynamic nature, with constantly shifting spatial and temporal interactions between players. These factors make analysis and prediction especially difficult.

The fine margins that separate wins from losses, along with the significant social and financial rewards of team success, drive clubs to seek any possible advantages. As a result, teams employ a variety of tactical approaches and decision-making strategies to maximise their chances of winning. An overview of these tactical processes is provided in the following subsection.

2.1.2 Tactics in Football

Teams employ several key tactical planning processes when aiming to maximise their probability of winning games and subsequently optimise league performance. Firstly, player scouting and recruitment are managed within the club to add players to the squad who possess the highest individual footballing ability and also fit the team's style. Team tactics are a multifaceted concept that aims to achieve optimal attacking performance from the players while also effectively neutralising the opposing team. A central element of tactics is the team's formation, which is the combination of defenders, midfielders, and attackers used in play. These categories are typically specified into role types, such as: goalkeeper (GK), wide defenders (LB/RB), central defenders (CB), defensive/central midfielders (CDM/CM), wide midfielders (LM/RM), wide attackers (LW/RW/LF/RF), and central attackers (CF). Figure 2.1 outlines two example team formations.

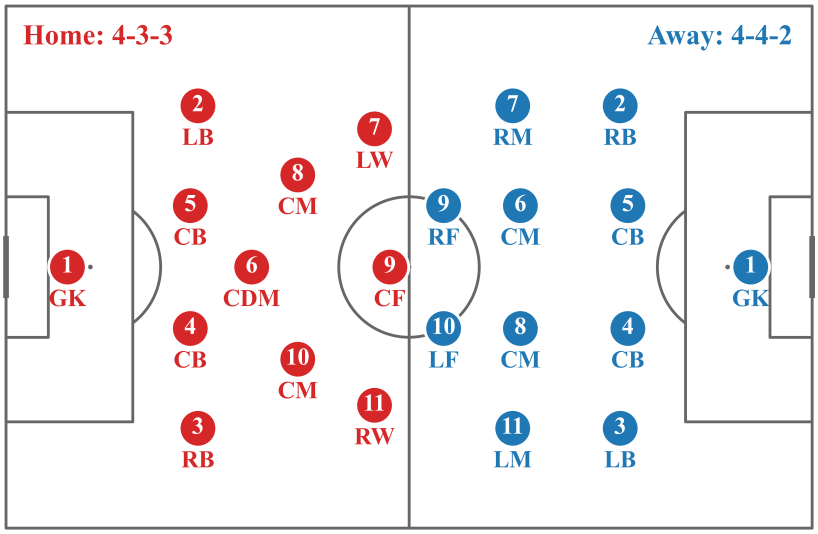

FIGURE 2.1: Example team formations.

Many formation variants exist, which may cater to a team's attacking or defensive style; for example, defensive-minded teams may play with five defenders. Additionally, teams use certain playing styles that suit their team best, whilst also considering their effectiveness against upcoming opposition. Examples of playing styles include high-intensity play, which involves fast one-touch passing and team pressing, or a more possession-based game, centred around many passes, gradually stretching the opposition's defence, and finding space. Set-piece optimisation can also give a team a considerable edge in a game since approximately one-third of all goals are scored from set-pieces (Yiannakos and Armatas, 2006).

Another key process within football analytics and tactical decision-making is opposition analysis, where the threats and weaknesses of an upcoming opposing team are identified using data from their previous games. Using this information, teams will design methods to exploit the weaknesses of the opposing team while neutralising their main threats. A fundamental part of this involves identifying space on the pitch, as many goal-scoring opportunities in a game arise from an attacking player having space to take an action under low pressure. The analysis of space can be applied from both an attacking and a defensive ideology. Attacking players may be presented with data regarding the opposing players who are most susceptible to leaving open space, whereas defending players will be made aware of the opposition's most dangerous threats and how these threats can be neutralised by limiting their space.

Some tactical decisions also occur in-game; for example, a manager will make substitutions in a game. A substitution occurs when one of the players on the pitch is replaced by a new player. Occasionally, substitutions are forced due to player injury. However, many substitutions are made with the intention of adding a tactical advantage to the team. These tactical advantages may include improving the team's energy, switching to a more defensive or attacking-minded approach, or reacting to an in-game observation that suggests the substituted player will have a positive impact. The timing of substitutions is also important, as strategically timed substitutions can proactively disrupt the opponent's momentum or reinforce an advantageous position.

Finally, managers will prepare their team to successfully execute the desired team style and perform optimally as a collective by improving player teamwork and communication. This preparation occurs through training sessions that utilise team-based exercises and challenges, thereby fostering team coordination. Furthermore, teams may often eat together to encourage socialising between teammates. Data analysis is increasingly being used in football to enhance various decision-making processes within clubs, facilitated by the collection of football datasets. We discuss the current datasets in football in the next section.

2.2 Data in Football

Current datasets that describe match play in football can be split into three distinct types. We list these as follows:

Match sheet data The earliest form of football data, consisting of high-level football game information such as the score, team line-up, goal scorers, number of cards and possession. This information is widely accessible, as it is collected for all professional football matches and is typically available to view on betting websites, football team and league websites, and sports websites such as FBRef[1] and WhoScored[2]. These data contain no spatiotemporal properties, as no location data is included in match statistics.

Event data Logs all in-game actions, such as passes, shots, dribbles and tackles, throughout a football game. For each event, the action type and location are stored together with the time of the event in the game and the player involved. These data are mostly collected using human annotators who watch a whole football game and log data whenever they observe a new event. Although less widespread than match data, event data are prominent in professional football, with many different providers covering over 80 leagues globally. These data are provided by data companies such as Statsbomb[3], Stats Perform[4] and WyScout[5]. The collection of event data has led to

[1] https://fbref.com/en/.
[2] https://www.whoscored.com/.
[3] https://statsbomb.com/.
[4] https://www.statsperform.com/.
[5] https://wyscout.com/.

the development of numerous AI models in football that utilise spatiotemporal data as input. A typical Statsbomb event dataset encompasses over 3,400 events in each game, resulting in spatial updates of on-the-ball play with an average time interval of less than 1.6 seconds. Figure 2.2 illustrates an example of an on-ball action recorded in event data format.

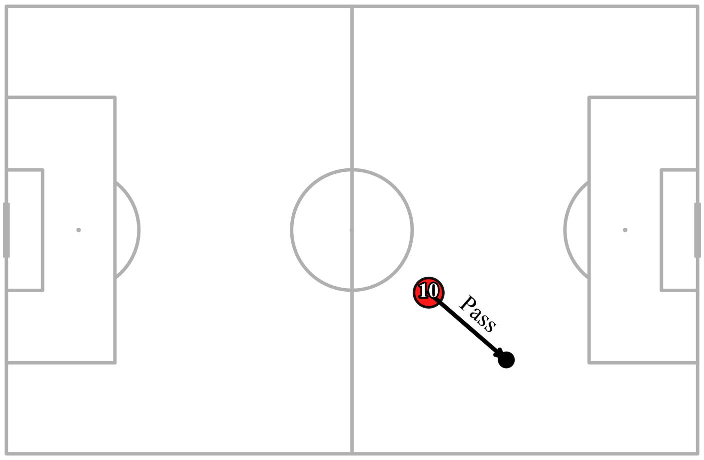

Time	Action	Start	End	Result	Player	Team
0m 33s	Pass	(64,25)	(76,14)	Success	M. Gibbs-White	N. Forest

FIGURE 2.2: Example visualisation of a football on-ball action recorded in event data format.

Tracking data Logs the location of all players and the ball throughout a game with a typical sample interval of less than 50 milliseconds. This generates a comprehensive image of the gameplay dynamics in a game. Data collection techniques for this data type vary, but most typically involve optical tracking systems (Liu et al., 2009; Manafifard et al., 2017; Najeeb and Ghani, 2021). Xu et al. (2004) detail the use of eight static cameras to form a multi-view tracker, whilst Beetz et al. (2007) describes the use of TV broadcasts to extract player tracking. Dearden et al. (2006) outline a tracking approach using a single moving camera and particle filtering. Rahimian and Toka (2022) give an extensive analysis of different optical tracking methods, and compares current methods for generating tracking data. The cost of suitable cameras and technology for these methods is very high, making the data expensive to collect. Consequently, these datasets are not as widespread as other data types and are mostly restricted to collection in the top 5 European leagues. These data are collected and provided by companies such as Stats Perform, Metrica[6], Ubitrack[7], Skillcorner[8], Inmotio[9], Bepro[10] and Gradient Sports[11],

[6] https://metrica-sports.com/.
[7] https://ubitrack.eu/.
[8] https://www.skillcorner.com/
[9] https://inmotio.eu/.
[10] https://www.bepro11.com/.
[11] https://www.gradientsports.com/

where provision of data often requires a specific enquiry on the teams or league. Figure 2.3 illustrates an example of a football tracking data frame.

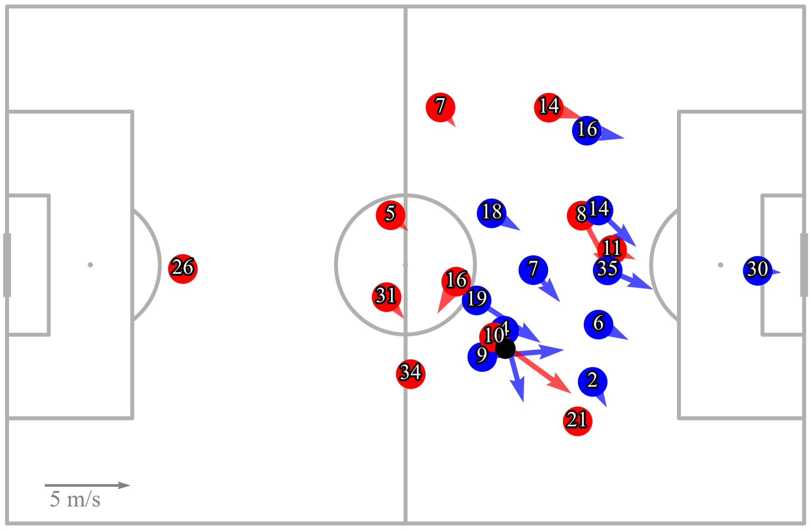

FIGURE 2.3: Example visualisation of a football tracking data frame.

Body pose data is an additional data type in football analytics that can be used to enhance tracking data. These data will include 3D motion capture of players, logging their body shape and position throughout a game (Nakano et al., 2020). However, it is in very early development stages, so its use is minimal. Another recent development in football data is a hybrid dataset named Statsbomb 360 data[12] created by Statsbomb. This data includes typical event data and additional coordinate information on the location of all players within the broadcast camera's view at the time of the event. Therefore, this data provides further context to event data, but does not contain as frequent samples of player locations as tracking data. Furthermore, this data only contains information on the team of the players in camera view, in contrast to tracking data, which includes unique identifiers for every individual player on the pitch. Many professional teams also use wearable athlete monitoring systems, such as Catapult[13], to collect physical and physiological load metrics during training and matches. These metrics include total distance, sprint count, accelerations and heart rate. While these data are highly relevant for physical performance, fatigue management and injury risk, they are largely separate from spatiotemporal analysis in the field. With this distinction in mind, Table 2.1 presents a summary of the key differences between the current core datasets in football, including the value they offer in performing spatiotemporal football research.

The spatiotemporal descriptiveness of tracking data is undoubtedly the most extensive of all current core data types. The frequent snapshots of all player locations during the game enable the development of models that consider off-ball information, such as defensive pressure and attacking runs. The production of tracking data has therefore led

[12]https://statsbomb.com/what-we-do/soccer-data/360-2/
[13]https://www.catapult.com/

TABLE 2.1: Comparison of football datasets.

Data Type	Key Features	Availability	Spatiotemporal Descriptiveness
Match Data	Basic game information such as the score, lineups, and cards.	Freely available on many platforms for all professional games.	N/A
Event Data	On-the-ball events in a game (e.g., shots). Contains the event location, player, and timestamp.	80+ professional leagues globally. Available for purchase from data companies.	Location and time of on-the-ball events.
Tracking Data	Player and ball locations with a sampling interval of 50ms or below.	Mostly restricted to major European leagues. Available from data companies upon enquiry.	Locations of all players and the ball throughout every moment in a football game.

to a surge in spatiotemporal research works in football (Beal et al., 2019; Kovalchik, 2023). However, the extent and speed of this research are still limited, as many researchers are unable to access tracking data due to budget constraints. The practical use of tracking data for player recruitment and scouting is also limited, as current data primarily resides in the major European leagues. Ideally, over time, tracking technologies will become cheaper and more accessible, and tracking datasets will have a similar price and availability to event data. Due to the key differences in spatiotemporal descriptiveness, price and availability of event and tracking data, we note the utilisation of a model to estimate tracking data using event data as a key research area. In the next section, we discuss current state-of-the-art spatiotemporal prediction models in football and across all domains, highlighting open research areas and their potential applications in football.

2.3 Spatiotemporal Prediction Models

Predicting agent locations can provide many use cases in multiple domains, such as path prediction for autonomous vehicles (Sriram et al., 2020), using drone footage to predict civilian locations in disaster response (Ramchurn et al., 2016; Wang et al., 2021) and predicting player locations in team sports (Lee and Kitani, 2016). Many research projects have investigated the use of AI to predict agent locations, with a significant portion focusing on future trajectories of agents (Kim et al., 2013). Some also focus on imputing agent trajectories, which fills in missing locations using past and future observations of an agent (Omidshafiei et al., 2022). Figure 2.4 visually demonstrates the difference between agent trajectory prediction and imputation, focusing on a single agent.

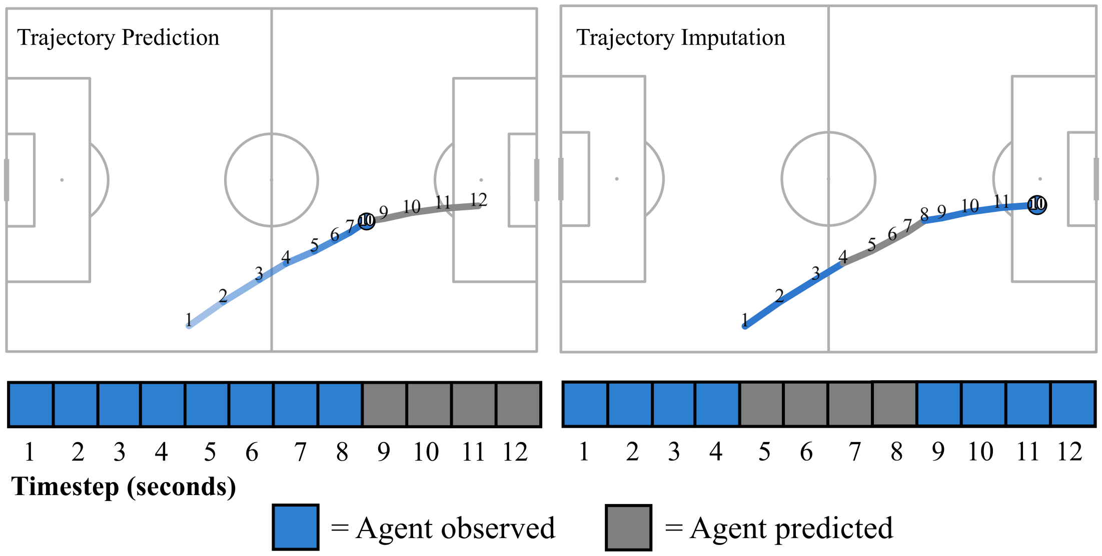

FIGURE 2.4: Example visualisation of agent trajectory prediction (left) and imputation (right) in football for a single player. Blue timesteps are where the agent location is observed, and grey timesteps are estimated by a prediction model.

In this section, we discuss the current spatiotemporal prediction and imputation models in both football and across all multi-agent domains. Areas for additional work, as well as applications and use cases in football, are also discussed.

2.3.1 Spatiotemporal Prediction Models in Football

The most common use case of spatiotemporal prediction models in football and other team sports is to predict future trajectories, where the future movement and positioning of a player are estimated given past observations of the player. Current approaches include the use of long short-term memory (LSTM) based architectures, which predict multiple possible trajectories for each player and assign probabilities to each outcome, as demonstrated in basketball (Hauri et al., 2021). In football, imitation learning (IL) has been applied to trajectory prediction by extending single-agent IL to the multi-agent domain using deep learning. This allows individual player policies to be learned while also capturing coordination and agent roles within the team, with models tested in both synthetic and real-world football scenarios (Le et al., 2017b). In American football, trajectory prediction of wide receivers has been approached using Markov decision process (MDP) frameworks that incorporate prior knowledge of the game and opposition positioning, such as in the work of Lee and Kitani (2016).

To address the challenges of modelling coordinated, multi-agent behaviour, recent research has increasingly turned to graph-based learning methods. The dynamic and relational nature of football makes it particularly well-suited to graph neural networks (GNNs), which can naturally represent the complex interactions between players on the pitch. For example, GNNs have been used to predict player movement and behaviour. Yeh et al. (2019) applied GNNs to forecast player trajectories in both football and basketball, while Monti et al. (2021) similarly used GNNs to model player objectives and interactions for trajectory prediction in basketball. Another notable approach is the use of variational recurrent neural networks (VRNNs) for basketball, where each player is associated with a VRNN and their interactions are captured through a fully connected graph network. This model leverages video frames to predict the current state of the game and subsequently predict future player trajectories (Sun et al., 2018).

In contrast to predicting future behaviour, imputation focuses on filling in missing information about agent behaviour using past and future observations of the agent. This is especially useful in domains where behaviour is observed intermittently. For example, an existing use case in football is the imputation of trajectories when players go out of camera view, which was achieved using GNNs and variational autoencoders (VAEs) (Omidshafiei et al., 2022). Player-specific information is encoded using LSTMs with shared parameters, and the hidden state of these LSTMs is represented as nodes in variational graph networks. Using this model, a distribution of each player's state is generated. For this problem, broadcast tracking data is used, where player locations are logged at regular timesteps when the player is in camera view, and the imputation model estimates player locations with the same timestep intervals when the player is out of camera view. Another example of imputation is presented by Qi et al. (2020), who build a model that predicts both future behaviours and imputes missing trajectories of players in basketball using IL.

While current work has studied the imputation of player behaviour and trajectory using tracking data with missing observations at some timesteps, to our knowledge, no work has studied the imputation of tracking data using only event data. This presents unique challenges due to the sparsity of player observations in event data and the non-uniformity of timesteps between on-ball events. The availability and price of event data compared to tracking data make this an important use case for sports, as clubs with limited resources can leverage the analytical strengths of tracking data, particularly for off-ball analysis, without the need to install expensive tracking systems. In the following subsection, we discuss spatiotemporal agent prediction models in MAS.

2.3.2 Spatiotemporal Agent Prediction Models

Spatiotemporal prediction models to predict agent locations have many use cases in a wide range of dynamic MAS. These domains include collision avoidance for autonomous vehicles (Xie et al., 2021; Sriram et al., 2020), pedestrian trajectory predictions in crowds (Marchetti et al., 2020; Alahi et al., 2016; Ivanovic and Pavone, 2019) and location habits of human populations (McInerney et al., 2012, 2013).

Similar to team sport, the focus of many of these studies is future trajectory prediction, with some utilising mathematical models to achieve this. For example, Kim et al. (2013) uses a physics-based approach to estimate a motion model for each human and subsequently predict human trajectories in crowded areas. Ensemble Kalman filtering is used to estimate the parameters of the human motion model that best match data from past observations. This motion model is then used to infer the human's future trajectory. Alahi et al. (2016) also predicts the future trajectory of humans in crowds given their past position data. In this case, each human is represented as an agent, and their past movement is used as sequential input to an LSTM. This model also handles interactions between humans when in close proximity to each other, using a social pooling layer where neighbouring agents share the hidden states of the respective LSTMs, representing their past movements. Sriram et al. (2020) also use deep learning to predict multiple human trajectories by utilising convolutional LSTMs and conditional VAEs. This model is tested on both simulated trajectory datasets and real traffic scene datasets. Some algorithms instead focus on estimating agent trajectories when they cannot be tracked temporarily due to obstacles. For example, Gong et al. (2011) use multiple hypothesis planning to learn human intentions from their trajectory data and predict multiple possible paths when the human is out of tracking view. Furthermore, Zhong et al. (2021) impute missing traffic data using graph convolutional networks (GCNs) and LSTMs.

Future locations of human populations are also predicted in the current literature. In these cases, data on human locations is collected using global positioning system (GPS) or WiFi data and the user's future locations are then predicted. McInerney et al. (2012) highlight that predicting new users in this system can be challenging. Therefore, they apply a hierarchical Bayesian model that models similarity between new users and current users, leveraging this information to predict the new user's future locations. This model is tested using a phone dataset for 38 real-world users. Liu et al. (2022b) utilise a GCN and gated recurrent unit (GRU) on data for the population density of crowds to predict future crowd density in major cities. These models can be used to learn more about human tendencies and daily habits, as well as to provide a useful estimator of population flow during an epidemic.

Trajectory prediction has also proven helpful in improving the efficiency and safety of shipping. For example, Siegert et al. (2016) use an Extended Kalman filter to track vessel trajectories and help reduce collision risk. Similarly, Perera et al. (2010) propose the use of an Extended Kalman filter to estimate the position and velocity of vessels using noisy positional data. Deep learning approaches have also been applied to the problem of ship tracking. For example, You et al. (2020) utilise an LSTM and GRU to work as an encoder-decoder model and consequently output future trajectories. This model uses a ten-minute sequence of a ship's past trajectories to predict its trajectory over the next five minutes. These predictions can be used to alert ships to potential collisions earlier.

A wide range of algorithms and models have been used to solve spatiotemporal prediction problems across many domains. These include mathematical tools such as the Kalman filter and Extended Kalman filter (Perera et al., 2010; Siegert et al., 2016), multiple hypothesis tracking (Blackman, 2004) and deep learning approaches. These deep learning architectures include convolutional models (Liu et al., 2022b), recurrent neural networks (RNNs) (You et al., 2020; Alahi et al., 2016) and graph-based architectures (Zhong et al., 2021; Yu et al., 2018). Whilst most of these studies consider future trajectory prediction, some studies focus on the retrospective imputation of missing data (Zhong et al., 2021). Whilst these studies cover a wide range of domains and problems, unique computational challenges arise from the problem previously discussed in Section 2.3.1 regarding the estimation of tracking data in team sport using only on-ball event data.

Event data leads to sparse player observations with non-uniform timesteps. This occurs as player behaviour is only gathered when the player performs an on-the-ball action (e.g., pass, shot, tackle), which can happen at any time. This often leads to a cluster of player events followed by an extended period with no information as the player is no longer on the ball. Furthermore, as only a single player can perform an on-the-ball event at a time, data is collected for just one player at each time step. To our knowledge, no previous work has studied the problem of imputing behaviour within MAS when information occurs at non-uniform timesteps and for only a single agent at a time. This problem can be visualised in Figure 2.5.

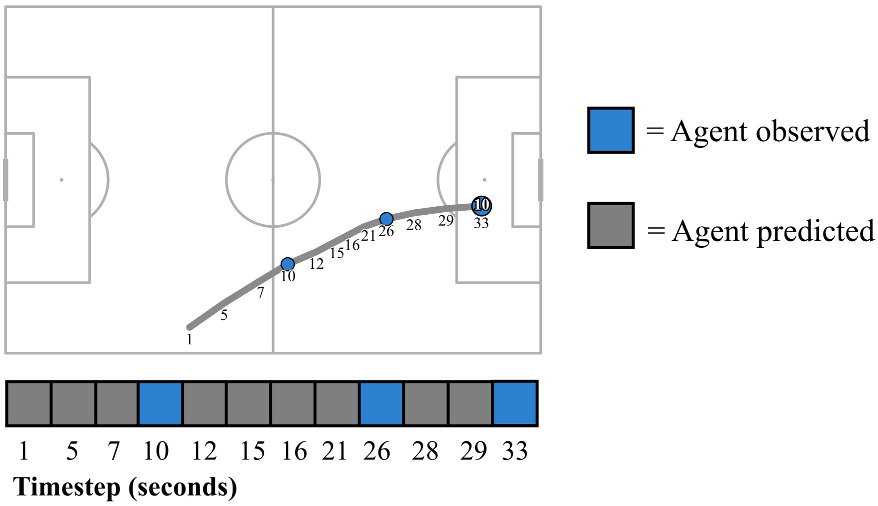

FIGURE 2.5: Example visualisation of agent position imputation in football for a single player using sparse event data with non-uniform timesteps.

In Section 2.6, we expand on this imputation problem as one of the research objectives in this thesis. In the next section, we discuss the application of AI models to football for analysing spatiotemporal player and team performance, and we present research and technical challenges in other domains that may also apply to football.

2.4 Spatiotemporal Performance Models in Football

AI modelling applications have been used to produce additional insight into many areas in football, such as player recruitment, scouting, player performance analysis and tactical analysis (Beal et al., 2019; Tuyls et al., 2021; Davis et al., 2024). Furthermore, the use of spatiotemporal features in football research has rapidly grown in recent years due to increased information on player and ball locations, providing better detail of gameplay. Many applications of these data have been evaluated, including the objective evaluation of player actions on and off the ball and the optimisation of player performance. We split spatiotemporal football models for players into two key areas: player performance and teamwork. These areas are identified as key for AI to optimise performance, training, and recruitment at the player level.

2.4.1 Player Performance

Evaluation of player performance and decision-making plays a factor in a club's strategy to pick the best line-up and tactical setup for their team and optimise their player recruitment. In football, player actions can be grouped and measured by their primary objective, with three broad categories:

Attacking actions Shots, passes and dribbles, evaluated by the extent to which they increase the team's probability of scoring.

Defensive actions Actions such as tackles, interceptions, and blocks, evaluated by the extent to which they reduce the opposition's probability of scoring.

Space creation Off-ball positioning and movement, evaluated by their effectiveness in disrupting defensive structures or reducing attacking space.

The following sections elaborate on each action category, including the existing literature for evaluating player performance in these areas. We also explore how player decision-making has been modelled and how this is used to understand player behaviour better and optimise decision-making for all types of actions and objectives.

2.4.1.1 Player Attacking Actions

Objectively measuring attacking actions in football allows for the evaluation of a player's efficiency and direct comparisons between players when deciding a team lineup. A widely adopted evaluation metric within the football community is expected goals (xG)[14], which objectively quantifies the probability of a shot resulting in a goal by using thousands of previous examples from data (Spearman, 2018). Various derivations of xG have been produced using different features, such as shot distance and shot angle. Building on these derivations, the rise of player tracking data fueled the development of xG models, which incorporated opponent positioning and pressure into the predicted expected value (Lucey et al., 2015). This metric can lead to many interesting interpretations for a coach to evaluate how their forward player is performing. For example, if a player has scored fewer goals throughout a season than their xG suggests they should have, this could imply that the forward is inefficient at converting goal-scoring chances. However, it may also indicate that they have been unlucky and that they are likely to score more goals in upcoming games.

Goal-scoring probability can be extended beyond just shot-taking, and attacking efficiency and potential can be evaluated further based on the number of relevant attackers and defenders, the space available to attackers, and the location of the ball. These factors justify the development of models that objectively measure the value added by other actions in football, such as passing. Decroos et al. (2019) use event data to create a 'Valuing Actions by Estimating Probabilities' (VAEP) value, designed to quantify the value of a player's actions in a game. This methodology considers all action types, such as passes, dribbles and tackles. The VAEP value is calculated by estimating the change in probability of a team both scoring and conceding a goal when the action occurred, such that a favourable action will have reduced the likelihood of the team conceding and increased the likelihood of the team scoring. Similar to the VAEP metric, an 'Expected Threat' (xT) metric was created by Karun Singh (Singh, 2019), which frames attacking sequences as an MDP where pitch zones represent states. This approach uses event data to value the goal-scoring potential of a possession situation by considering both the probability of shooting and scoring from the current zone, as well as the probability of alternative on-ball actions, such as passes, to further the scoring potential later in the attacking move. These probabilities are learned from historical data.

Despite both VAEP and xT sharing a common task of valuing actions in football using event data, both metrics vary in how they quantify value. A critical comparison of the two models found that xT valued creative play (e.g. passes and dribbles) higher with a stronger correlation to assists per 90 minutes. In contrast, VAEP had more bias towards shot-takers with a stronger correlation to shots per 90 minutes (Decroos et al., 2020). The comparison concludes that both approaches deviate from previous metrics and therefore add value to the current state-of-the-art in football analysis.

The value of event sequences can also be evaluated by considering the success of similar sequences in the past. For instance, Decroos et al. (2017) apply event data to rate phases of play, which they define as continuous intervals without a possession turnover. Using

[14] Introduced in: https://www.statsperform.com/resource/assessing-the-performance-of-premier-league-goalscorers/ (Last accessed: October 2025)

a K-nearest-neighbours approach, they identify comparable phases in historical data and assess how frequently these phases resulted in goals. In this way, actions are ranked based on both the overall effectiveness of the phase of play and the specific contribution of each action within that phase, thereby offering a measure of an individual player's influence on the outcome.

The emergence of tracking data has enabled the incorporation of off-ball information, advancing the state-of-the-art in action valuation. Before reviewing research that uses these data, we first introduce the concept of pitch control. Pitch control is a popular term in football analytics that quantifies the degree of control each team exerts over different areas of the pitch based on player locations. If a player is likely to gain possession should the ball arrive in a given area, the model designates that area as controlled by their team. An example is shown in Figure 2.6, illustrating the regions of the pitch controlled by each team, based on which players are most likely to reach the ball first.

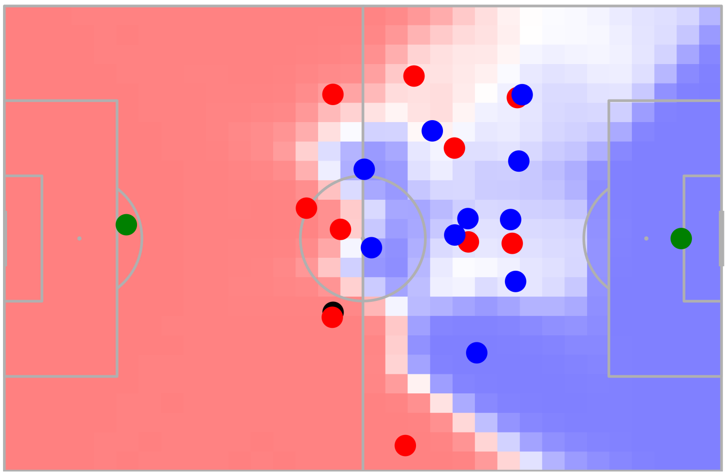

FIGURE 2.6: Example visualisation of team pitch control.

Fernández et al. (2019) use tracking data to calculate player pitch control and introduce a risk–reward ratio that assigns values to all possible actions available to the ball-carrier. Specifically, the model integrates the probability of pass completion, the expected value of the completed action, and the expected value of an opponent interception. These features are combined to calculate an expected possession value (EPV), estimated through machine learning (ML) techniques such as logistic regression and deep neural networks. EPV is a state-of-the-art metric that analysts can use to evaluate decision-making by assessing whether players in possession select the action with the highest EPV.

Building on this, Fernández et al. (2021) extend EPV by splitting it into sub-components. They develop neural network models to separately estimate the value of player dribbles, passes and shots and combine these predictions using an ensemble of these models. This framework also utilises player tracking data for spatiotemporal context, enabling more detailed analyses and visualisations of specific aspects of play. For example, it can highlight situations where a player chose a low-probability shot despite more promising passing options. Visualisations of these metrics also provide coaches with tools to communicate tactical feedback and support decision-making guidance for players.

In a similar vein, Link et al. (2016) introduce the *dangerousity* metric to quantify the goal-scoring probability at each moment of a match and thereby assess the potential of an attack using tracking data. This metric is derived using four key factors: ball location, the degree of control the player has over the ball, the probability of the defending team preventing a player's action, and the density of defenders in the attacking zone. They also demonstrate how this metric can be applied to value player actions and present a related *Dominance* metric, defined as the difference in cumulative dangerousity between the two teams. The effectiveness of Dominance in capturing team strength is validated by its stronger correlation with win probability from betting odds (r = .82) compared with classic indicators such as goal difference (r = .55) and shots (r = .58).

The central aspect of attacking build-up play is passing, and players must account for numerous factors when selecting passing options. Passing valuation models have been developed to evaluate passing options based on spatial features, such as opponent positioning, thereby enabling the evaluation of player decision-making. Such models can also inform opponent analysis and defensive coaching. Power et al. (2017) use tracking data to measure the risk and reward of potential passes, defining risk as the probability of pass completion and reward as the likelihood of a shot occurring within 10 seconds of a successful pass. These values are estimated using contextual features from tracking data, including the receiver's speed and the position of the nearest defender. Similarly, Goes et al. (2021a) model pass risk and reward with a LightGBM classifier to predict pass outcome probabilities at the moment of passing. Their approach incorporates multiple reward functions, such as the probability of scoring or the number of opponents bypassed by the pass.

Beyond risk and reward, Robberechts et al. (2023) propose an evaluation of pass creativity in football, defined by both originality (how much it differs from typical player passes) and their contribution towards goal-scoring. Using an XGBoost model with hand-crafted features derived from player location data, this framework provides a means of assessing players based on the creativity of their passing decisions.

Goes et al. (2019) note the scarcity of goals in football and instead value passes by the degree of defensive disruption they cause towards the opposing team. They present a quantitative model that uses tracking data to measure how far opposing players are displaced by a pass. This approach contrasts with previous approaches that focus on scoring probabilities, instead focusing on team interactions and spatial dynamics. Additionally, it avoids overvaluing forward passes and highlights the importance of passing sequences in creating space. Spearman et al. (2017) present a physics-based model that estimates the time required for a player to arrive at a pass trajectory and subsequently control the ball, which is then used to calculate a pass probability. From this, the model derives a pass value and allows evaluation of a player's capacity to intercept passes. The validity of this framework is demonstrated through a strong correlation (r = 0.83) between pass value and passes completed in the final third.

The works discussed above only consider a limited range of data types. Several important factors and action subtypes, such as headers, strong-footed shots, and weak-footed shots, remain largely unexplored. The value of each action subtype could be evaluated separately, enabling more nuanced assessments of attacking play. Similarly, richer features of passing and player characteristics, such as pass elevation and the height of the target player, could be considered when identifying optimal passing options. Including these factors would allow coaches and researchers to analyse decision-making at a finer level, moving beyond player positions and velocities. A key bottleneck in these developments is the current lack of descriptive data for such factors. However, these tasks will become more achievable as new datasets, such as body pose data and physical player features, become more common.

Another research direction for valuing attacking actions could apply inverse reinforcement learning (IRL) approaches from other domains to learn player intentions and build a custom reward function using expert demonstrations rather than goal-scoring probabilities. These approaches could be applied to football to rate player actions based on player behaviour without bias towards just goals. Examples of IRL approaches to model human motivations and build custom value functions exist in the popular game World of Warcraft (Wang et al., 2019a) and Ice Hockey (Luo, 2020). Further detail on these modelling approaches is given in Section 2.4.1.4.

2.4.1.2 Player Defensive Actions

Measurement of a defender's ability to prevent goal-scoring chances is a key factor in a coach's decision-making when deciding a team line-up and recruiting new defensive players. However, research into the objective evaluation of defensive actions is scarcer than attacking actions (Forcher et al., 2022). This leaves significant potential for novel research and tools for coaches to directly compare the performance of defenders using objective statistics which consider spatiotemporal game features.

Research in defensive evaluation has used different extents of spatiotemporal data. For example, Merhej et al. (2021) measure the value of defensive actions using event data and a neural network to model the expected threat of an attacking situation if the defensive action did not occur. This approach predicts the value of the defensive action based on what it prevented and is used to objectively rank defenders based on their cumulative defensive score. However, this is limited to on-ball actions, overlooking the substantial influence of off-ball defending, which is particularly significant given that players spend approximately 95% of match time off the ball.

Recent advances in spatiotemporal tracking data have enabled the development of more context-aware defensive metrics. For example, Stöckl et al. (2021) leverage tracking data and GCNs, which are well-suited for modelling interactions and relationships between players, to evaluate team defensive performance. Their approach estimates the likelihood of opposing players receiving a pass at any moment, the probability of a shot occurring within the next 10 seconds, and the success rate of passes to each attacker. It also allows identification of defensive strategies such as man-to-man and zonal marking, providing insights into both individual and collective defensive behaviour.

Similarly, Forcher et al. (2024) use features derived from tracking data as input to a logistic regression model to predict the probability of successful defensive actions (i.e., ball recoveries). Their approach can be applied to evaluate players and teams by, for example, identifying whether a player wins possession more frequently than expected or whether team pressing actions significantly increase recovery probabilities. However, while these models and other recent studies (Fernández et al., 2021; Robberechts et al., 2023) incorporate off-ball defensive positioning, to our knowledge, no existing approach directly evaluates the influence of individual off-ball defenders on attacking outcomes.

A key gap in current research is the use of off-ball spatial data to value the defensive actions of individuals. Most existing approaches concentrate on on-ball events, overlooking contextual factors such as the number and positioning of defenders behind the ball. For instance, a last-man tackle is typically more valuable than a challenge made when the defensive unit is already organised, yet many existing metrics fail to capture this distinction. Incorporating fine-grained context from tracking data, such as the scoring potential of the situation, would allow for more accurate and informative quantification of defensive contributions. Metrics like EPV or dangerousity (Link et al., 2016) could be adapted to estimate the danger of an attack prevented by a defender, thereby improving comparability between individuals.

Developing a quantifiable metric for individual off-ball defensive play using tracking data would also enable coaches to objectively assess how effectively defenders limit space - an essential aspect of defensive performance. For instance, such a metric could capture the value of successfully marking a dangerous run in the box or covering a space left open by a teammate. By evaluating the danger of an attacking situation had that space not been covered, new metrics could provide a more nuanced understanding of defensive contributions beyond on-ball actions. Inspiration for this approach can be drawn from basketball analytics, where Franks et al. (2015) use tracking data to analyse the spatial relationship between defenders and attackers, quantifying the impact of defensive positioning on an attacker's shot frequency and efficiency. To learn which defenders are guarding which attackers, their method employs an expectation–maximisation algorithm to infer latent defender–attacker assignments over time by iteratively inferring probabilistic assignments based on spatial positioning and updating a model of defensive behaviour until the assignments and parameters are consistent. This is combined with generalised least-squares to estimate the impact that defender positioning has on attacker shooting outcomes while accounting for correlation in the error structure resulting from repeated measurements of the same defenders and attackers. Adapting a similar methodology to football would allow analysts to directly compare defenders based on their spatial impact on attackers, rather than solely on traditional statistics like tackles and interceptions. These comparisons could enhance decision-making in scouting and team selection.

Building on the use of GCNs for modelling player interactions (Stöckl et al., 2021), recent advances in ML, particularly the development of attention mechanisms, also offer new opportunities for more nuanced defensive analysis. Attention mechanisms enable models to focus on the most relevant input data, thereby improving both predictive performance and interpretability (Niu et al., 2021). These techniques have become widely adopted in fields such as natural language processing (Galassi et al., 2020), computer vision (BenTaieb and Hamarneh, 2018), and, increasingly, in graph-based modelling (Simeunović et al., 2022; Xing et al., 2021). Their potential to enhance the evaluation of off-ball defensive contributions makes them especially relevant in the context of football.

In football analytics, GNNs have proven particularly effective for modelling player relationships, such as tactical setups during corners (Wang et al., 2024) or predicting player actions (Rahimian et al., 2023). Traditional GNNs like GCNs aggregate information equally from connected nodes through message passing (Hamilton et al., 2017), which can limit their ability to capture the varying importance of different players in a given context. Graph attention networks (GATs) address this limitation by incorporating attention mechanisms that assign context-dependent, learnable weights during aggregation (Veličković et al., 2018; Hu et al., 2021). This allows the model to identify and focus on the most influential nodes, such as key defenders in a particular phase of play, thereby enhancing both model performance and interpretability (Hu et al., 2021; Zhang et al., 2024). In the context of evaluating off-ball defensive contributions in football, GATs could significantly advance previous approaches by enabling the extraction of interpretable defensive influence metrics for individual players, even when they are not directly involved in on-ball actions. A visualisation for modelling football players as a graph, both in the form of a GCN and a GAT, is shown in Figure 2.7.

Despite these advantages, extracting meaningful interpretability from attention weights remains challenging. High attention weights do not always clearly indicate positive or negative impacts, and some studies suggest they may not reliably reflect true importance in certain settings (Bai et al., 2021). To address this, methods such as node masking have been proposed to better extract node influence in GNNs (Ying et al., 2019). Other recent

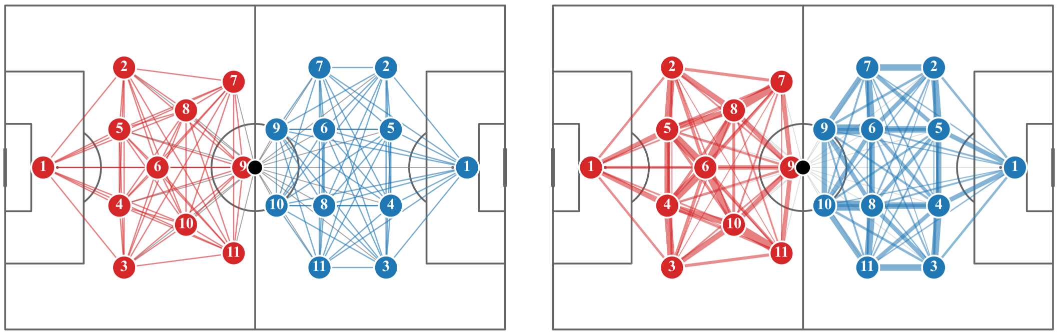

FIGURE 2.7: Example of modelling football players as a graph. Left: a graph convolutional network (GCN), where all player relationships are weighted equally. Right: a graph attention network (GAT), where relationships are assigned learnable, context-dependent weights based on their importance.

advances in deep graph learning have also made progress in knowledge graph reasoning, interpretability, and model robustness. Zhang et al. (2023) propose adaptive propagation paths for GNNs that identify semantically related entities, improving the interpretability and efficiency of reasoning in complex relational domains. Zhou et al. (2023) present the Robust Graph Information Bottleneck method to enhance the robustness of link prediction in GNNs when there is bilateral edge noise, a common issue occurring from noisy data in real-world multi-agent domains. Additionally, Zhou et al. (2024) propose scalable, query-dependent subgraph extraction to improve link prediction efficiency on large-scale graphs.

Taken together, these advances highlight both the potential and the challenges of applying GATs for defensive analysis in football. While attention mechanisms can help identify influential defenders, it is likely that complementary methods, such as node masking (Ying et al., 2019), may be required to improve model interpretability and to evaluate the effectiveness of defender influence on attackers. Building on these ideas, a dedicated framework could be developed in football to derive interpretable defensive influence metrics from a GAT model, enabling a more comprehensive evaluation of individual off-ball defensive relationships and contributions. Section 2.6 summarises this challenge as one of the research objectives in this thesis. The following section provides a detailed discussion on the importance of space creation and the current approaches for modelling it.

2.4.1.3 Player Space Creation

The creation and exploitation of space is a fundamental aspect of football play. Off-the-ball movement determines the available options for the on-ball player, impacts the organisation and positioning of defenders, and plays a vital part in the success of attacking sequences. Additionally, minimising dangerous attacking space will be a key objective of the defensive team. Since players spend the majority of a game off the ball, their ability to create and occupy valuable space is a critical aspect of player performance. While coaches have recognised its importance in training and tactical preparation, as well as in player evaluation and recruitment, objective quantification of space creation has only recently become feasible with the emergence of spatiotemporal tracking data.

Early work in this area focused on probabilistically modelling the attacking potential of areas on the pitch. For example, Spearman (2018) uses tracking data to estimate the probability that a player's positioning leads to them receiving and controlling the ball at a specific location on the pitch and the likelihood of them scoring from that location given that they receive the ball. By combining these probabilities, the model can assign a value for their off-ball positioning based on their scoring potential. Similarly, Llana et al. (2020) propose an off-ball metric derived from the deconstructed EPV model introduced in (Fernández et al., 2021), that quantifies a player's probability of getting valuable goal-scoring chances given their off-ball positioning.

Other work has also utilised the concept of pitch control to measure the value of space. Fernandez and Bornn (2018) model pitch control parametrically using logistic functions combined with bivariate Gaussian distributions to compute individual player influence on space, incorporating key spatiotemporal information such as player velocity and distance from the ball. Using this model, methods of defining space generation and occupation are presented to measure how players use space efficiently. For instance, a defender being dragged away to allow a teammate to move into high-value space can be explicitly captured as effective space creation. Consequently, players can be directly compared based on how well they generate space for teammates or effectively occupy space themselves.

More recent developments have used ML techniques that directly assign value to off-ball runs. Gregory et al. (2024) employ an XGBoost model, trained on spatiotemporal features extracted from tracking data, to assign value to player runs based on how it changes the team's scoring probability from the start to the end of the run. Interestingly, their study highlights correlations between run speed and increased offensive value. This offers a scalable, direct approach to evaluating space creation.

Taken together, these models demonstrate that spatiotemporal tracking data can help evaluate how players strategically create and occupy space. Many applications of these methods could provide great value to football coaches and analysts. For example, Spearman (2018) suggests that these models could help teams extract key footage for training without having to manually re-watch games by identifying high-value situations where a team had a dangerous attack or defensive vulnerability. Coaches and analysts may also use these metrics to evaluate individual player performance and contributions, identify successful attacking patterns, or pinpoint particular defensive vulnerabilities. This is particularly useful for key football tasks such as player recruitment and opposition analysis.

Despite these advances in valuing space creation and player positioning, most current research evaluates space retrospectively, identifying where space was created after the fact or where it currently exists. An interesting research gap lies in predicting where and when valuable space may occur in the near future. For example, anticipating that an attack will reach the penalty area within five seconds could allow a striker to preemptively occupy threatening zones rather than reactively adjusting once the ball arrives. An effective approach to this challenge may be to predict an upcoming game sequence and determine the probability of space becoming high-value as the game progresses. Developing anticipatory models of space creation would provide actionable insights for coaching, enabling proactive rather than reactive positioning, and represents a promising direction for future research.

While we've focused so far on evaluating player actions and positioning in football, another interesting area is the optimisation of player decision-making. In the following section, we examine how AI can be used not only to evaluate past actions but also to imitate player behaviour and optimise future player decisions.

2.4.1.4 Modelling Player Decision-Making

Modelling player decision-making represents a promising application of AI and football analysis, aimed at inferring player intentions behind on- and off-ball actions and identifying actions that optimise performance. The advancement of events and tracking data, as well as the development of football simulators, provides opportunities to design agents that can learn directly from real-world match scenarios. Figure 2.8 illustrates the general framework, where spatiotemporal events and tracking data, as well as simulated environments, can be used as inputs to different learning paradigms. These learning paradigms include IL, IRL, and reinforcement learning (RL). These methods can provide unique and valuable outputs, including policies that mimic real-world player behaviour, reward functions that model player intentions, and action recommendations to optimise performance.

IL is typically used to model agent policies directly from expert demonstrations, making it highly suitable for modelling football player behaviour using player actions extracted from spatiotemporal data. An early example for modelling player intentions includes work by Beetz et al. (2005), who developed a multi-functional model that automatically recognises activities in a game using real-time player and ball tracking data. Their system integrates a game model database, an interpreter, a situation miner, and a game analyser, and is formalised using first-order interval temporal logic. This enables the model to infer individual intentions, recognise game situations, and analyse team behaviour

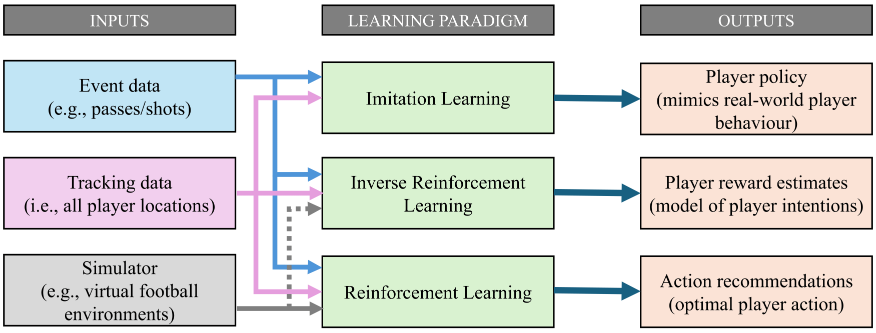

FIGURE 2.8: Framework illustrating learning paradigms (imitation learning, inverse reinforcement learning, and reinforcement learning) applied to the modelling and optimisation of player decision-making in football.

based on play transitions in real time. Deep IL has also been applied to predict expected player movements and to characterise general team patterns, particularly in defensive situations, as discussed further in Section 2.5.1.3 (Le et al., 2017a,b).

More generally, IL has been applied in other domains with spatiotemporal data, such as video games (Harmer et al., 2018) and navigation (Hussein et al., 2018; Ollis et al., 2007; Silver et al., 2008), highlighting its utility in learning from expert trajectories. For example, Silver et al. (2008) use IL to train agents to navigate environments using overhead spatial data, where routes derived from expert demonstration are considered optimal decision-making. These approaches, when applied to football events and tracking data, present the possibility of modelling individual player characteristics and replicating their typical decision-making across various scenarios. Coupling IL with football simulators could allow us to evaluate how different players react under very similar spatial conditions. By incorporating evaluation metrics such as VAEP (Decroos et al., 2019) and EPV (Fernández et al., 2019), comparative tools could be developed to assess and contrast player performance in specific tactical contexts, providing a valuable tool for coaches in team selection and recruitment.

IRL focuses on inferring underlying reward functions that explain observed decisions. Instead of being trained to imitate actions like in IL, IRL learns why players act as they do. For instance, Luo (2020) applied IRL within a Markov game framework for ice hockey, using this approach to estimate fine-grained reward values from player behaviour. This methodology addresses the challenge of goal sparsity in sports like ice hockey and football and demonstrates stronger alignment with key performance metrics, such as points and assists, than traditional action values like xG and VAEP. The methodology is flexible, supporting both single-agent and multi-agent reward function learning, and offers strong potential for adaptation to football to advance the understanding of player behaviour and action values. In football, IRL has been used to extract team intentions from spatiotemporal data. For example, Rahimian and Toka (2021) demonstrated how IRL can uncover the tactical rationale behind actions like passes by leveraging player trajectories from tracking data, providing insights into the tactical decision-making of elite football teams.

In contrast to these approaches, RL aims to optimise decision-making by directly learning action policies that maximise long-term rewards. In football, batch RL and novel football state representation have been applied to recommend optimal actions for maximising xG (Rahimian et al., 2021). The authors expanded upon this method in later work, focusing on evaluating optimal pitch locations using pitch probability surfaces rather than discrete passes (Rahimian et al., 2024b). It was also found that comparisons between real actions and recommended optimal actions serve as an effective predictor for in-game match outcomes (Rahimian et al., 2024a). Beyond individual behaviour, multi-agent deep RL methods (Nakahara et al., 2023) have been introduced to value both on- and off-ball actions in attacking football situations using contextual information from tracking data and an RNN to predict action q-values. These actions include types of on-ball actions (e.g., shots or passes) and discrete movement directions (e.g., left or right). Additionally, Wang et al. (2024) utilise a GAT to predict football corner outcomes, such as receiver and shot predictions, and suggest adjusted player positions for both attacking and defending teams to influence shot probability, further demonstrating the applicability of RL-based methods to tactical recommendations.

The development of simulators has played a key part in advancing these learning paradigms, and in particular, RL methodologies for football. An early simulation environment was the Robocup soccer server (Noda et al., 1998b,a), which facilitated multi-agent research by providing a 2D virtual football setting where agents interact and compete. A more recent simulation environment is Google Research Football (Kurach et al., 2020), a highly optimised 3-D physics engine for football simulation that offers a variety of tasks for agents to complete. However, the complexity of football means that training agents on physics-based simulators requires significant computational resources to achieve high-level decision-making. To mitigate this, Mendes-Neves et al. (2021) introduce a simulator fundamentally built on football event data, enabling agents to focus on learning optimal decision-making rather than basic football skills.

Collectively, these approaches demonstrate how AI can be used not only to evaluate past player decision-making but also to imitate and understand player intentions and propose optimised actions. However, while individual decision-making has received significant attention, the modelling and optimisation of coordination and decision-making on a team level is comparatively underexplored. Understanding how players coordinate and interact as a team is fundamental to overall performance and can lead to opportunities for more effective training processes and tactical optimisations. The following section, therefore, shifts focus to methods for modelling and evaluating teamwork.

2.4.2 Teamwork

Teamwork is a vital determinant of success in football and a natural extension of individual player decision-making. It has been defined as a dynamic process in which team members exhibit both independent and interdependent behaviours to maximise the team's chances of successfully completing a shared goal (McEwan and Beauchamp, 2014). In football, such behaviours include creating passing options, maintaining defensive organisation to limit attacking danger, or communicating information to teammates. For instance, a player may alert a teammate in possession of the ball to an approaching opponent outside their line of sight, increasing the likelihood of ball retention. Importantly, teamwork also involves recognising and compensating for the strengths and weaknesses of individual players, and tactical structures are often designed to maximise the impact of each player while mitigating their weaknesses. For example, a defender with strong attacking ability but weaker defensive positioning may be supported by a teammate who covers their role when they advance forward, thereby enhancing collective attacking potential without compromising the team's defensive stability.

Coaches recognise that cohesive teamwork not only enables teams to execute tactical plans but also provides the potential for individuals to thrive in ways that may not be possible if they were acting completely individually (Carron et al., 2002). As a result, training sessions, player analysis, and player recruitment will carefully consider these factors. For instance, when evaluating potential signings, clubs often assess a player's capacity to integrate into a new environment and develop effective on- and off-field relationships with teammates (Beal et al., 2019). Although teamwork is widely considered a vital aspect of a successful football team, it is more challenging to quantify objectively than individual performance. This is because many of its underlying interactions occur off-the-ball, making them harder to capture in football event data.

Several recent studies have attempted to evaluate teamwork quantitatively and optimise teams based on how effectively players coordinate. Beal et al. (2020b) learn the value of teamwork by modelling chains of directed interactions (i.e., passes) between players using weighted directed graphs. Each chain of passes ends in a terminal event (e.g., goal, shot on target, shot off target, loss of possession), and the contributions of players or pairs of players towards these events are extracted using various graph metrics listed as:

- **Centrality:** The sum of the weights of all incoming and outgoing edges for the player in a passage of play.
- **Distance from event:** The average number of subsequent players in the passing chain after the player until the terminal event.
- **Walk frequency:** The frequency with which a player, or pair of players, appears in the passing chain leading up to the event.

These metrics are weighted by event importance towards match outcome, which is estimated using a logistic regression model. Players and player pairs are assigned an overall teamwork value by summing their weighted event contributions across all events they participated in. Mixed-integer programming is used to form team line-ups that maximise teamwork. This approach significantly outperformed methods that ignore player interactions when predicting real-world team outcomes, achieving a 46% improvement in accuracy.

Bransen and Van Haaren (2020) adopt a similar approach, defining player chemistry based on the joint offensive and defensive contributions of two players on the team. The offensive contributions are derived from the VAEP metric (Decroos et al., 2019) for in-game play sequences involving both players, and the defensive contributions are calculated as the impact the players had on reducing the offensive impact of the opposing attacker. ML regression models are also used to predict the chemistry of player pairs who haven't played together, and the mixed integer model introduced by (Beal et al., 2020b) is similarly used in this work to form an optimal team that maximises player chemistry. Both studies demonstrate that teamwork can be formally modelled and exploited for team selection and player decision-making optimisation; however, these models both rely on the use of event data, restricting the evaluation of teamwork to on-ball actions.

The challenge of modelling teamwork extends beyond football into MAS, where team performance similarly depends on interdependent interactions between agents. A framework that effectively evaluates interactions and coordination in MAS can also provide insight for team-based domains such as disaster response and search and rescue. Despite the importance of teamwork on performance in these domains, research on evaluating agent interactions is relatively limited. Existing models mainly focus on areas like task allocation, agent skills and roles, and spatiotemporal constraints such as task location and completion times (Wu and Ramchurn, 2020; Matthews et al., 2012; Ramchurn et al., 2010; Capezzuto et al., 2020; Amador et al., 2014).

For example, Gaston and DesJardins (2005) explore various network structures of agents and finds the most effective agent networks to maximise performance in dynamic environments, while Arnold and Schwalbe (2002) describe a framework in which individual agents make decisions on which teams to join based on their own strategies and value functions. Scerri et al. (2005) develop a task allocation algorithm for 'extreme teams' with interdependent tasks and strict time windows, such as UAV coordination or hospital resource management. In adversarial contexts, Mutzari et al. (2021) model multi-defender Stackelberg security games, applying cooperative game theory to optimally disperse agents to independent targets.

While these approaches are valuable, they all predominantly focus on task allocation rather than assessing the value of teamwork or spatial interactions between agents.

In football, a team aims to dominate an area against an adversary under significant uncertainty, both in terms of the decisions of teammates and of opponents. These uncertainties differentiate football from scenarios like Stackelberg games, where adversaries are assumed to have complete knowledge (Mutzari et al., 2021).

The first benchmark for valuing teamwork using MAS concepts was the graph-based model introduced by Beal et al. (2020b), which learned from real-world football data. However, this approach focused solely on one-to-one directed interactions between agents using on-ball actions in football, leaving many other forms of teamwork that can exist between agents unexplored. Real-world examples include multiple defenders forming a compact shape to effectively cover space and neutralise an opponent's attack, or a set of agents (e.g., robots) coordinating to maximise space coverage in a patrol task. Furthermore, the approach in (Beal et al., 2020b) does not consider spatial proximity and constraints, which are shown to be important factors for collective performance in many real-world teams (O'Leary and Cummings, 2007).

There are opportunities to advance teamwork modelling by designing approaches that can identify and classify a wide range of interaction types as well as explicitly account for spatial context and its impact on agent decision-making and teamwork. Due to the use of event data as the primary data source in (Beal et al., 2020b), the study is limited to on-ball actions to assess how players communicate and interact, even though most teamwork occurs off-the-ball. The use of tracking data enables richer evaluation of agent interactions such as defensive runs, pressing coordination and space creation. Developing models of teamwork grounded in spatiotemporal tracking data could deliver more accurate insights into key player partnerships and contributions, which, to our knowledge, has yet to be explored. Football analysts and coaches may use these models and insights to inform team selection and tactical planning, as well as enhance post-match analysis by identifying patterns of suboptimal player and team decision-making. We discuss this challenge as one of our research objectives in Section 2.6.

In the next section, we shift focus to team and tactical planning in football. We also review current state-of-the-art TF algorithms in AI and discuss their relevance and potential application to football.

2.5 Team Planning

In this section, we review the evaluation of teams and tactical planning in football, split into pre-game tactical planning (Section 2.5.1.1), in-game tactical changes (Section 2.5.1.2), modelling of team playing styles (Section 2.5.1.3) and injury prediction and prevention (Section 2.5.1.4). In addition, we focus on TF as a broader MAS problem and review current algorithms for TF and planning (Section 2.5.2).

2.5.1 Team Planning in Football

Similar to player performance analysis, team evaluation and tactical planning are fundamental processes taken by coaches and analysts in football that shape how clubs prepare and adapt during matches. While broader organisational factors, such as appointing a head coach, forming a recruitment strategy informed by club budget, designing a player scouting process, and incorporating academy players determine the long-term identity of a team, the immediate tactical setup for any given match follows an iterative cycle of inputs (e.g., team and opponent information), pre-game planning, in-game tactical changes and post-match analysis (see Figure 2.9).

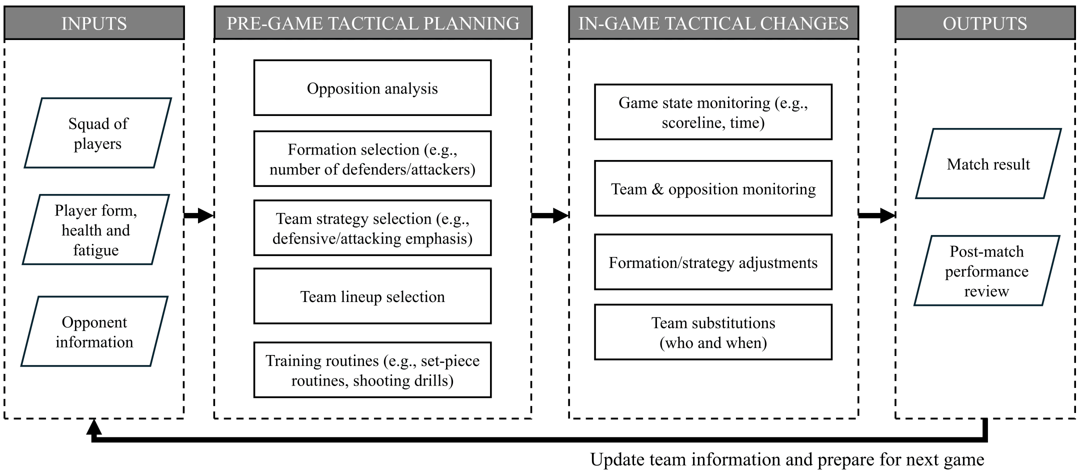

FIGURE 2.9: Framework diagram of the team planning process in football.

Inputs The team planning process begins with assessing various inputs regarding the squad. These inputs include the players in the squad, the current form of each player, player fatigue and fitness levels, and detailed information on the opponent. These contextual factors constrain and inform the planning and decisions available to coaches. Inputs regarding past player performance may include derivations and analytical metrics from the models discussed in Section 2.4.

Pre-game tactical planning Given these inputs, coaches and analysts complete several tasks to prepare for the upcoming match optimally. These tasks include an opposition analysis to learn their strengths, weaknesses and key threats. Additionally, they select a team formation (e.g., number of defenders or attackers) that best suits the desired playing style and players in the squad. Furthermore, they define strategic emphases such as defensive compactness or attacking intensity. They also choose a starting lineup. Training drills, such as rehearsed set-piece routines, are also designed and practised to help the team prepare, given the anticipated tactical requirements for the game.

In-game tactical changes Throughout the duration of a game, coaches continuously monitor the game state (e.g., scoreline, time remaining) as well as the performance and behaviour of the team and opponent, and will consider making tactical changes. In-game tactical optimisation involves adjusting formations or strategies and making player substitutions in response to the opponent, which are expected to influence team performance positively. Substitutions, in particular, can be made for various purposes, including strategic (e.g., adding attacking players when losing the game) and restorative (e.g., replacing fatigued players who may be at risk of injury).

Outputs and Post-Match Analysis Tactical decisions ultimately contribute to the final match result, and they can be reviewed post-match to evaluate their effectiveness. These evaluations highlight what tactics or players had a positive effect, reveal team or tactical weaknesses, and inform future team planning. This creates an iterative cycle in which each match contributes knowledge that shapes the team's preparation for future games.

Research into AI methods has been applied across all stages of this process. Event and tracking data have facilitated advances in opposition analysis, strategy and formation selection, in-game decision-making (e.g., optimal substitutions), and post-match performance evaluation. In the following subsections, we examine state-of-the-art methods in four key areas: pre-game tactical planning, in-game tactical changes, modelling team styles, and injury prediction and prevention. The practical benefits of these methodologies for football teams, and key areas for future research, are also discussed.

2.5.1.1 Pre-Game Tactical Planning

Pre-game tactical planning involves preparing a team to maximise performance against a specific upcoming opponent. The process typically includes an analysis of the squad, including an evaluation of player availability, health, fatigue, form and chemistry. This process also includes an analysis of the opposition, including their key strengths and weaknesses. During this planning phase, using information about the squad and opponent, the manager must decide upon a game strategy, formation, starting lineup and set of player instructions. The modelling and optimisation of each of these decision points have been focused on in football analytics, especially with the increased information on team performance and dynamics provided by the emergence of spatiotemporal event and tracking data.

One line of research has modelled the tactical preparation of teams for games as a decision problem under uncertainty. Beal et al. (2020a) formulate pre-game planning as a Bayesian model, where each team chooses actions related to their tactical style and formation under incomplete information about the opponent's tactics. Multiple approaches are reviewed for this Bayesian model to optimise tactical decision-making, such as the best response method, which maximises win probability given the uncertainty around possible opponent actions. ML approaches such as k-means clustering and deep neural networks are used to classify team playing styles and predict match outcome probabilities. The key objective of this work is to determine optimal tactics by integrating information on team styles extracted from spatiotemporal event data. The model is validated using two seasons of EPL data, finding that team tactics selected by the best-response optimisation method increase win probability by an average of 16.1%.

Beal et al. (2021) extend on the pre-match Bayesian model to optimise tactics long-term by setting fluent season objectives for teams. These objectives update iteratively as each new game provides new evidence to the Bayesian model. This approach captures how real-world football clubs realistically adapt their team strategies over time as more information on the team's league performance and their strengths and weaknesses becomes apparent. The authors found that this approach can improve teams' mean expected league position distribution by up to 35.6%. These models encompass many aspects of pre-game tactical planning, including opposition analysis, team formation selection, and strategy selection.

Event and tracking data have been widely used to analyse team strategies and formations. Decroos et al. (2018) use event data to design an automated model for reviewing team tactics by detecting common patterns of play. This model clusters event sequences according to ball movement, identifying common team attacking strategies and ranking these strategies based on how regularly they lead to shots. Such models provide coaches with an automated way of assessing which play strategies are most effective and identifying an opponent's most preferred attacking patterns. Goes et al. (2021b) investigate the synchrony of player interactions within play patterns using player tracking data. This approach divides a team into smaller sub-groups and analyses which interactions between these sub-groups are most frequently associated with successful attacking outcomes, thereby providing insights into effective tactical configurations. A key finding is that successful attacks frequently involved defenders moving asynchronously to create space for attacking players and disrupt the opposition's organisation. These insights have practical implications for coaching, as training could emphasise the rehearsal of movement patterns that are empirically shown to improve attacking success probabilities.

Tracking data has also been used to automatically detect formations and playing styles. Bialkowski et al. (2014a) use expectation–maximisation to infer latent player roles, assigning players probabilistically to roles based on the learned spatial distributions of their mean positions in tracking data. They show that teams often play the same formation both home and away, but alter their attacking emphasis within the formation. In similar work, Bialkowski et al. (2014b) categorise team styles and formations using k-means clustering from tracking data-based features, finding that this methodology can represent team behaviour three times more descriptively than traditional match statistics (e.g., goals, passes, shots on target). These methods offer potential as effective opposition analysis tools for identifying expected formations and playing styles for upcoming games.

The studies reviewed above show how event and tracking data can be used to evaluate teams and tactics in pre-game analysis; however, there remain some areas for future work. In particular, while event data has been used to provide information for tactical optimisation, its descriptive power is constrained by the lack of information on off-ball players and dynamics. Tracking data offers richer spatiotemporal features, which provide many advantages for formation detection, structural analysis of teams, and the evaluation of player coordination. However, these methods have not yet been extensively integrated into tactical optimisation frameworks. A promising area for future research lies in combining the Bayesian decision-making approach presented in (Beal et al., 2020a) with the automated formation and team-style detection model in (Bialkowski et al., 2014a), providing increased descriptiveness of opponent style for the tactical optimisation process. Furthermore, current tactical optimisation approaches pay relatively little attention to player workload, fatigue and injury risk, and how these can inform pre-game team selection, despite these being critical constraints in real-world team management, particularly in congested match schedules. Incorporating physical considerations into the optimisation process would strengthen the practical applicability of pre-game tactical planning models.

Designing AI systems to recommend optimal tactical setups is a complex computational challenge, both pre-game and in-game, due to the inherent uncertainty and adversarial dynamics of football. However, sophisticated models in other domains may help provide valuable insights. Game-theoretic approaches, for example, explicitly model interactions against adaptive opponents with the aim of maximising game reward. A particularly useful area may be Stackelberg games, which represent leader–follower scenarios where one agent commits to a strategy and the adversary responds with knowledge of that choice (Sinha et al., 2018). Stackelberg games have been successfully applied in areas such as wildlife protection, where optimal patrol routes are determined using terrain and animal distribution data to maximise defensive coverage (Fang et al., 2016), and transport security, where patrols are scheduled under spatiotemporal constraints to optimally deter commuters from evading fares (Yin et al., 2012). These approaches may present a benefit to the football domain by recommending optimal defensive spatial setups in response to predicted opponent attacking strategies given by formation and style detection models (Bialkowski et al., 2014a; Decroos et al., 2018). These areas have clear potential to inform coach decision-making in lineup and formation selection. Existing TF models across the broader AI literature are discussed further in Section 2.5.2. The following section reviews how AI modelling has been applied to model game progression and support the evaluation and optimisation of in-game tactical changes.

2.5.1.2 In-Game Tactical Changes

In-game tactical planning requires coaches to adapt rapidly to a highly dynamic environment, where decisions must be made quickly and effectively before the game state shifts. The process involves continuously assessing team performance and player condition, monitoring opposition behaviours, tracking evolving game states (e.g., scoreline or game momentum), and responding to these aspects through appropriate tactical adjustments or substitutions. The emergence of spatiotemporal event and tracking data has led to opportunities for deeper real-time analysis of match progression when optimising in-game actions, extending beyond just core match statistics to measures of spatial dominance, pressing intensity and workload management. Developing AI systems capable of supporting or even automating elements of in-game tactical decision-making represents a major challenge in football analytics, with many significant potential benefits.

In-game tactical optimisation approaches have been designed using different levels of football data. For example, Hirotsu et al. (2009) propose a game-theoretic framework for evaluating tactical changes, such as formation adjustment, using only match-level data. In their approach, football is modelled as a zero-sum game based on the tactics employed by both teams, and scoring probabilities are derived from Poisson distributions that represent each team's attacking and defensive strength, with parameters estimated using the maximum likelihood method. This model demonstrates that tactical adjustments, which increase a team's likelihood of winning, simultaneously reduce the opponent's probability of winning. The authors provide an extensive analysis of the impact of formation changes within the J-League (Japan's top football division). While insightful, this work is restricted to aggregate match data. Incorporating event or tracking data could help capture the spatiotemporal dynamics of player and team behaviour, allowing for more detailed evaluation of the impact of tactical changes on team structure and performance.

Spatiotemporal data has allowed for more in-depth monitoring of player condition, performance, and risk factors such as fatigue and injury in-game. Additionally, it provides for the evaluation of the opposition, including notable tactical changes, dangerous players, or exploitable weaknesses. For example, Castellano et al. (2013) use player tracking data to review team structure based on the length, width and surface area of a team formation at given phases of a game, showing that more dominant teams often have wider and longer team structures. These insights could help coaches determine which team may be in the most advantageous position in-game. Goes et al. (2021c) adopt a slightly different approach, highlighting the necessity for tactical changes through in-game match prediction. To calculate this, models are trained on spatial match features such as pass length and shot distance to predict binary match outcomes in real time. These models show strong predictive performance, ranging from 67.1% to 74.1%

accuracy in match prediction depending on match time. Such systems demonstrate the potential to automatically alert coaches when tactical adjustments may be required.

One of the most influential and important in-game decisions available to coaches is the timing and personnel of substitutions. Early work by Hirotsu and Wright (2002) modelled football as a four-state Markov process using match data and used a log-linear model to estimate state transition probabilities. Dynamic programming is then used to find the optimal substitution timing by maximising the expected match points. Myers (2012) and Rey et al. (2015) applied data mining on decision trees to evaluate factors such as the quality of opposition or the scoreline when recommending optimal times for substitutions. Beal et al. (2020a) extend their pre-match Bayesian model by representing a football match as a stochastic game. In this framework, states represent the current scoreline, and transitions, whose probabilities are estimated from historical data on team strengths and styles, represent goals scored. Teams can select actions, corresponding to player substitutions, to maximise their likelihood of transitioning to more favourable match outcomes. While these models rely mostly on match-level and event data, they highlight the potential value of in-game substitution optimisation. The increasing availability of tracking datasets could significantly improve these models by giving more granular overviews of the current game state. For example, pitch control metrics (Fernandez and Bornn, 2018; Fernández et al., 2019) could be used to identify which areas of the pitch are being most heavily dominated, and potentially exploited by the opposition, and passed as input features to ML models trained on historical substitution data. Linking substitution timings with their spatiotemporal impact across previous matches would allow such models to learn patterns of effective tactical changes, providing coaches with data-driven recommendations on both which players to substitute and when.

The challenge of in-game tactical optimisation in football can be modelled as a real-world representation of real-time decision-making in MAS, where agents must respond adaptively to dynamic and uncertain environments. Concepts extracted from MAS research, particularly for coalition formation, role adaptation and task reallocation under uncertainty, can therefore offer valuable ideas and mechanisms for in-game decision-making models. For instance, Arnold and Schwalbe (2002) introduced a dynamic coalition formation model in which agents choose the team they want to join using a revised individual value function that updates based on the current coalition structure and observations of other agents. Similarly, Jones and Barber (2008) evaluate various TF strategies in highly dynamic environments with incomplete information, such as open markets or complex business organisations. Katayanagi and Sugawara (2011) study an agent network where tasks emerge spontaneously and communication delays exist between agents. The authors propose a reorganisation mechanism that adapts network topology based on workload to balance agent activity effectively. Other examples have been presented for sequential coalition formation where tasks arise spontaneously and agents are uncertain about their team members' skills and roles (Chalkiadakis and Boutilier, 2012), and for agent role reallocation in response to agent failures in search and rescue contexts (Gunn and Anderson, 2015). These approaches address challenges similar to those in in-game football decision-making, where events such as injuries, changes in game momentum, or player performance require dynamic adaptation.

In addition to exploring applications of current MAS approaches, extending the in-game stochastic game model of football presented in (Beal et al., 2020a) by increasing the granularity of in-game states from scorelines to in-depth team performance and match dynamics metrics extracted from tracking data could improve in-game decision-making performance. For example, integrating metrics that measure the quality of scoring opportunities (Spearman, 2018) and evaluating player teamwork and performance could enable earlier identification of emerging tactical vulnerabilities, allowing for intervention before they lead to negative outcomes, such as conceding a goal. Moreover, a framework that considers both agent performance and long-term fatigue and injury risks could further improve decision-making, with potential applications extending beyond sport to domains such as emergency response and search and rescue. Interesting parallels also exist with adaptive AI agents in video games who must use experience to improve their strategies in dynamic environments with opposing players (Demasi and Adriano, 2003; Spronck et al., 2006). For example, Spronck et al. (2006) use dynamic scripting, a stochastic optimisation method that adapts an agent's policy based on their in-game experience to optimise their responses to changing environments. Modelling football teams as adaptive agents in this way could lead to match strategies that evolve proactively based on in-game experience.

A key barrier to the production of in-game tactical analysis and optimisation models is the requirement for live gathering and feature extraction of spatiotemporal data, which is still uncommon in football. However, as real-time data collection becomes more widespread in the sport, algorithms like those discussed in this section could provide coaches with actionable recommendations during matches. This could help support decisions on substitutions, formation alterations, and attacking strategies, thereby offering the potential for a competitive edge in matches. In the next section, we focus on how AI can be applied to model team styles and structures, forming a deeper perspective of team behaviour.

2.5.1.3 Modelling Team Styles and Structures

Modelling team styles in football using AI typically involves learning patterns and policies adopted by teams through spatiotemporal data. Football tracking data is particularly valuable as it enables fine-grained analysis of team structure and setup, resulting in deeper insights into collective play. Such modelling supports more effective opposition analysis and a better understanding of team tactics. As with player modelling discussed in Section 2.4.1.4, football simulators can act as test beds for AI agents, where learning approaches are applied to identify optimal team styles and structures in given situations. When modelling a whole team, this becomes a multi-agent problem in which the interactions between agents must also be considered.

Several studies use AI learning approaches applied to player and team trajectories to model the style of play. Le et al. (2017a) apply deep IL to anticipate teams and opponent movements throughout a match, using an LSTM with feature vectors containing absolute player coordinates and their positions relative to the ball and goal. The learner retains information on past positions and roles, using it to project future positions over subsequent time steps. This approach is particularly useful for anticipating team actions and preparing for these, whether by finding effective attacking routes to break down defensive structures or by defensively preparing for common attacking patterns. Building on this work, Le et al. (2017b) focus specifically on predicting defensive behaviour. Using LSTMs trained on data from 45 league matches, they learn policies for ten defensive roles (the entire team excluding the goalkeeper), building a model of an entire team's defensive patterns. In contrast, Van Roy et al. (2023) focus on offensive behaviour, modelling play as an MDP where pitch zones represent states and actions consist of passes or shots. By learning policies from historical data, this approach captures team-specific attacking tendencies and provides a probabilistic outlook on likely future actions. Bialkowski et al. (2014c) model player positioning and team tactics by applying expectation-maximisation to estimate player role distributions across matches using tracking data, predicting team formations with an accuracy of 75.33%.

Other work has focused on using football event data (e.g., passes and shots) to model team style. For example, Lucey et al. (2012) model team behaviour using fixed time windows of ball trajectories extracted from event data for a team. Team behaviour is categorised using entropy maps, which are used to compute team predictability. This offers practical applications for match preparation, enabling teams to anticipate opponent styles or to evaluate their own unpredictability in disrupting opposing defences. Additionally, Rahimian and Toka (2021) employ IRL with convolutional RNNs to infer team intentions behind actions (e.g., passes) from player trajectories. Moving to an approach applied in ice hockey and basketball, Mehrasa et al. (2018) use convolutional neural networks on player tracking data to classify events such as shots and categories teams based on their collective movement patterns. While this approach is tested in ice hockey and basketball, the approach could be extended to football to improve the state-of-the-art for modelling team behaviour based on their collective spatial dynamics.

Together, these works discussed in this section demonstrate how spatiotemporal data can offer deeper insights into modelling team tactics and formations, not only by identifying structural patterns but also by inferring underlying strategic intentions. These models can support match preparation by enabling coaches to evaluate their own team's and the opponents' strategies. This may also enable opportunities to test different strategies for combating opposition tactics in simulated environments. Smaller clubs may also benefit by modelling their team around tactical intentions extracted from elite teams through publicly available spatiotemporal datasets. There are also open research areas for these approaches; for example, team behaviour models that use match events could be applied to estimate player tracking data using only event data, providing a cost-effective alternative to expensive tracking systems.

Advancements in multi-agent learning from other domains further highlight opportunities for football. RL applied in video games has achieved professional-level performance, as shown in (Berner et al., 2019) and (Vinyals et al., 2019), where agents learned complex game strategies through their experiences of spatial surroundings and in-game objects. Similar approaches in air traffic management (Brittain and Wei, 2019) and multi-agent path planning (Sartoretti et al., 2019) demonstrate how multi-agent RL and multi-agent IL can enhance spatial coordination among agents and reduce conflicts. These examples show how agent learning can improve spatial awareness, highlighting applications to football due to the importance of spatial positioning and coordinated defensive and offensive interactions for team success.

Football simulators could be deployed to create environments where agents can test tactical strategies, enabling experimentation with different team styles against automated models of opposition behaviour. For example, automatic formation and team style detection methods (Bialkowski et al., 2014b) could generate opposition profiles for simulated match testing. Similarly, IL methods such as those in (Sartoretti et al., 2019) could be used to model team play styles from spatiotemporal data, simulating player actions based on observed behaviour. This combination offers the potential to optimise team setups and tactical styles by testing strategies against predicted opponent styles in simulation. A further step would also be to integrate multi-agent learning and tracking data into the creation of a digital twin of a football team, simulating matches in real time and trialling potential strategic decisions. Such a system could provide coaches with automated in-game tactical insights, including recommendations for substitutions or formation changes based on detected sub-optimal team patterns. In the next section, we focus on injury prediction and prevention.

2.5.1.4 Injury Prediction and Prevention

Across all contact team sports, including football, injuries to players can reduce long-term team performance significantly (Hägglund et al., 2013). If a team's most vital players are injured for many games across the season, their expected performance in these games decreases. Furthermore, a team with multiple injuries may be forced to play the same players repeatedly due to a constrained squad, which could lead to additional team fatigue and injury. The financial impact of injuries is also significant due to large players' wages and the financial reward received for achieving team objectives (e.g., winning the league or avoiding relegation). A report by JLT found that player injuries in the 2018/19 EPL season cost a total of £221 million in wages.[15] In professional football teams, it is common to have congested season schedules with multiple matches each week, increasing the risk of injuries to players due to fatigue and the requirement for the team to rest players and improve long-term team performance (Howle et al., 2020; Bengtsson et al., 2013; Dellal et al., 2015).

The impact of injuries on team performance and finances highlights the importance of team planning to minimise the long-term issues caused by player injury. In football, managers may rest a tired player to allow them to recover for the next game. Furthermore, managers may play a weakened team against an opposition they are expected to beat, resting multiple key players for a more difficult game later. Despite this, it is difficult to know if these team planning decisions are optimal, as there are many factors involved that are unknown. These factors include the injury probability of players, the match outcome probabilities based on the selected team, and the long-term benefit of resting a player. The use of AI to optimise decision-making in this area may help significantly reduce injury in key players and improve team performance. However, to date, there is limited research in this area (Beal et al., 2019). To optimise team and injury management, an effective and reliable injury prediction model must first be implemented.

Injury prediction is a difficult problem due to the many factors that may cause injury. Orchard and Powell (2003) study the effect of weather conditions on injury risk in American football, and find that knee and ankle injury risk is lower in colder conditions. Kucera et al. (2005) study the impact of injury history on future injury risk in youth soccer using multivariate generalised Poisson regression. Results of this study indicated that injury risk increased significantly for players with an injury history, with the risk being three times greater for those with two or more previous injuries. Several studies consider the impact of the Acute:Chronic Workload Ratio (ACWR) on injury risk (Malone et al., 2017; Bowen et al., 2020; Cummins et al., 2019; Hulin et al., 2016). The ACWR is widely used in sports science to model the fitness and fatigue of players by comparing a player's recent load (acute workload) with their longer-term load (chronic workload). At day t, the ACWR is defined as:

$$\mathrm{ACWR}_t = \frac{\mathrm{AcuteLoad}_t}{\mathrm{ChronicLoad}_t} = \frac{\frac{1}{7}\sum_{i=t-6}^{t}\mathrm{Load}_i}{\frac{1}{28}\sum_{i=t-27}^{t}\mathrm{Load}_i} \tag{2.1}$$

Where $\mathrm{AcuteLoad}_t$ is the average player load over the previous seven days and $\mathrm{ChronicLoad}_t$ is the average load over the previous 28 days. Player load can be calculated using various methods, such as the product of the player's session Rate of Perceived Exertion (sRPE) (Haddad et al., 2017) and the session duration, or by their total distance covered

[15] https://www.si.com/soccer/2019/08/17/how-much-each-top-6-club-spent-injured-players-during-201819-season.

during the session. Physical metrics, such as total distance covered and the total number of sprints, are also typically collected for players in games using tracking systems or wearable technology. These metrics are also commonly used as a measure of player fitness and load. The ACWR compares player fatigue to their fitness base, giving insight into the risk of abnormal body stress levels. Bowen et al. (2020) find that when a player's ACWR spikes above two, their injury risk increases dramatically. Malone et al. (2017) present that an ACWR of between 1 and 1.25 reduces the injury probability of players most. Other factors, and their impact on injury risk, have also been studied, such as in-game playing actions (Rahnama et al., 2002) (e.g., tackles or passes), playing surface (Ekstrand and Nigg, 1989) and training program (Kraemer and Knobloch, 2009).

Using AI to reliably learn the relationships between these factors and injury risk in players could be very useful as an advisory tool for managers and coaches in football. The emergence of in-depth tracking datasets, wearable sensors and medical data in recent years makes the task of curating features to train an ML model to predict injury risk more solvable (Beal et al., 2019; Kempe, 2025). These features can capture key factors such as a player's injury history, workload, intensity and fatigue. Due to a relatively low injury incidence rate in football, many games of data would likely be required to train an ML model to predict injury probability reliably without overfitting.

Several papers have studied the use of ML to predict injury in football (Majumdar et al., 2022). These approaches vary in dataset, features, and the model used. Rossi et al. (2018) use GPS data from training sessions as input to a decision tree to predict injury for twenty-six players over a single season and report a 'Compute Area Under the Receiver Operating Characteristic Curve' (AUC) score of 0.76. Rommers et al. (2020) predict the injury risk of 734 players at the youth level using physical features such as the player's height, weight, speed and flexibility. An XGBoost model is trained using this data as input and predicts injury on test data with an accuracy of 85%. Jauhiainen et al. (2022) predict the risk of anterior cruciate ligament injuries in female football and handball players using eight years of data. The athletes underwent many tests and a questionnaire to assess strength, balance, flexibility and other physical attributes. This data was used as input to a support vector machine (SVM), which was compared to Logistic regression and random forest as baselines to predict injury, with an AUC score of 0.63 reported. Lu et al. (2022) predict injury risk in basketball using twenty seasons of historical data and compare predictive performance for various ML models, including XGBoost, a neural network, random forest and logistic regression. They note that XGBoost achieves the highest predictive performance with an AUC score of 0.84. They also study the features used as input for the model and identify that past injury is a key factor in future injury risk. We summarise the studies identified here towards modelling injury risk in sports using various key features and approaches in Table 2.2.

There remains scope to study further the best features and ML models to use to integrate an injury risk model into a club's decision-making processes. It is key that the model

TABLE 2.2: Summary of the discussed injury modelling studies by sport, primary features examined, model type, and citation.

Sport	Key features	Model used	Paper
American football	Weather conditions (e.g., temperature, rain)	Statistical analysis	(Orchard and Powell, 2003)
Football	In-game actions (e.g., tackles, passes)	Statistical analysis	(Rahnama et al., 2002)
Football	Playing surface (e.g., grass, artificial turf)	Statistical analysis	(Ekstrand and Nigg, 1989)
Football	Balance training program duration	Statistical analysis	(Kraemer and Knobloch, 2009)
Football	Overall training workload (e.g., weekly load, ACWR)	Logistic Regression	(Malone et al., 2017)
Football	GPS workload (e.g., sprint distance, ACWR)	Logistic Regression	(Bowen et al., 2020)
Rugby	GPS workload (e.g., total distance, ACWR)	Logistic Regression	(Hulin et al., 2016)
Football (youth)	Injury history (e.g., sessions and games missed)	Poisson Regression	(Kucera et al., 2005)
Football (training)	Training workload (e.g., sprints, distance, ACWR)	Decision tree	(Rossi et al., 2018)
Football & handball	Physical tests (e.g., balance, strength, flexibility)	SVM	(Jauhiainen et al., 2022)
Football (youth)	Physical tests (e.g., endurance, agility, speed)	XGBoost	(Rommers et al., 2020)
Basketball	Prior injury history	XGBoost	(Lu et al., 2022)

is explainable to ensure that coaches and players can trust and understand its outputs. Furthermore, it allows coaches to comprehend the factors contributing to injury risk and may enable them to organise targeted training sessions to improve player welfare efficiently. As an example, we propose that a Bayesian network may be a suitable injury risk model due to its natural ability to show cause-and-effect relationships between features influencing injury risk. Furthermore, the model could balance expert knowledge with evidence from new data to form its predictions.

While there have been some studies into the use of AI to predict injuries in football, to our knowledge, the use of AI to plan team selections and substitutions to proactively and optimally prevent injury is yet to be investigated. Football may greatly benefit from an objective, data-driven approach to team selections that uses player injury probabilities to maximise players' long-term contributions to the team. This would provide a method for a club to utilise physical data on their players and apply it directly to their decision-making processes to improve team performance and reduce the financial cost of injuries. In the next section, we turn our attention to state-of-the-art TF models across all domains.

2.5.2 Team Formation Models and Algorithms

In many real-world domains, agents with individual properties (e.g., roles and skills) are teamed up and distributed to complete tasks that contribute to a common objective (e.g., maximising reward or minimising risk). TF has been studied across a variety of settings, including the assembly of optimal response teams in disaster response (Ramchurn et al., 2010; Capezzuto et al., 2021) or the selection of optimal Fantasy Premier League teams (Matthews et al., 2012). Team sport shares similarities with these domains, where teams must be chosen repeatedly from a set of players each week to maximise a season objective such as total points (Beal et al., 2019).

Multi-agent TF has been extensively studied in dynamic and adversarial environments, both in deterministic settings and under uncertainty. Some approaches focus on optimising team performance under various constraints, such as spatio-temporal limitations or task requirements (Ramchurn et al., 2010; Zhang et al., 2025). Previous work has also addressed dynamic scenarios where tasks emerge spontaneously and agents can form teams in a decentralised manner based on individual value functions and network structure, with applications such as supply chains and sensor networks (Gaston and DesJardins, 2005). Chalkiadakis and Boutilier (2012) examine the trade-off between exploration and exploitation in coalition formation, employing Bayesian RL and partially observable MDPs to address environment stochasticity and develop effective coalition policies. In simulated search-and-rescue scenarios, Gunn and Anderson (2015) evaluate a TF framework for dynamic environments in which desired team composition and task suitability are defined by human experts, and robots can switch roles and merge or redistribute teams during an operation to match these targets.

Coalition formation under spatial and temporal constraints has also been studied extensively. Ramchurn et al. (2010) introduce the "Coalition Formation with Spatial and Temporal Constraints Problem", which involves assigning tasks to agents in dynamic environments while considering factors such as spatial distribution and task completion deadlines, with applications in disaster response scenarios. A mixed integer program and a set of anytime heuristics were used to solve the constrained optimisation problem. Building on this foundation, Capezzuto et al. (2020) propose the Cluster-based Task Scheduling algorithm, an efficient anytime method that guarantees convergence. Further extensions to this problem have been conducted in (Capezzuto et al., 2021), which improves the scalability of solutions by presenting a distributed adaptation of the Cluster-based Task Scheduling approach. These methods use deterministic formulations and react to emerging events; they do not plan explicitly under uncertainty, such as probabilities of agent loss or unavailability.

The use of search algorithms and RL techniques to solve TF problems has been explored in the literature (Matthews et al., 2012; Chalkiadakis and Boutilier, 2012; Wu and Ramchurn, 2020; Kartal et al., 2016). For example, Matthews et al. (2012) use Bayesian Q-learning for sequential team selection in fantasy football, modelling the process as a partially observable MDP with action constraints reflecting game rules such as squad size, budget, and role requirements (e.g., a goalkeeper is required). This constrained optimisation problem is treated as a multi-dimensional knapsack packing problem and sets a benchmark for sequentially optimal TF. In simulated disaster response, Wu and Ramchurn (2020) use Monte Carlo tree search (MCTS) as a scalable and anytime method for selecting optimal sets of coalitions over many combinations of agent teams and team sizes.

Robustness to agent failures and unavailability in TF has been explored in simulated settings (Gunn and Anderson, 2013). For example, robust TF has been formally defined as the problem of selecting efficient teams that can cover a range of skills for a task, even after a number of agents have been removed from the team (Okimoto et al., 2015). A team is considered k-robust if k agents can be removed, and the team can still complete the appointed task. Similarly, recoverable TF considers the cost of recovering a team's ability to complete a task if a number of agents were to be removed from the team due to failure or injury (Demirović et al., 2018). Algorithms are given in these studies to find optimally robust and recoverable teams. Gunn and Anderson (2015) address robot failure by allowing the remaining team members to adapt to fill missing roles, and permitting lost robots to join new teams. Anagnostopoulos et al. (2012) consider workload management by minimising the number of tasks assigned to each agent in a social network.

In adversarial settings, strategic TF is often modelled using Stackelberg games, where a decision-making agent commits to a strategy that an adversary observes and responds to (Paruchuri et al., 2008). These games have found application in domains such as security and wildlife protection. For example, Shieh et al. (2012) optimise port security patrol schedules to anticipate potential attacks, while Wang et al. (2019b) employ repeated Stackelberg games to dynamically adjust wildlife patrols in response to adversary behaviour. Research on repeated security games further explores how defenders adapt to attacker strategies over time. For instance, Nguyen et al. (2019) examine how attackers can deceive defenders, highlighting the importance of adaptability, while Alcantara-Jimenez and Clempner (2020) model decision-making uncertainties in security environments using partially observed Markov games.

Several TF models address dynamic agent reallocation in domains without adversaries. For example, Shriyam and Gupta (2018) proposed a model that reallocates teams to new tasks to minimise task completion times, allowing robots to be reassigned mid-task. Similarly, Liu et al. (2022a) utilise multi-agent RL to optimise the dynamic reallocation of UAVs to tasks as new information and objectives emerge. Cohen and Mouaddib (2018) employ MCTS for team reformation in disaster response scenarios, enabling teams to rearrange agents and roles as tasks evolve. These studies advance dynamic TF, but generally do not focus on adversarial domains where agent fatigue, unavailability risk, and adaptive opponents must be managed.

The TF literature spans a variety of approaches, with existing models capturing important aspects of real-world TF problems. Some focus on adaptation to opponents in adversarial settings (e.g., Stackelberg games), others model repeated interactions across multiple incidents, while robustness-focused approaches ensure teams can complete their objectives following agent loss or failure. Reactive methods, by contrast, reconfigure teams after agent faults have occurred. However, there remains a gap in research for a model that integrates pre-event team selection with in-event team adaptation, and explicitly accounts for stochastic processes such as fatigue and injury risk from workload that impact agent availability and team performance over time.

These challenges are vital in team sport, but also extend to many real-world domains beyond sport. For example, emergency response teams must schedule shifts to avoid exhaustion while maintaining high task completion success, while in search and rescue, agents must maximise coverage while limiting the probability of battery depletion. In these settings, teams may need to change mid-event to meet the event's requirements. For example, a patient's health might deteriorate quickly during an emergency response, or battery levels could drop faster than anticipated, requiring UAVs to reroute in search and rescue. Similar requirements exist in adversarial domains such as team sports and security, where opponents adapt and the team must respond quickly and effectively. In such settings, teams must be selected proactively and be flexible enough to adapt effectively during events. There is potential for models that balance unavailability risk with the reward of deploying skilled agents both pre-event and in-event across repeated interactions to improve long-term team outcomes and reduce agent injury or failure. We address this challenge as one of our research objectives in Section 2.6.

Football provides a suitable real-world testbed for these ideas. Managers face congested schedules and varying opponents, and must plan pre-game selections and in-game substitutions to maximise performance while accounting for player workload and fatigue. The sport offers rich data on workload, injury history, and outcomes, with player injury risk varying based on these factors and the game context. Player injuries in football cost clubs millions of pounds every season in the EPL[16] and can lead to reduced season performance due to long-term injuries to key players. The sport, therefore, combines adversarial dynamics, repeated interactions, pre-game and in-game team selections, and the practical need to manage player risk, making it an ideal context for developing models of TF under uncertainty. In the next section, we summarise the literature review and discuss how open areas of research relate to the research objectives addressed in this thesis.

[16] https://www.howdengroupholdings.com/news/howden-2022-23-mens-european-football-injury-index

2.6 Discussion

This chapter has reviewed the landscape of spatiotemporal prediction models, performance evaluation frameworks, and team planning methodologies, with a particular focus on their application in football and related domains such as other team sports. A recurring theme in football research is the predominant focus on on-ball actions, with many opportunities remaining to leverage player tracking data for greater contextual richness and analytical depth. Across the research areas we surveyed, several open research questions have emerged, many of which reflect broader challenges in AI and MAS that remain underexplored in the academic literature. Furthermore, applying existing AI methods to football requires careful consideration of the sport's inherent complexity and the adaptations necessary for effective implementation.

Advancing state-of-the-art AI models and algorithms in football offers substantial benefits to coaches and practitioners. These include streamlining training by identifying strengths and weaknesses in both players and team structure, enabling more objective player recruitment through robust comparative metrics, and enhancing opposition analysis by identifying opponent team styles and key tactical areas. The importance of improving these processes is underscored by the significant financial and social incentives associated with team success. Additionally, the insights gained from AI research in football benefit the wider AI community by providing a platform to test spatiotemporal models in a real-world environment with clear objectives and large amounts of data spanning multiple years, thereby informing potential advances in team-based AI across other domains.

Below, we summarise the key open research areas identified in this review, highlighting football as a valuable test bed for advancing these areas.

Spatiotemporal Prediction Spatiotemporal prediction models have been widely applied to forecast and impute agent trajectories across various domains, including team sports (Yeh et al., 2019), vehicle collision avoidance (Xie et al., 2021), and pedestrian movement in crowds (Marchetti et al., 2020). A range of modelling approaches have been explored for these tasks, such as Kalman filtering (Kim et al., 2013), LSTMs (Sriram et al., 2020), Bayesian models (McInerney et al., 2012), GNNs (Liu et al., 2022b), and VRNNs (Sun et al., 2018). Despite this diversity in approaches, several important research areas remain underexplored.

A primary open challenge is the development of spatiotemporal prediction models that can perform effectively in environments characterised by sparse and irregular agent observations (Figure 2.5). Many real-world domains, such as disaster response scenarios with limited sensor coverage or team sports with event-based data streams, provide only intermittent and non-uniform agent observations. This issue is particularly important in football, where on-ball event data yields sparse player information with non-uniform timesteps. Imputing agent trajectories under these conditions is an important challenge, as it would enable player movement analysis for clubs and researchers who lack access to comprehensive tracking data. To date, there has been little work addressing the imputation of multi-agent behaviour when observations are both temporally irregular and limited to a single agent at each timestep.

Another significant research direction is the interpretability and explainability of spatiotemporal prediction models. While deep learning approaches have achieved impressive predictive performance, they often function as black boxes, offering limited insight into the underlying decision-making processes or the interaction patterns between agents. Extracting interpretable outputs from spatiotemporal prediction models, such as highlighting influential features or environmental factors, would be highly valuable for domains where model explainability and trust are important, including autonomous driving and disaster response.

Finally, most existing models are predominantly data-driven, relying solely on historical location and trajectory data. There is substantial potential to enhance spatiotemporal prediction by integrating domain knowledge and behavioural constraints, such as physical limits, environmental rules (e.g., speed limits), or tactical positions in sports. Incorporating such information could improve both the accuracy and plausibility of predictions, particularly in structured environments like sports or transportation systems.

Spatiotemporal Performance Evaluation Recent advances in spatiotemporal data collection have enabled a wide range of AI-driven performance models in football, supporting applications in player recruitment, tactical analysis, and player development. By leveraging both event and tracking data, these models provide objective evaluations of player performance, including attacking and defensive actions, space creation, and decision-making processes. Despite significant progress, several important research challenges remain, many of which represent valuable directions for further study in both football analytics and the broader fields of AI and MAS.

A key challenge lies in the objective evaluation of defensive actions, especially those occurring off the ball. Most existing approaches in football focus on on-ball events (Merhej et al., 2021), overlooking the substantial influence of off-ball positioning and spatial coverage. Similarly, in other team-based domains, such as security and patrol, the indirect contributions of agent positioning to adversary decision-making remain underexplored. There remains a need for metrics and models that can robustly quantify the indirect spatial impact and influence of individual agents, particularly in complex, dynamic environments. The advancement of attention-based models could be leveraged to produce these metrics. Developing such models would enable more comprehensive performance evaluation and enhance player selection and team optimisation processes.

Another open research area is how both on-ball and off-ball behaviours can be evaluated through richer, more detailed data. In football, metrics such as xG, xT, and EPV have advanced the objective assessment of attacking actions. However, both in football and more generally, there is limited integration of nuanced action subtypes (e.g., headers, weak-footed shots) and deeper contextual features (e.g., pass elevation, player physical attributes, or environmental conditions). In the wider AI literature, action evaluation in MAS often does not handle such in-depth human data, suggesting that approaches using these rich data could help advance these fields. The development of more comprehensive datasets, such as body pose and eye tracking data, along with context-aware models, would enable more granular and actionable insights for coaches, analysts, and decision-makers in football.

Space creation and occupation represent a further area of opportunity for research. In football, recent models have begun to quantify the value of space and off-ball movement (Fernandez and Bornn, 2018; Llana et al., 2020). However, predicting the emergence of valuable space in future states, especially in highly dynamic, adversarial, or partially observed environments, remains an open problem in both football and the general AI literature. Models capable of forecasting where and when high-value space will occur could provide a powerful tool for proactive tactical planning and player coordination.

IL approaches from other domains, such as video games (Harmer et al., 2018) and navigation (Silver et al., 2008), could also demonstrate the potential for modelling individual player behaviour in football using spatiotemporal data. Applying these methods to football could enable the development of tools that use simulations to compare player decision-making and effectiveness in similar tactical scenarios, offering valuable insights for team selection and performance analysis.

Teamwork The objective evaluation of teamwork remains a significant challenge in both football analytics and the broader field of MAS. Although the importance of effective collaboration and communication within teams is widely acknowledged, objectively quantifying these aspects is difficult due to the complex and often subtle nature of team interactions. In football, recent studies have begun to address this gap by modelling player interactions using network analysis and event data, which has enabled the identification of key partnerships and the optimisation of team selection based on learned teamwork metrics (Beal et al., 2020b; Bransen and Van Haaren, 2020). These approaches have shown that incorporating measures of player interaction can substantially improve the prediction of team performance compared to models that focus solely on individual player skills.

Despite these advances, several open research areas remain. Current models are limited to on-ball actions and direct one-to-one interactions, overlooking the rich context provided by off-ball movement, spatial proximity, and indirect forms of collaboration.

The increasing availability of tracking data presents new opportunities to capture these off-ball dynamics, such as coordinated defensive positioning and space creation, which are critical to team success. In the broader AI literature, agent interaction modelling has primarily focused on task allocation (Chalkiadakis and Boutilier, 2012) and coalition formation to meet agent objectives (Bistaffa et al., 2017), placing less emphasis on the spatial and temporal coordination between agents to maximise spatial control in real-world teams, especially in adversarial environments like football and security. Developing models that can optimise spatial coordination between agents in dynamic, adversarial contexts using agent location data remains an essential and open challenge.

Advancing teamwork modelling will require the integration of richer spatiotemporal data and the development of metrics that capture both direct and indirect forms of collaboration. These advancements could be used to inform coaching, player recruitment, tactical planning and post-match analysis. It would provide information on which players coordinate well together and predict strong future player relationships. This would ultimately support better decision-making in both sports and other domains where teamwork is essential.

Team Planning The application of AI and ML to team planning in football has significantly advanced tactical analysis, TF, and injury prediction, providing coaches and analysts with new tools to optimise both short-term match outcomes and long-term team performance. By leveraging event and tracking data, recent research has enabled more objective, data-driven approaches to pre-game tactical planning, in-game decision-making, and the modelling of team styles and structures. Despite these advances, several important research challenges remain, some of which are shared with the broader AI and MAS literature.

One notable area of progress is the use of AI models for pre-game tactical planning. Bayesian and ML approaches have been employed to optimise team formations and playing styles based on historical data and opponent analysis (Beal et al., 2020a, 2021), demonstrating practical value in improving win probabilities and league performance. However, most existing models primarily rely on event data, and integrating richer spatiotemporal information, particularly tracking data, remains an open challenge for achieving deeper tactical insights.

In-game tactical optimisation also presents several open research questions. While game-theoretic (Hirotsu et al., 2009) and dynamic programming (Hirotsu and Wright, 2002) approaches have been used to inform tactical changes and substitution timing, existing approaches typically do not integrate in-depth spatiotemporal features to address emerging issues on the pitch in real-time. Research in MAS, such as dynamic coalition formation, offers promising directions for developing more in-depth in-game decision-making frameworks. The potential for live data gathering and feature extraction to support these models is significant, but practical implementation in football is a challenge.

Injury prediction and prevention are another critical aspect of team planning, with substantial implications for both team performance and financial outcomes. Although ML models have shown promise in predicting injury risk using features such as workload, injury history, and physical attributes (Kucera et al., 2005; Bowen et al., 2020), their integration into team selection and workload management processes is still underexplored. Developing explainable and reliable injury prediction models, along with frameworks that balance immediate performance with long-term player availability, remains an open research area. This research is relevant across domains where agent failure or unavailability is a concern, such as other team sports and emergency response.

Finally, the broader problem of TF and planning in MAS has been extensively studied in domains such as disaster response (Capezzuto et al., 2021) and wildlife patrol (Wang et al., 2019b). Techniques such as coalition formation under spatiotemporal constraints (Ramchurn et al., 2010) and robust team selection (Okimoto et al., 2015) have been developed to address challenges related to agent skills, dynamic task environments and agent failures. However, proactive management of agent workload and failure risk through TF to improve long-term team performance and reduce unavailability is underexplored. Additionally, explicitly connecting and optimising pre-event team selection with the dynamic replacement of agents to handle fatigue and adapt to evolving opponent strategies in adversarial environments remains an open challenge. Football is a suitable testbed for advancing these models and validating their effectiveness due to the availability of in-depth spatiotemporal, physical and opponent data.

Our Research Objectives From these open research areas that have been identified from our literature review, we identified a set of research objectives to focus on to contribute to the state-of-the-art. Extending from our research questions in Section 1.1, these research objectives are:

- **1. Spatiotemporal Agent Imputation in Limited Observability:** While many spatiotemporal models predict agent trajectories in domains like crowd movement (Kim et al., 2013; Alahi et al., 2016), autonomous vehicles (Xie et al., 2021), and ship tracking (Siegert et al., 2016), some address imputing missing agent locations from intermittent and incomplete data (Omidshafiei et al., 2022; Qi et al., 2020). However, the challenge of imputing missing agent locations under limited and non-uniform observations, such as when only a single agent is observed at each timestep, remains underexplored.

  Our first objective is to develop a model that can accurately estimate agent locations and movements under these challenging constraints, with a particular focus on generating player tracking data from sparse football event data. This will make advanced off-ball and physical player analysis, such as those completed in Chapters 4, 5 and 6, more accessible to clubs and researchers with restricted data resources. The model's effectiveness will be evaluated against baseline spatiotemporal and ML models. This objective will be addressed in Chapter 3.

- **2. Learning and Optimising Spatial Teamwork in Multi-Agent Teams:** A key challenge in MAS is objectively quantifying and optimising teamwork, especially the complex spatial interactions that drive effective collaboration in spatiotemporal domains. While previous work in domains like social ride-sharing (Bistaffa et al., 2017) and disaster response (Ramchurn et al., 2010; Chalkiadakis and Boutilier, 2012; Capezzuto et al., 2020) has explored TF, the quality and dynamics of spatial agent interactions are rarely evaluated. In football, most frameworks focus on one-to-one, on-ball interactions (Beal et al., 2020b), overlooking the broader spectrum of off-ball player collaboration that is essential to team success.

  Our objective is to harness player location data in football to construct a model that learns spatial interactions between team members, and optimises spatial decision-making to maximise the team's probability of winning the game. By capturing the impact of spatial proximity, coordination, and collective movement on team performance, this approach also aims to provide actionable insights for coaching and training and to advance the understanding of teamwork as a central component of football tactics, space creation, and team modelling. The methodology for this objective will be presented in Chapter 4.

- **3. Evaluating Indirect Agent Contributions in Multi-Agent Defence:** Assessing individual contributions in team-based defence is particularly challenging when those contributions are indirect and not directly observable in outcomes. While previous research in security and sports-based defence domains has primarily focused on quantifying direct actions (Shieh et al., 2012; Decroos et al., 2020; Merhej et al., 2021), the influence of indirect defensive behaviours, such as off-ball positioning in football, remains underexplored. Recent advances in spatiotemporal data and deep learning, particularly graph-based models, provide new opportunities to capture these complex interactions.

  This objective seeks to develop a GAT-based framework to generate interpretable metrics for off-ball defender contributions, identifying and quantifying the influence of individual defenders on attacking outcomes. By introducing novel attention-based metrics and validating the approach with real-world event and tracking data, this work aims to advance the measurement of indirect agent contributions in multi-agent defence, offering practical tools for more granular scouting, tactical analysis, and player evaluation in sports analytics. This objective will be explored in Chapter 5.

- **4. Optimising Short- and Long-Term Team Selection:** Existing TF models have been applied across various domains (Ramchurn et al., 2010; Gaston and DesJardins, 2005), often focusing on agent skills, spontaneous task emergence, or robustness to agent failure (Gunn and Anderson, 2013; Okimoto et al., 2015; Demirović et al., 2018). However, these approaches do not account for the influence of workload, fatigue, and long-term agent availability. These factors are especially critical in football, where player injuries and fatigue can significantly impact team performance and financial outcomes. Current football team selection models primarily address tactics and player performance (Beal et al., 2020a, 2021), but do not integrate real-time monitoring of player fatigue or the long-term risk of injury into team selection and substitution strategies.

  Our objective is to develop a sequential TF framework that combines pre-game and in-game models, leveraging real-world football data to monitor player workload, fatigue, and opponent strategies. This framework will enable proactive and flexible team selection and substitution decisions, optimally balancing immediate match performance with long-term player welfare and team objectives. By utilising estimated tracking data (from Chapter 3) to compute player workload features, the model will be available to clubs with limited data resources. The approach will be validated against real-world team selection strategies, with the aim of reducing injuries, improving season-long performance, and providing a generalisable solution for other dynamic, team-based domains. This objective will be addressed in Chapter 6.

The first of these objectives is expanded on in the next chapter, which will focus on "Multi-Agent Spatial Imputation from Event Data to Infer Player Movement".

Chapter 3

Multi-Agent Spatial Imputation from Event Data to Infer Player Movement

Existing MAS problems that predict agent locations typically use uniform timesteps with observations for all agents. In this chapter, we analyse the problem of agent location imputation, specifically posed in environments with non-uniform timesteps and limited agent observability ($\sim$95% missing values). Our novel approach uses long short-term memory (LSTM) and graph neural network (GNN) components to learn temporal and inter-agent patterns to predict the location of all agents at every timestep. We apply this to football by imputing the location of all players in a game from sparse event data (e.g., shots and passes). Our model estimates player locations to within $\sim$6.9m; a $\sim$62% reduction in error from the best performing baseline. This approach facilitates downstream analysis tasks such as player physical metrics, player coverage, and team pitch control. Existing solutions to these tasks often require tracking data, which is expensive to obtain and only available to elite clubs. By imputing player locations from easily obtainable event data, we increase the accessibility of downstream football analysis.

3.1 Introduction

Predicting the behaviour of such agents is an important task in many areas, such as path prediction for autonomous vehicles (Sriram et al., 2020), predicting locations of civilians in disaster response scenarios using phone data or drone footage (Ramchurn et al., 2016; Wang et al., 2021), and predicting player movement in sports analytics (Omidshafiei et al., 2022). However, situations may arise where the observability of the environment is limited. For example, an autonomous vehicle may have an obstructed view and lose sight of nearby pedestrians, or drone footage may have limited observability of an area when searching for injured civilians. In these situations, reasonable estimations of agent locations need to be made in order to improve the efficacy of response systems. In this work, we build a multi-agent time-series imputation model, named *Agent Imputer*, which learns typical spatiotemporal interactions between agents to estimate the location of agents when system observability is limited. We apply this to football by predicting the location of all players on a pitch based solely on observing the location and action (e.g., a pass or shot) of the ball-carrying agent and in other words, estimating tracking data using only on-ball event data. Figure 3.1 describes the workflow of our model and its application to football.

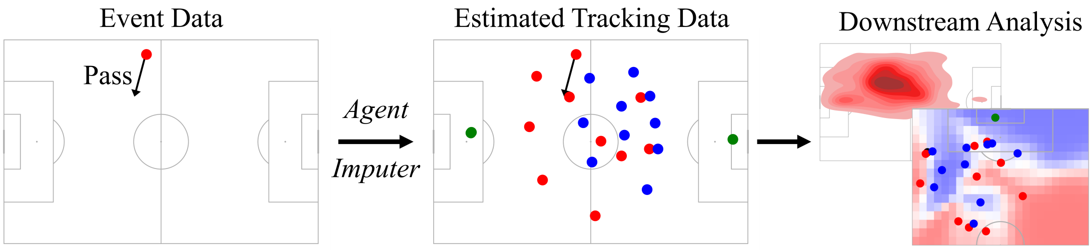

FIGURE 3.1: Comparison between event and tracking data for a single timeframe, and how our model estimates tracking data to facilitate downstream analysis tools.

Football represents an interesting domain for modelling teams as dynamic MAS, with each player having individual roles and behaviours. Whilst all players exhibit different behaviours, the system is centred around the ball. Therefore, in this work, we use data regarding the ball and the ball-carrier to estimate all player locations when an on-ball event occurs. We build a model that includes LSTM and GNN components to learn the temporal and inter-agent relationships between players in a football game, enabling us to make generalised predictions about unobserved player locations during a game.

As discussed in Chapter 2, analytical data in football has many uses. These include modelling team tactics, calculating physical metrics (e.g., distance covered), and measuring pitch dominance of teams (Spearman, 2018; Link et al., 2016; Fernández et al., 2019). It can also be used by football clubs to perform player evaluation, scouting, and opposition analysis. In this work, we examine how sparse football event data of on-the-ball actions can be used to impute knowledge of player positions. This facilitates the desired downstream analysis tasks without requiring expensive tracking data.

Thus, this chapter makes the following novel contributions:

1. We propose a novel *Agent Imputer* model that predicts agent locations when constrained to limited system observability. This model learns both temporal and inter-agent interactions using LSTM and GNN components to make estimations of agent positioning applied to football.

2. We use event and tracking data from real-world football games to train and test our model, and find that it predicts agent positioning with a mean Euclidean distance error of ∼6.9m, outperforming numerous baselines by ∼62%.

3. We apply our model predictions to real-world downstream applications in the football domain, such as calculating estimations of player physical metrics, pitch dominance, and player coverage heatmaps.

To our knowledge, no previous work has studied the problem we approach in this work: imputing agent behaviour within a multi-agent system when information occurs at non-uniform timesteps and for a single agent at a time. This challenge is further explained in Section 2.3 and illustrated in Figure 2.5.

The rest of this chapter is structured as follows. Section 3.2 formally defines the imputation problem. Section 3.3 then introduces our novel *Agent Imputer* model. Section 3.4 contains our empirical evaluation, and Section 3.5 presents downstream applications. In Section 3.6, we discuss model outcomes and future work. Finally, Section 3.7 summarises the work.

3.2 Problem Formulation

Our multi-agent system is a set of N agents, $C = \{c_1, \ldots, c_N\}$. In our target domain of football, $N = 22$ (two teams of 11 on-field players). The time-series is a sequence of T events $E = \{e_1, \ldots, e_T\}$. Note this time-series is non-uniform, as each element $e_t \in E$ is recorded when an agent performs an on-ball action (e.g., a pass, dribble or shot), leading to varying gaps between each timestep.

For each event $e_t \in E$, we have a set of observations $\Phi_t^C = \{\phi_t^1, \ldots, \phi_t^N\}$, where ϕ_t^n is the observation of agent n at timestep t. This gives a complete set of observations over time of $\mathbf{\Phi}^C = \{\Phi_1^C, \ldots, \Phi_T^C\}$. However, in our configuration, only one value in each $\Phi_t^C \in \mathbf{\Phi}^C$ is known, i.e., there are $N - 1$ missing values for each Φ_t^C, and $T(N - 1)$ missing values in total. This occurs as we only make an observation of one agent at each timestep — in our target domain of football, this is the position of the on-the-ball agent. As the observed agent changes over time, we construct a one-hot encoded mask $\mathbf{M}$, where $\mathbf{M}_t^n = 1$ if agent n is observed at timestep t, and is 0 otherwise. This $(T \times N)$ binary matrix fully captures the information regarding known and unknown observations across the time-series problem.

In this configuration, the goal of an imputation model is to predict values for the unknown observations. Formally, for each $e_t \in E$, the model makes a prediction $\hat{\phi}_t^n$ for every $n \in [1, \ldots, N]$. This leads to a complete set of predicted observations $\hat{\mathbf{\Phi}}^C = \{\hat{\Phi}_1^C, \ldots, \hat{\Phi}_T^C\}$, where $\hat{\Phi}_t^C = \{\hat{\phi}_t^1, \ldots, \hat{\phi}_t^N\}$. Note that in this setup, the model is also predicting observations for cases it has already observed (i.e., players on the ball); however, when we apply downstream analysis (Section 3.5), we instead use the actual locations for already observed agents.

Relating the above to our target domain of football, the predicted observations $\hat{\mathbf{\Phi}}^C$ are the estimated locations of all players (both on-ball and off-ball) for every event $e_t \in E$. The known observations are derived from the locations where the events occur. So, for event $e_t \in E$ with an on-ball agent $c_n \in C$, the assigned observation ϕ_t^n is $e_t^{x,y}$, where $e_t^{x,y}$ is the x, y position at which e_t occurred.

To summarise, the set of known observations $\mathbf{\Phi}^C$ (containing missing data) is a $(T \times N \times 2)$ tensor, and the set of imputed observations $\hat{\mathbf{\Phi}}^C$ is also a $(T \times N \times 2)$ tensor. Note, for this thesis, each observation is two-dimensional as these observations refer to the x, y positions of players on the pitch — in other domains, these observations could be a different size. In the next section, we explain our model to solve this imputation problem.

3.3 Agent Imputer Model

The problem we describe in Section 3.2 is highly complex as over 95% of the values in $\mathbf{\Phi}^C$ are missing. We aim to extract useful information on agents' spatial and temporal behavioural patterns from event data, and apply it to a model which can consider how agents move over time and in relation to other agents. In this section, we outline the feature engineering (Section 3.3.1), model architecture (Section 3.3.2), and training process (Section 3.3.3) used in our approach. For implementation details, see Appendix A.1.

3.3.1 Feature Engineering

The features available at a given timestep are strictly derived from on-the-ball event data, consisting of the event location, player on the ball, time at which the event occurred, and other information such as the current scoreline. This provides little to no information on the spatiotemporal context of an off-ball agent, such as when their location was last observed (i.e, when they were last on the ball). Therefore, we perform feature engineering on the event data to generate our own feature set, which captures a more comprehensive view of each agent and the general movement of play. We create these features for each agent, so they are a combination of agent-specific and global features. For a timestep t with event $e_t \in E$, and observed agent $c_n \in C$, we compute the following features:

Agent-Specific Features

- `prevAgentTime,prevAgentX,prevAgentY`: time since the agent was last observed, and their location at that time. Formally, this uses the most recent previous timestep where the agent was on the ball: t' s.t. $\mathbf{M}^n_{t'} = 1$ and $t' \leq t$, which can be found by iterating backwards in time from t. Note $t' = t$ if the agent is currently on the ball. Given t', (`prevAgentX`, `prevAgentY`) = $e^{x,y}_{t'}$ and `prevAgentTime` = $t - t'$. If the agent is yet to be observed, we impute these features with the values at the first timestep where they are observed.

- `nextAgentTime,nextAgentX,nextAgentY`: time until the agent is next observed, and their location at that time. Similarly to the previous time and location features, this uses the soonest future timestep (including the current timestep) where the agent is on the ball: t' s.t. $\mathbf{M}^n_{t'} = 1$ and $t' \geq t$, which can be found by iterating forwards in time from t. Given t', (`nextAgentX`, `nextAgentY`) = $e^{x,y}_{t'}$ and `nextAgentTime` = $t' - t$. If the agent has no future observations, we impute these features using the last time they were observed.

- `avAgentX,avAgentY`: mean location of the agent across the entire game, i.e., the mean position of all events where $\mathbf{M}^n = 1$.

- `agentRole`: the agent's role in the team (e.g., central defender or goalkeeper). Different data providers label roles differently — for our data source, there are 16 possible roles, see Appendix A.2.2.

- `agentSide`: binary indicator of whether an agent is on the same team as the current ball-carrying agent.

- `agentObserved`: binary indicator of whether the agent is the one performing the current on-ball action.

- `goalDiff`: the difference in score between the teams (number of goals for this agent's team - number of goals for the other team).

Global Features

- `eventX,eventY`: location of the current event (i.e. $e^{x,y}_t$).

- `eventType`: the type of the current event (e.g., a pass or shot). For more details, see Appendix A.2.1.

These 15 features, for each agent at each timestep, represent a transformation of the original event data. Note that agent positions are relative to the agent's own goal-line, irrespective of the direction of play (as opposed to absolute position on the pitch). This feature set is designed to cover a variety of contextual factors relevant to position imputation: (i) short-term dynamics through previous/next observation times and locations, providing information on current movement trends; (ii) longer-term agent information using average positions and role labels, which reflect typical spatial occupation of agents; (iii) team context and game state using team side and goal difference, which influence expected behaviours and agent interactions; and (iv) global player context through the current event location and type, which indicate where and how play is developing. Together, these features provide an machine learning (ML) ready summary of agent- and global-level context from on-ball events alone. In the following subsection, we detail our *Agent Imputer* model.

3.3.2 Model Architecture

Our *Agent Imputer* model estimates the set of agent locations $\hat{\Phi}^C$ using the engineered features from the set of events E. The goal of this model is to capture agent movement over time, the relationship between on-ball events and off-ball agent locations, and how agents interact as a team within the multi-agent system. In football, this is often defined as team structure. We build an architecture that learns across time and agents within the multi-agent system. Below, we explain this model using a step-by-step process, with corresponding elements of the architecture marked in Figure 3.2.

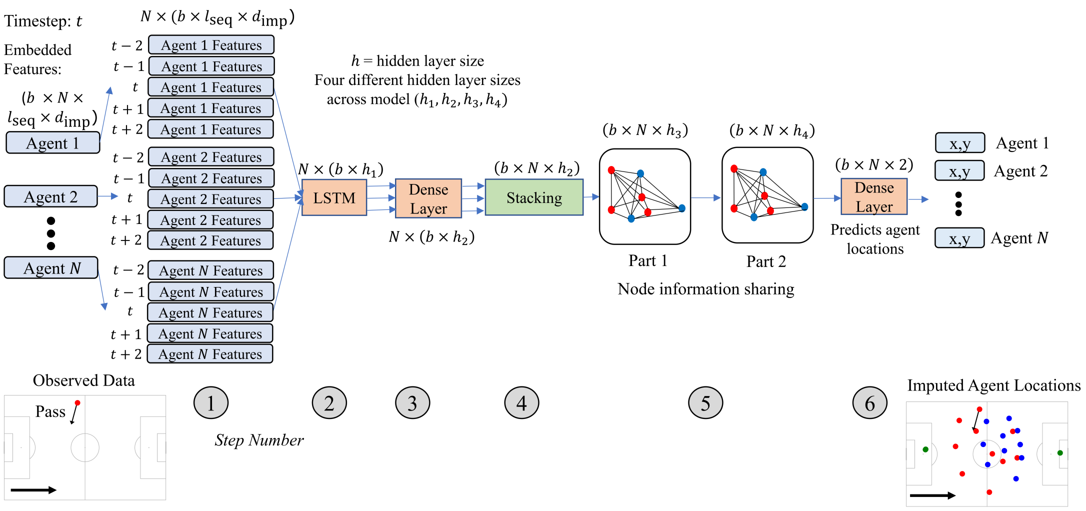

FIGURE 3.2: *Agent Imputer* model architecture. All agent location estimations for an event are made using agent feature information described in Section 3.3.1, with batch size b, number of agents N, sequence length l_{seq}, and number of features d_{imp}.

Step 1: Input Formatting To capture the ball and agent locations across time, we input a window of event data. We extract the feature set defined in Section 3.3.1 for a sequence of $l_{seq} = 5$ events. This sequence is centred around a particular timestep t, i.e., we compute the features for timesteps $\{t-2, \dots, t+2\}$.[1] Each categorical feature (`agentRole`, `agentSide`, `agentObserved`, `goalDiff`, `eventType`) is converted to an embedding (see Appendix A.1.2 for details). Our imputation model takes all the agent feature sets as a single simultaneous input, such that the model learns the spatial structure of

[1]If timesteps before or after the current timestep don't exist due to being at the start or end of a game, the data for the current timestep is used instead.

the whole system (as opposed to processing each agent independently). Therefore, the input data has shape $(b \times N \times l_{\text{seq}} \times d_{\text{imp}})$, where b is the batch size, N is the number of agents, l_{seq} is the sequence length, and d_{imp} is the post-embedding dimensionality of our feature set including both the numerical features and embedded categorical features. In this chapter, $d_{\text{imp}} = 24$ and $b = 128$.

Step 2: LSTM Component Each agent is independently passed into a shared bidirectional LSTM component, i.e., the input data is split into N segments of size $(b \times l_{\text{seq}} \times d_{\text{imp}})$. The LSTM is then able to learn the temporal relationship between the engineered features and agent location. Importantly, the LSTM is shared across all agents. This overcomes issues with the sparsity of agent observations, as the LSTM is able to learn agent positioning through common movement patterns for agents with similar roles. Due to the irregular time intervals between timesteps, it is important to augment the LSTM architecture to deal with non-uniformity. Therefore, we use a Time-Aware LSTM (Baytas et al., 2017) which adjusts cell memory to alter the discount rate of previous or future actions in the sequence based on the difference in time from the current event. In this work, an LSTM with a single hidden layer of size $h_1 = 100$ is used.

Step 3: Dense Layer with ReLU Activation Outputs from the LSTM model for each agent at the current timestep t (i.e., the middle of the input sequence) are extracted. This output is passed through a dense layer with a rectified linear unit (ReLU) activation function to give an output of size $h_2 = 50$, resulting in a latent representation of the time-aggregated features for each agent.

Step 4: Stacked Agent Embeddings The temporally aggregated latent representations are stacked together, giving an output tensor of size $(b \times N \times h_2)$. This data is now a suitable input for a GNN, with each tensor row representing agent information that becomes node features for the GNN.

Step 5: GNN Component The LSTM component of the network models the temporal aspect of agent behaviour. We now aim to model the inter-agent relationships within the multi-agent system and how these impact the behaviour of each agent. We construct a fully connected graph with the temporally aggregated agent representations as node features. The GNN uses this graph structure to allow information sharing across all agents. The GNN architecture consists of two message-passing layers with feature sizes of $h_3 = 64$ and $h_4 = 32$, using the SAGEConv operator (Hamilton et al., 2017). This updates each node's feature representation by using a mean aggregation scheme on its neighbours' features. The aggregated neighbourhood features are then transformed and combined with the node's own transformed features to create an updated node representation. These transformations are executed using a learned linear mapping, with non-linearity added by Step 6 of the model. This process allows the GNN to learn information about agent neighbourhoods and interactions. The use of two layers enables information to propagate beyond immediate neighbours, which allows us to capture more complex, higher-order agent interactions within the MAS.

Step 6: Dense Layer with ReLU Activation The final output of the GNN component is of size $(b \times N \times h_4)$, i.e., latent representations of size 32 for each agent, which captures both temporal and inter-agent interaction information. These final representations are then passed through a single dense layer with a ReLU activation function to make the final position (x, y) predictions for each agent. Therefore, the final model output is of size $(b \times N \times 2)$.

3.3.3 Model Training

We use tracking data as target variables during training, and as ground truth labels to test the predictive accuracy of our *Agent Imputer* model. We use the mean Euclidean distance between the predicted positions and actual positions as the loss function to train the model. The model was trained for 150 epochs with a batch size of 128. The AdamW (Loshchilov and Hutter, 2019) optimiser was used with an initial learning rate of 0.002. These hyperparameters and the exact model architecture (e.g., hidden layer sizes) were found through trial and error, i.e., no formal hyperparameter tuning was carried out. In the next section, we empirically evaluate the *Agent Imputer* model.

3.4 Empirical Evaluation

In this section, we describe the data used to train and evaluate our *Agent Imputer* model (Section 3.4.1), introduce the baselines (Section 3.4.2), and give our results (Section 3.4.3).

3.4.1 Datasets

We train and evaluate our model using 34 games of event and tracking football data collected from K League 1[2] and supplied to us by Bepro Group Ltd. These are gold-standard industry datasets that allow us to rigorously evaluate our approach. Each game provides event sequences, which we use as model input, and tracking data, which we use as training targets and for evaluation. In total, there are ∼64,000 events and ∼1.4 million tracking locations. We use a 31/3 (∼91.2%/8.8%) train/test split and five-fold cross-validation (CV) to evaluate our model. For each fold, three games are randomly sampled without replacement to form the test set, and the remaining games are used

[2]The top men's professional football division in South Korea.

for training. This procedure is repeated five times, so each fold uses a different subset of three games as the test set, and a total of 15 games are tested across the five folds. All geometric data is scaled to a standard football field size of 105x68 metres. Further dataset details are given in Appendix A.2.

3.4.2 Baselines

To evaluate whether our model improves performance compared to other imputation methods, we consider the following baselines:

Naïve imputation baselines:

- **Baseline 1:** Predicts the agent location as the average on-ball location of the agent during the match.
- **Baseline 2:** Predicts the centroid between the last observed and the next observed location of the agent.
- **Baseline 3:** Predicts a time-scaled position along the straight-line trajectory between the last and next observed location of the agent.

Machine learning model baselines (with the same input feature set):

- **XGBoost regression model.**
- **Time-Aware LSTM model:** The same model as in Figure 3.2, but without step 5.
- **GNN model:** Corresponds to steps 3, 4, 5, and 6 of Figure 3.2, using the middle of the original input sequences as input features.

These baselines allow us to disentangle and evaluate the contributions of the constituent parts of our *Agent Imputer* model towards performance.

3.4.3 Results

Five experiments are used to evaluate the performance of our *Agent Imputer* model and relevant baselines. We evaluate player position prediction, performance over time, performance over different positions, how performance varies with observation, and performance over various pitch locations.

TABLE 3.1: Predictive performance of models averaged across five folds, along with 95% confidence intervals. We report errors separately in the X and Y directions, as well as the overall 2D Euclidean error, which we label as XY error. All errors are in metres. Bold results indicate the best performance (lowest error).

	Train			Test		
Models	X Error	Y Error	XY Error	X Error	Y Error	XY Error
Baseline 1	14.20 ± 0.02	9.69 ± 0.02	18.81 ± 0.05	14.22 ± 0.24	9.70 ± 0.14	18.82 ± 0.22
Baseline 2	14.10 ± 0.02	9.79 ± 0.02	18.91 ± 0.05	14.04 ± 0.48	9.78 ± 0.17	18.80 ± 0.50
Baseline 3	13.37 ± 0.05	9.63 ± 0.01	18.15 ± 0.05	13.27 ± 0.44	9.57 ± 0.14	18.01 ± 0.46
XGBoost	5.88 ± 0.01	5.14 ± 0.02	8.67 ± 0.02	6.42 ± 0.35	5.50 ± 0.09	9.26 ± 0.26
Time-Aware LSTM	4.46 ± 0.02	4.26 ± 0.02	6.89 ± 0.02	4.48 ± 0.08	4.49 ± 0.10	7.09 ± 0.06
GNN	5.28 ± 0.05	5.08 ± 0.10	8.18 ± 0.09	5.42 ± 0.35	5.13 ± 0.14	8.32 ± 0.25
Agent Imputer	**4.06 ± 0.02**	**4.12 ± 0.02**	**6.47 ± 0.02**	**4.29 ± 0.09**	**4.41 ± 0.11**	**6.88 ± 0.10**

Position Prediction We evaluate the predictive performance of the model and baselines in Table 3.1. Note that the X direction is along the length of the pitch (105m) and the Y direction is along the width of the pitch (68m), and that we do not normalise the error in these different directions by pitch size.

We find that the *Agent Imputer* model predicts agent location with the highest accuracy (lowest distance error). We also find that the Time-Aware LSTM model outperforms the GNN model, suggesting the essential part of our *Agent Imputer* is the LSTM component. However, we still see an improvement in performance by including the GNN component in the *Agent Imputer* model, demonstrating the value of modelling agent interactions. We note a 61.8% decrease in error using the *Agent Imputer* compared to the best performing naïve baseline. We also find the *Agent Imputer* model has roughly equal error in both X and Y directions.

Interestingly, for the *Agent Imputer*, we find that the sum of the squares of the mean X and Y errors is smaller than the mean Euclidean error. This is because the Euclidean error is computed for each sample before averaging. Because the square root is a nonlinear operation, the sum of squared mean X and Y errors does not necessarily equal the mean Euclidean error. Figure 3.3 shows the distribution of per-sample X, Y and Euclidean errors across our testing dataset.

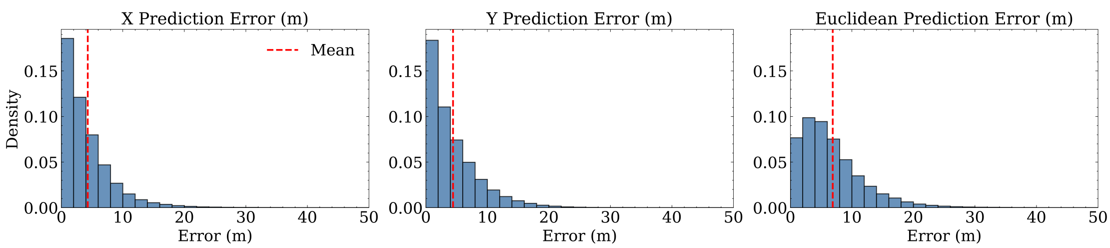

FIGURE 3.3: Distribution of the X, Y and Euclidean predictive errors for the *Agent Imputer* model across the entire testing dataset.

Predictability over Time We evaluate the predictive accuracy of the model across different periods of a match. In football, the game has two 45-minute halves, so we evaluate how model accuracy changes during these periods. We compute this using a rolling average of model error; see Figure 3.4. Additional time at the end of halves is accounted for — first half added time is merged with early second half predictions, and values beyond 90 minutes facilitate second half added time.

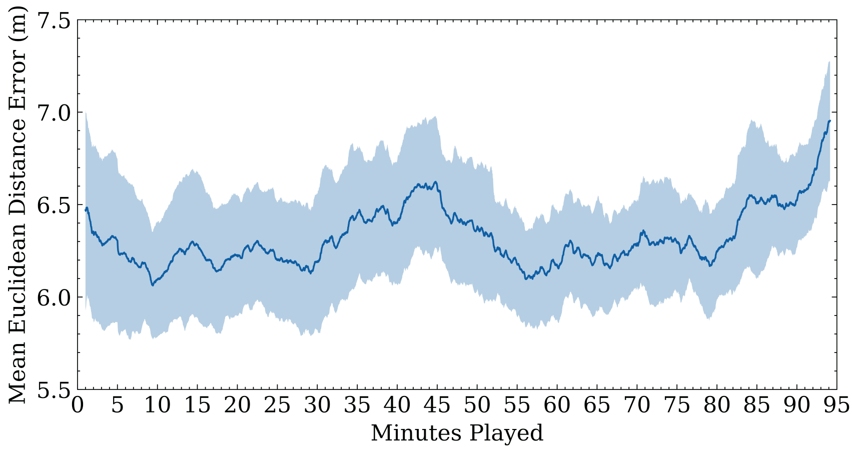

FIGURE 3.4: Predictive performance of the *Agent Imputer* model across time using a rolling average of Euclidean error (5-minute windows with a 1-second step size) on the combined training and testing data (all 34 games). Shading shows the standard deviation across means for each game.

From these results, we suggest that the most unpredictable moments of a game are at the end of both halves. Due to the error margin, we perform t-tests to investigate the hypothesis that the mean of the underlying sample distribution in the middle of a half is lower than the underlying distribution at the end of a half. When doing this for the first half (all events between 20-25 minutes compared to all events between 42.5-47.5 minutes), this is found to be significant ($p < 0.01$). The difference in the second half (comparing all events between 65-70 minutes with all events between 87.5-92.5 minutes) is also found to be significant ($p < 0.01$).

This supports our theory that the game is more unpredictable at the end of halves, which could be a result of players getting tired and teams becoming less structured during this time. We also perform t-tests to assess whether the mean error differs between the first and second halves of a match. Comparing all events in the first half with those in the second half, we find a statistically significant increase in error in the second half ($p < 0.05$). This suggests that play is generally more unpredictable later in matches, which is consistent with the trend observed in Figure 3.4, where error increases towards the end of the second half.

A plausible explanation for these results may also be that teams are more likely to take risks when chasing a result or defending a lead, leading to less structured and more variable player movement patterns. This type of temporal analysis may therefore help identify match phases associated with changes in team structure and could be extended to detect unusual movement behaviour that coincides with shifts in performance during a match.

Predictability over Roles We evaluate the performance of our *Agent Imputer* model across agent roles in Figure 3.5. This gives insight into how an agent's goals and responsibilities within a team affect the predictability of their behaviour. For simplicity, in this experiment and downstream applications (Section 3.5), we group roles into wide and central positions, reducing the total number of positions from 16 to 7 (see Appendix A.2.2). Goalkeepers are expectedly the most predictable role, as their range of movement is usually limited to their own box. We also find that defenders are more predictable than attackers, highlighting that defensive agents behave in a more structured way than attacking agents.

FIGURE 3.5: Mean distance error of player location estimations for each agent role using the *Agent Imputer* model. Error bars represent the 95% confidence intervals calculated across all predictions for players in each role.

Interesting findings can also be drawn from the X and Y distance errors for each player role. A general trend is that the model finds it harder to predict the X location of wide outfield players in comparison to central outfield players (an average of 5.01m vs 4.15m), which is expected as wide players cover a lot of ground along the wings of the pitch. However, the model generally predicts the Y location of wide players better than central players (an average of 4.62m vs 4.99m), which is also expected as wider players usually stick to a single side of the pitch. These comparisons help identify where the model could be improved, and could be used to help convey uncertainty in downstream analysis.

Predictive Performance over Observation Offset In Figure 3.6, we evaluate the predictive performance of our proposed models with respect to "time since last observed" — the amount of time that has passed since a particular agent was last involved in an event. We also do the same for "time until next observed" — the amount of time until the agent is next involved in an event. We find that the performance of the XGBoost and GNN models rapidly decreases as time increases. In comparison, the *Agent Imputer* and Time-Aware LSTM models show a slower decay in performance, as well as a lower plateau. This demonstrates the necessity of modelling temporal aspects of predictions. However, the time-aware models still decrease in performance as time increases, supporting the intuition that observation data will become less reliable as more time passes since the observation. Similarly to the error by player role (Section 3.4.3), this may be useful in measuring the uncertainty of model predictions, i.e., there is greater certainty in predictions for agents that have been on the ball more recently.

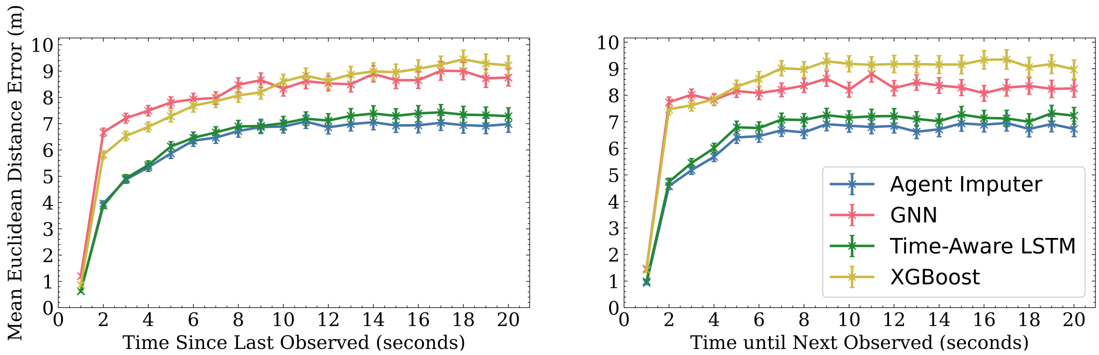

FIGURE 3.6: Model performance compared to the last (left) and next (right) time the agent was observed. We use one-second intervals with a one-second window, e.g., at 2s, we evaluate for all events between 1 and 2 seconds. Error bars represent the 95% confidence interval calculated across all predictions within each one-second interval.

On average, agents are observed performing an on-ball event every 60 seconds in a game. While the figure only shows up to 20 seconds since or until the next observation, the error curves plateau after around 10 seconds, suggesting that prediction uncertainty stabilises mid-range. This range is likely where the structural, inter-agent component of the model becomes most important for prediction due to the long intervals between observations.

Predictability over Ball Location In Figure 3.7, we evaluate the *Agent Imputer* model error conditioned on the ball location. For each zone, we compute the mean error over all events in which the ball resides within that zone. The results indicate that player positioning is least predictable when the ball is in dangerous attacking areas or deep defensive zones. This may reflect high-intensity pressures that displace players from their typical positions, opportunistic attacking runs intended to escape defenders, and defensive adjustments (e.g., man-marking) that temporarily disrupt the overall team structure. Increased error may also be attributable to chaotic phases in which teams are scrambling to create or prevent chances when the ball is not under secure possession in dangerous areas. Conversely, positioning is most predictable when the ball is in central midfield zones, which suggests that build-up play in football is the most structured period of play. We also note that errors across the left and right wings of the pitch (corresponding to the top and bottom of the y-axis in Figure 3.7) show minor differences; however, these differences largely fall within the 95% confidence intervals. This indicates that there are no statistically significant asymmetries in player predictability when play is on different wings of the pitch, which aligns with the expectation that teams generally mirror tactical structures on the left and right sides.

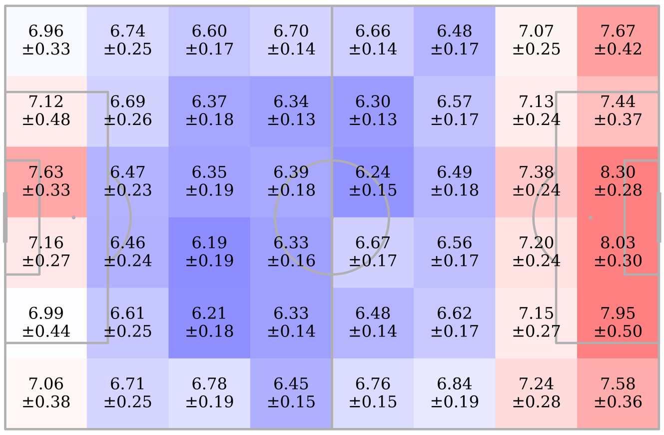

FIGURE 3.7: Mean distance error of agent location predictions for each zone of the pitch that the ball occupies using the *Agent Imputer* model. Error is averaged across all events where the ball was located within each zone, and ± values represent the 95% confidence interval. Teams attack from left to right.

In the next section, we demonstrate how the *Agent Imputer* model outputs can be utilised for downstream analysis.

3.5 Model Applications

In this section, we demonstrate downstream analysis that can be performed using outputs from our *Agent Imputer* model. In football, these types of analyses are often implemented by elite-level clubs using tracking data. However, we demonstrate that lower league clubs with limited resources could use our model to perform processes similar to those of elite-level clubs without requiring hard-to-obtain tracking data. We begin by visualising example predicted player trajectories using the *Agent Imputer* model (Section 3.5.1), and then apply our model to analysing player physical metrics (Section 3.5.2), pitch control (Section 3.5.3), and positional heatmaps (Section 3.5.4).

3.5.1 Team Trajectory Predictions

We present example visualisations of *Agent Imputer* predictions for agent locations over a sequence of five events, alongside ground-truth positions, in Figure 3.8. To aid interpretability, we focus on the attacking team and use annotated player identifiers to directly compare predicted and observed trajectories. The mean Euclidean errors for the examples in Figure 3.8 (4.61m and 6.28m) are slightly below the overall mean test set error of 6.88m. These sequences were selected for visual clarity, where players are clearly separated and not clustered, rather than for optimal predictive performance.

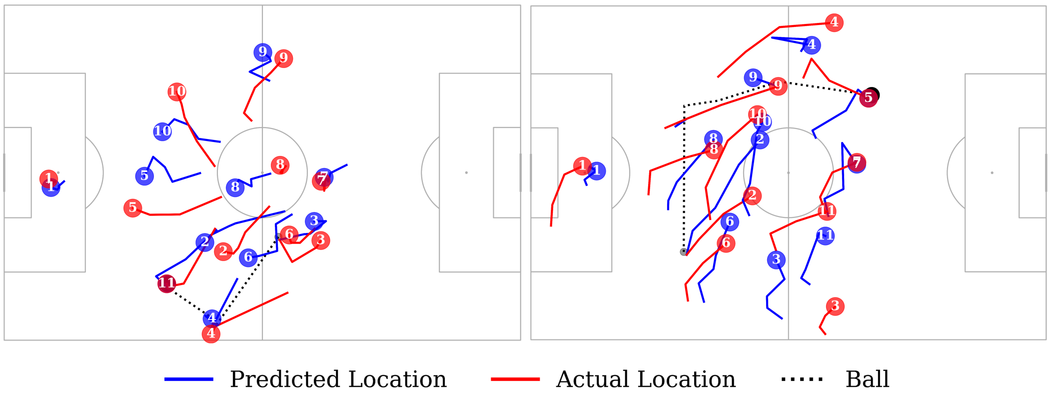

FIGURE 3.8: Examples of predicted (using estimated locations from the *Agent Imputer* model) and actual trajectories of players for an attacking team across 5 events. Dots show the end of the trajectory for each player. The team attacks from left to right. The mean Euclidean errors between the predicted and actual player positions in these examples are 4.61m (left) and 6.28m (right).

These examples illustrate how *Agent Imputer* can be used to estimate team-level spatial context from on-ball event data alone. In particular, the model seems to generate similar directional movement and team structures as well as role-consistent player behaviour. This suggests that the model could also be useful for simulating play sequences, with a potential use case for testing new tactics in simulated environments.

3.5.2 Player Physical Metrics

We show how predicted player locations from our *Agent Imputer* model can be used to estimate player physical metrics. We focus on distance covered, which is calculated by summing the Euclidean distances between a player's predicted locations throughout a match. We initially found that the predicted distance covered is consistently higher than the actual distance covered, showing an average overestimate of 11.5% for *Agent Imputer*, 16.8% for the Time-Aware LSTM, 31.7% for the GNN, and 29.7% for XGBoost. This overestimation bias can be attributed to the large number of events that occur soon after another event — 29.9% of events happen within one second of the previous event. Typically, these are duels between a player on each team or a player receiving the ball and then quickly passing. In these instances, the players' predicted positions shift markedly, i.e., moving faster than the maximum player speed. This implies the models are not considering realistic player trajectories between events that occur in quick succession. Fixing this model issue is future work. For now, we use a post-processing step that combines events that occur within a second of each other, using the model output for the initial event as the player location when grouping. This leads to more accurate estimations of the distance covered. Results for the predicted distance covered averaged across different player roles are shown in Table 3.2.[3]

[3]Distance is normalised as if all players play for 90 minutes (i.e., accounting for extra time and substitutions). Players who played for less than 20 minutes were excluded from these results. Results are averaged across all games for players in a certain role.

TABLE 3.2: Player distance covered results. Distance is averaged over the test data and given in kilometres. Absolute error is averaged over each game for each player role, and the 95% confidence interval is provided.

Role	Pred. Dist.	True Dist.	Abs. % Error
Goalkeeper	3.23	3.12	4.05 ± 2.52
Central Defender	8.26	8.42	4.31 ± 1.38
Wide Defender	8.97	9.08	4.99 ± 1.89
Central Midfielder	9.37	9.48	3.66 ± 1.26
Wide Midfielder	8.73	8.97	2.79 ± 0.76
Central Attacker	8.34	8.35	6.03 ± 2.32
Wide Attacker	8.26	8.70	5.33 ± 0.76

3.5.3 Pitch Control Analysis

To test a model application which considers agent locations relative to each other within the system, we calculate pitch control of teams. Pitch control (Spearman et al., 2017; Spearman, 2018; Fernández et al., 2019) is a popular downstream analysis tool within football analytics that uses player locations and trajectories to compute the area in which a team dominates. It splits the pitch into a grid of zones and computes the probability of the attacking team controlling the ball if it arrived in that zone. For example, if a defending player is closer to the zone than the nearest attacker, it is unlikely the attacking team will control the ball in that zone. We perform a similar analysis to that used in (Omidshafiei et al., 2022) — we compute pitch control using *Agent Imputer* outputs and compare it with the pitch control computed from actual player locations in tracking data. Table 3.3 shows the mean absolute error of pitch control for each model, averaged across all pitch zones, and shows that the *Agent Imputer* performs best.

TABLE 3.3: Pitch control performance compared to the ground truth. Calculated over test data. The ± values represent 95% confidence intervals.

Model	Mean Average Error
Baseline 1	0.272 ± 0.002
Baseline 2	0.304 ± 0.002
Baseline 3	0.305 ± 0.002
XGBoost Regression	0.150 ± 0.001
GNN	0.149 ± 0.001
Time-Aware LSTM	0.135 ± 0.001
Agent Imputer	**0.130 ± 0.001**

We provide an example of pitch control outputs in Figure 3.9. These visualisations can help clubs review their team's dominance in particular pitch areas in game scenarios.

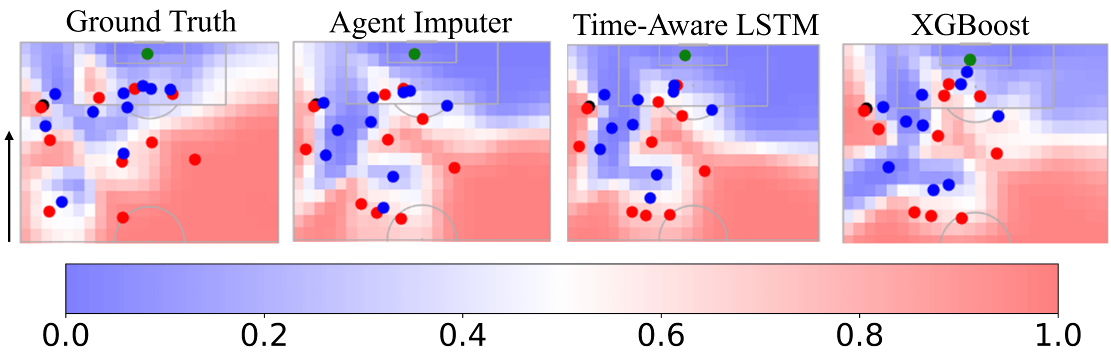

FIGURE 3.9: Pitch control diagrams comparing three models with ground truth. Red regions are in control of the attacking team, and blue regions are in control of the defending team. We only show the attacking half as the remainder of the pitch is entirely under the attacking team's control (red). Arrow refers to the attack direction.

This plot also highlights the differences in predicted player locations between the different models. It can be seen that, for the defending team, players are further apart and a structured horizontal defensive line is shown in the *Agent Imputer* predictions. Whereas for the other models, some defenders can be seen to be stacked vertically - a very unlikely prediction given the usual defensive setup and strategy employed by football teams. This suggests that the GNN module is correctly learning how agents spatially interact.

3.5.4 Player Heatmap Analysis

In this section, we use predicted player locations from the *Agent Imputer* model to generate player heatmaps over an entire game. This provides coaches with useful information on the areas most frequently covered by their players. Note, while it is possible to generate heatmaps from event data, due to the nature of the data, they will be missing ∼95% of player positions. Figure 3.10 presents some comparisons between heatmaps of a central defender, central midfielder, central attacker, wide defender, wide midfielder, wide attacker and goalkeeper using imputed positions.

The imputed heatmaps show many similarities to the ground truth heatmaps. Note the model has picked up interesting features, such as the small, distinct region for the central defender near the opponent's goal. This most likely stems from set-pieces (such as a corner), demonstrating the *Agent Imputer* has learnt features of the game state and how this influences player position. Furthermore, the central defender heatmap is offset to the left-hand wing of the pitch. In comparison, the central attacker heatmap is more central. This shows the model has learnt role-specific information — typically there are two or three central defenders, allowing them to focus on particular sides of the pitch, whereas a central attacker is likely more free to roam in an unstructured manner. Expected differences are also evident between central and wide players, with wide players

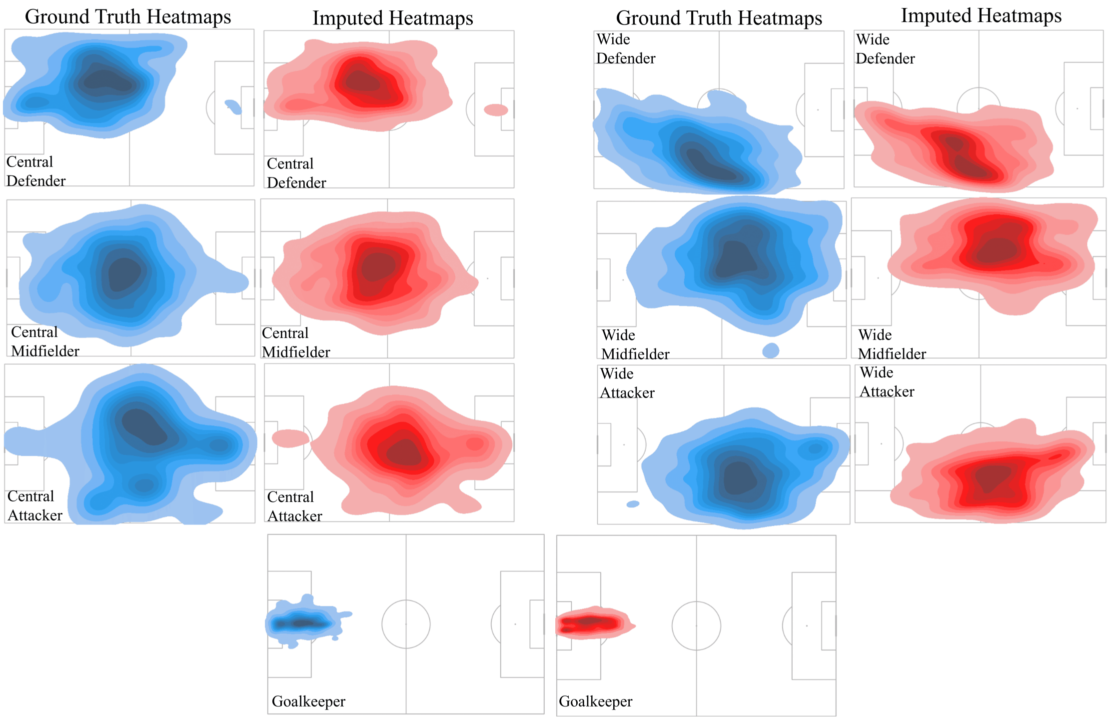

FIGURE 3.10: Example *Agent Imputer* heatmaps compared with ground truth heatmaps. Darkest regions indicate the most frequently occupied areas. Players are playing from left to right.

predominantly occupying areas near the pitch's flanks. Finally, the goalkeeper heatmap is highly concentrated around the goal area, accurately reflecting the constrained spatial behavior inherent to this position.

3.6 Discussion

In this section, we further discuss our findings (Section 3.6.1) and highlight limitations and areas of future work (Section 3.6.2).

3.6.1 Analysis of Results

As shown by our results (Section 3.4.3), our model is able to impute agent locations using very sparse data, and do so with greater accuracy than naïve baselines and other ML models. The increase in performance compared to XGBoost suggests that the temporal and inter-agent dynamics captured by the LSTM and GNN components offer value in learning agent location.

Our *Agent Imputer* model also enables easier access to insightful use-cases in football, such as player physical metrics, player heatmaps, and team pitch dominance analysis (see Section 3.5). These use cases could be implemented by lower league clubs with fewer resources to perform player evaluation, scouting, and game analysis processes similar to those of elite-level clubs. This would help close the performance gap between these teams.

As an example, the player heatmaps are shown to be strongly correlated to the ground truth heatmaps exhibited in-game. This can provide value to an opposition analytics scout who could monitor an opponent and get their players to exploit areas that are frequently left uncovered by the opposition team. It may also allow scouts to find players who occupy dangerous space more often. Additionally, the analyses on off-ball spatial teamwork and individual off-ball player contributions presented in Chapters 4 and 5 could also be performed using the estimated tracking data generated by the *Agent Imputer* model. Indeed, an adapted version of the *Agent Imputer* underpins the simulations of future player movement in Chapter 4. Moreover, Chapter 6, which focuses on long-term team selection and player replacement informed by player fatigue and injury risk, directly relies on the *Agent Imputer* to estimate in-game physical workload metrics, making the model accessible to a much wider range of clubs.

The predictive ability of our model demonstrates that positioning in football is somewhat predictable. This is likely due to tactical positioning, where teams will be set up to defend and attack in consistent shapes. Work that considers the game theoretic implications of team tactics may introduce further insight into the predictability of teams and how this varies over time (Beal et al., 2020a, 2021). Furthermore, players in specific roles will be given instructions on where to position themselves in specific patterns of play. Our model can learn these patterns to predict future occurrences of such positioning.

3.6.2 Future Work and Limitations

Our *Agent Imputer* approach could also be applied to other MAS with limited observability, such as searching for injured civilians in disaster response systems or predicting human movement from phone location data (McInerney et al., 2012, 2013). We plan to validate the model in these domains in future work. A change of domain would inevitably lead to some differences in problem structure, such as the bounds of the environment or the sparsity of the dataset. These differences must be considered when applying our method to other domains. Another research direction would be to make the model probabilistic so that uncertainty in agent locations can be quantified.

Our model performs well at estimating location. However, as discussed in Section 3.4.3, extracted trajectories from these estimations can sometimes be unrealistic when events happen in quick succession (although we note this issue is worse in the non-temporal models we compared our approach to). Future work could consider predicting the trajectory of an agent over a number of timesteps simultaneously to mitigate this issue.

Further studies could also compare our proposed *Agent Imputer* model against state-of-the-art models that use tracking rather than event data, such as the methods presented in (Omidshafiei et al., 2022) and (Le et al., 2017b). In this work, we have only compared with other models that use event data. Comparing to tracking data models would facilitate a better assessment of the effectiveness of our imputation approach in downstream analysis tasks. Similarly, we could compare the prediction accuracy of player locations for our model to other spatiotemporal models that are not designed for environments with non-uniform timesteps and limited agent observability.

Future work could extend our dataset to include a wider variety of leagues or apply our model to a different team sport. This would help the model learn a wider range of scenarios beyond the ones that occur within our current dataset. Furthermore, this would facilitate the development of more specific models, such as those focusing on a particular player or team. Finally, beyond the analyses in Figures 3.5 and 3.6 on the importance of observation time and player role, a full feature importance study, covering all inputs and their contribution to model predictions and performance, could provide further insight into the key drivers of agent position prediction.

3.7 Summary

This chapter presents a novel *Agent Imputer* model to address a multi-agent behavioural prediction problem in dynamic environments with limited system observability. We apply this task to football by imputing off-ball player locations using only on-ball data. We find that our model can impute player location to within $\sim$6.9 metres, outperforming multiple baseline imputation models. We perform a deeper analysis of model performance, examining accuracy over different game times, player roles, observation rates and pitch zones. We also present off-ball football analytics applications facilitated by our model, which allow lower-league clubs to perform processes similar to those of elite-level clubs without requiring access to expensive data. This novel work could be applied to other real-world domains involving limited agent visibility, such as disaster response or tracking daily human movement from phone data.

The next chapter leverages off-ball player locations in football to introduce a method for quantifying spatial coordination between agents and demonstrates how collective defensive positioning can be optimised to reduce attacker threat in defence-based scenarios.

Chapter 4

Learning and Optimising Spatial Teamwork in Multi-Agent Teams

In this chapter, we introduce a novel method for assessing agent teamwork based on their spatial coordination. Our approach models the influence of spatial proximity on team formation (TF) and sustained spatial dominance over adversaries using a Multi-agent Markov Decision Process (MMDP). We develop an algorithm to derive efficient teamwork strategies by combining Monte Carlo tree search (MCTS) and linear programming. Applied to team defence in football using real-world data, our approach reduces opponent threat by 21%, outperforming optimised individual behaviour by 6%. Additionally, our model enhances the predictive accuracy of future attack locations and provides deeper insights compared to existing teamwork models that do not explicitly consider the spatial dynamics of teamwork.

4.1 Introduction

TF problems in MAS involve heterogeneous agents collaborating to achieve common goals such as task completion or risk minimisation. In TF problems, the environment is often dynamic, and agents face uncertainties about the positions and intentions of other agents. In some settings, agents compete against adversaries, adding complexity to decision-making. Understanding these uncertainties and choosing the characteristics of optimal teams is crucial for developing efficient TF systems.

Recent research explores various criteria for team formation, such as task proximity (Capezzuto et al., 2021), agent waiting times (Amador et al., 2014), and the alignment of agent roles to tasks (Aswale and Pinciroli, 2023). For example, in social ridesharing, agents form teams to minimise travel times (Bistaffa et al., 2017), while in fire response, teams are formed within spatial and temporal constraints to minimise job completion time (Chen et al., 2021). These methods assess team value using predefined metrics such as task completion time or team suitability for nearby tasks (Parker et al., 2016), often overlooking specific agent-to-agent interactions. Beal et al. (2020b) studies directed agent interactions in past teams to predict future team performance. However, the approach does not consider spatial constraints commonly found in TF problems (O'Leary and Cummings, 2007).

In many real-world teams, agents form spatial arrangements to optimise team dominance across an area. For example, Orcas encircle prey to minimise the chance of escape, and officers spread to cover criminal hot spots in security settings. While previous TF models have explored agent proximity and spatial constraints, such as information limitations (Bulka et al., 2007) or the impact of spatial dispersion on agent communication (O'Leary and Cummings, 2007), they do not focus on optimising agent communication. Studies in UAV teams consider agent coordination to maximise task coverage in an area (Baker et al., 2016), however, no existing model studies optimising agent interactions to maximise spatial dominance against an adversary. We define this concept as spatial teamwork.

Against this background, we develop a novel spatial teamwork model focused on nearby agents, addressing agent uncertainties regarding the intentions of teammates and adversaries. We use a stochastic MMDP to model these dynamics. To navigate these uncertainties, MCTS is used to explore future scenarios and determine effective spatial teamwork in dynamic conditions. This approach allows agents to act individually or as subteams, where they coordinate actions with a shared objective function.

We validate our approach using real-world football data, focusing on defensive strategies to minimise goal concession probability. Players are modelled as agents with their own actions, ability to coordinate with teammates and objectives. They operate with incomplete information regarding future ball and player positions. Football transitions, including player movements, are learned from real-world data.

Thus, this chapter makes the following novel contributions:

- We propose a novel model of spatial teamwork where a dynamic defence domain with incomplete information is modelled as an MMDP.
- An algorithm based on MCTS and linear programming is proposed to optimise individual and teamwork-based decision-making.
- We learn agent behaviour from real-world data and use it to simulate future scenarios in a dynamic environment.
- We compare our algorithm to real-world outcomes and show that our model reduces opponent threat by $\sim$21%, outperforming optimised individual behaviour by $\sim$6%.

By so doing, we set a benchmark for optimising spatial teamwork in dynamic real-world domains with adversaries.

The rest of the chapter is structured as follows. Section 4.2 introduces the teamwork model and Section 4.3 discusses MCTS to optimise agent contributions. In Section 4.4, we apply our model to football and evaluate our approach in Section 4.5. We critically review our findings in Section 4.6 and then summarise the work in Section 4.7.

4.2 Spatial Teamwork Model

In this section, we model an attack-defence scenario inspired by real-world systems, including team sports and security settings, such as port security (Shieh et al., 2012) and nature security (Xu et al., 2017). We define the environment, agents, objectives, and agent subteams.

4.2.1 Basic Definitions

We define a defence scenario E as a sequence of timesteps $E = \{e_1, ..., e_T\}$, where T is the number of timesteps and may vary for different scenarios. There are variable time intervals between timesteps. These scenarios may represent waves of attacks in security settings or sequences of on-ball events (possession phases) in football. Each scenario results in an outcome Ω_E, such as the loss of possession in football or a failed attack in security.

Agents operate in a two-dimensional Euclidean space Λ, defined as a set of (x, y) coordinates on a plane. In contrast to typical defence scenarios with independent targets, this scenario requires defence of the entire plane, where each location $(x, y) \in \Lambda$ has a value depending on strategic and temporal factors. Formally, the value of a location (x, y) at timestep t is denoted as $V_t^{x,y} \in \mathbb{R}$.

Each scenario has a set of agents C. The agents are partitioned into two adversarial teams, a defending team C_α, and an attacking team C_β. These designations are consistent throughout the scenario timesteps. At each timestep t, each agent $c \in C$ has a set of possible actions $a_t^c \in A_t^c$. These actions refer to the movement of an agent from their current location, defined as ϕ_t^c, to a new location ϕ_{t+1}^c. We also define agent-specific characteristics, such as role and physical condition, that may impact agent behaviour as ζ_t^c. The collective locations of a set of agents, defined as their spatial structure, are formed by both teams at each timestep t based on the actions chosen by each agent at the previous timestep. These spatial structures can be denoted as Φ_t^α and Φ_t^β where Φ_t^α contains all defending agent locations $\phi_t^\alpha \forall \alpha \in C_\alpha$. Unlike in Chapter 3, where Φ included incomplete agent locations due to sparse event data, the remainder of this thesis assumes full observability of agent positions at every timestep. We build on these definitions to explain team objectives in the next subsection.

4.2.2 Team Objective

Attackers and defenders aim to maximise control over the space Λ. At timestep t, the control that defenders have over a location (x, y) is represented as a probability, $\Pr(\vartheta_t^{x,y} | \Phi_t^\alpha, \Phi_t^\beta)$, where $\vartheta_t^{x,y}$ indicates that the defenders C_α control the location (x, y) at time t, which depends on the spatial structures Φ_t^α and Φ_t^β. Defensive utility is defined as a penalty corresponding to the attacking team's dominance over Λ, weighted by the value of the locations $(x, y) \in \Lambda$. Defending agents aim to minimise this penalty. We represent team C_α's utility at time t as:

$$U_t(C_\alpha) = -\int_{(x,y)\in\Lambda} V_t^{x,y}(1 - \Pr(\vartheta_t^{x,y} | \Phi_t^\alpha, \Phi_t^\beta))\, dx\, dy \tag{4.1}$$

The utility function is defined as a penalty to incentivise defenders to end scenarios (e.g., attacks in security or team sports) faster while minimising the valuable space available to attackers, with the spatial integral and value weighting ensuring that focus is applied to controlling high-impact regions across the entire space. The optimal spatial structure for defenders C_α at timestep t is defined as $\Phi_t^{\alpha,*} = \arg\max_{\Phi_t^\alpha} U_t(C_\alpha)$. The defenders aim to minimise attacking control (i.e., maximise their utility) throughout an entire scenario, defined as:

$$U_E(C_\alpha) = \sum_t^T U_t(C_\alpha) \tag{4.2}$$

Defenders face multiple challenges when optimising $U_E(C_\alpha)$. Firstly, there are spatiotemporal movement constraints requiring agents to move feasible distances between locations within timestep intervals, bounded by physical and tactical limitations (see Section 4.4). Additionally, agents do not know the actions that the attackers or other defenders will choose. Hence, defenders must anticipate future states and plan movement strategically to maximise expected utility. While attackers share similar goals and complexities, this work focuses on the defensive perspective.

4.2.3 Agent Contribution and Subteam Formation

At each timestep t, defending agents can form subteams with nearby teammates to optimise spatial coverage. A subteam is defined as $\Psi \subseteq C_\alpha$, with $\boldsymbol{\Psi}$ as the set of all subteams. Each agent can only belong to one subteam at t, determined by proximity-based K-means clustering, with the number of clusters chosen by maximising the

Silhouette coefficient (Rousseeuw, 1987) (see Appendix B.3.1 for implementation details). Subteam members share an objective function and action at t, providing awareness of their subteam members' actions, similar to human or intelligent systems communication. The objectives of agents and subteams are defined as:

Individual Agents Agents choose actions to improve their contribution to team utility. Using cooperative game theory, we value an agent's impact by their marginal contribution. The marginal contribution of agent α to the team C_α at timestep t is defined as $\Theta_t^\alpha = U_t(C_\alpha) - U_t(C_\alpha \setminus \{\alpha\})$ where $U_t(C_\alpha \setminus \{\alpha\})$ is the utility achieved by the team if the agent α was not in the team C_α at timestep t.

Agent Subteams Instead of maximising individual contributions, agents may spatially coordinate with teammates in a subteam Ψ to maximise their joint marginal contribution. The marginal contribution of a subteam Ψ at timestep t is $\Theta_t^\Psi = U_t(C_\alpha) - U_t(C_\alpha \setminus \Psi)$. Agents must use context to decide whether to act individually or coordinate in a subteam. Subteams offer certainty about each member's actions and can improve defensive efficiency, such as in patrols or football, where teammates may surround adversaries. However, individual actions may be more effective for defending nearby vulnerable locations.

4.2.4 Spatial Teamwork Multi-Agent MDP

Markov decision process (MDP) models can be used to define how a stochastic environment changes as a decision-maker interacts with it. An MMDP extends this to multiple decision-making agents, each with their own action sets. We formulate a scenario E as an MMDP defined as $\mathcal{M}_{\mathrm{MA}} = \langle S, \{A^c\}_{c \in C}, P, \{R^c\}_{c \in C}, \gamma \rangle$ with a set of states S, a set of actions A^c for each agent c, P is the transition function, R^c is a reward function for each agent c and γ is a discount factor. Each MMDP state $s \in S$ is a tuple $s = \langle \Phi_t^C, \zeta_t^C \rangle$ where Φ_t^C and ζ_t^C are the agent locations and characteristics defined previously in Section 4.2.1.

The action space for each agent includes all possible spatial movements within the realistic distance they can cover between timesteps (see Section 4.4; this may differ depending on the application domain). The transition function $P = \Pr(s'|s, \{a^c\}_{c \in C})$ represents the probability distribution over possible next states s' given the current state s, which contains information on agent locations and characteristics, and each agent's selected action. The reward of action a for agent c chosen at state s depends on the actions of other agents and the resulting state s', denoted as $R^c(s, a^c, s')$. This reward reflects the agent's immediate contribution to team utility at state s', $\Theta_{t'}^c$. The discount factor determines the preference for immediate versus future rewards.

The MMDP reaches a terminal state when the scenario E ends, such as when possession is lost in football or an attack fails in security. The probability of a scenario concluding is captured in the transition function P. In the next section, we discuss techniques for optimising agent and team rewards.

4.3 Optimising Team Decision-Making

This section outlines how team reward is maximised by optimising agent actions and coordinating spatial teamwork. We model decision-making in three steps: first, at the level of individual agents, then at the level of subteams, and finally at the level of the entire defending team.

4.3.1 Agent Objective

Agents and subteams aim to maximise their long-term contribution, extending beyond just the current timestep. We formally define these objectives below.

Agents An agent c aims to maximise expected long-term reward over a scenario from state s. This is formalised by the Bellman equation:

$$V(s) = \max_{a^c}(\sum_{s'} \Pr(s'|s, a^c) \cdot (R^c(s, a^c, s') + \gamma V(s'))) \tag{4.3}$$

Where $V(s)$ is the value at state s. The agent must therefore balance immediate rewards against future utility, while reasoning under uncertainty about the future actions of teammates and opponents. This requires agents to predict behaviour and future state transition probabilities.

Subteams Agents in a subteam choose a joint action to maximise their collective long-term reward, $V(s)$. This is formally defined as:

$$V(s) = \max_{a^\Psi}(\sum_{s'} \Pr(s'|s, a^\Psi) \cdot (R^\Psi(s, a^\Psi, s') + \gamma V(s'))) \tag{4.4}$$

Where a_Ψ denotes a joint action taken by subteam Ψ. More details on subteam actions are given in Section 4.4. Due to the many possible actions for each agent, there are $|\mathcal{A}|^{|C|}$ possible next states from a current MMDP state s. Despite the use of action discretisation in Section 4.4, there are still approximately 1×10^{26} possible new states from the current state. Given this complexity, computing exact solutions for Equation 4.3 is infeasible.

To address this, we approximate agent behaviour (in Section 4.4), and use MCTS to estimate Q-values of actions by simulating many possible scenario outcomes resulting from actions.

4.3.2 Monte Carlo Tree Search

MCTS is an anytime search algorithm used to optimise actions. We apply MCTS at a timestep t for all agents $\alpha \in C_\alpha$ and all potential subteams $\Psi \in \mathbf{\Psi}$ by iteratively selecting actions and simulating the rest of the scenario E. The MCTS algorithm runs iteratively as follows:

1) **Selection -** Select the most promising action from the root using UCB1 (Auer et al., 2002) until an unexplored or terminal node is reached. Each node, η_s, with state s, leads to a child node with state s', determined by action a and transition function P.

2) **Expansion -** Expand a node η_s by randomly selecting an unexplored action.

3) **Simulation -** At leaf node η_s, approximate $V(s)$ using cumulative reward in an MMDP simulation until a terminal state is reached, using our real-world model of state transition probabilities (see Section 4.4).

4) **Backpropagation -** Backpropagate the value of the new child node η_s up the tree to the root.

Our transition function, which we compute using a deep learning-based model (as described in Section 4.4), can be inefficient for single simulations. To improve MCTS convergence speed, we parallelise MCTS expansion and simulation. Unlike traditional leaf parallelisation (Cazenave and Jouandeau, 2007), our approach uses a single thread and instead uses parallel tree nodes to batch process state transitions using P.

At node η_s, we expand a random action and perform ϖ transitions (100 in this work) to reach ϖ leaf nodes $\eta_{\mathbf{s}'} = \{\eta_{s_1}, ..., \eta_{s_\varpi}\}$. We batch process these nodes using the transition model P to produce new states $\mathbf{s}'$. Each new leaf node is then simulated until reaching a terminal state, generating estimated values $V(\mathbf{s}') = \{V(s_0), ..., V(s_\varpi)\}$. In backpropagation, the value of the parent node η_s is updated with the average cumulative reward across all expanded nodes $V(\mathbf{s}')$, adjusting for the node count ϖ.

4.3.3 Optimising Decision Selection

Given that we have optimised actions for each defending agent $\alpha \in C_\alpha$ and subteam $\Psi \in \mathbf{\Psi}$ at timestep t, the goal is to maximise the long-term utility of the entire team

$U_E(C_\alpha)$. As agents and subteams independently select actions without knowing the rest of the team's actions, it is crucial to evaluate whether each agent should act individually or within their subteam. To find the combination of individual and subteam actions that maximises $U_E(C_\alpha)$, we treat this as a linear programming problem:

$$\begin{aligned} \text{maximise} \quad & U_E(C_\alpha) \\ \text{subject to} \quad & \sum_{\alpha \in C_\alpha} \alpha_\Psi + \sum_{\Psi \in \mathbf{\Psi}} |\Psi| = |C_\alpha| \\ & \sum_{\Psi \in \mathbf{\Psi}} |\Psi| > 1 \end{aligned} \tag{4.5}$$

Where α_Ψ is a binary indicator variable, set to 1 if agent α acts individually and 0 if they are within one of the subteams. To solve this linear program, we extract all combinations of individuals and subteams and approximate $U_E(C_\alpha)$ for each using one thousand MMDP scenario simulations, similar to MCTS simulation. The action set that maximises $U_E(C_\alpha)$ is chosen as the team's joint action. We term this optimisation strategy *spatial teamwork*. The entire optimisation process, which includes MCTS for agents and subteams to identify the joint action of the team, is shown in Figure 4.1 and is called the teamwork optimiser. The teamwork optimiser allows for inter-agent interaction to maximise joint subteam contributions, extending beyond a simple individual strategy where each agent selects their own optimised action.

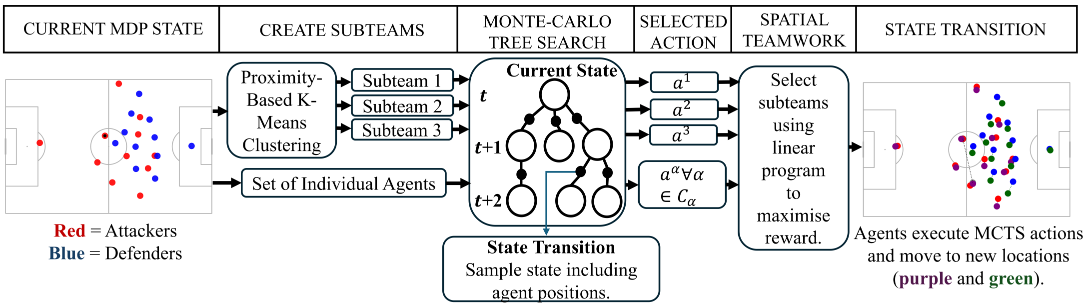

FIGURE 4.1: The agent optimisation process. Starting from the current MMDP state, K-means clustering determines subteams. MCTS approximates optimal actions for each subteam and agent, and the linear program determines which subteams should form, resulting in an action for each agent. Executing these actions causes state transitions, as illustrated by a football example.

We note that subteams are fixed during both MCTS and the MMDP scenario simulations: agents remain in their K-means–assigned subteams (based on spatial proximity at the current state) for the duration of rollouts. Consequently, each agent optimises a consistent objective aligned with its initial assignment (either as an individual or within its subteam), promoting long-term coordination at that granularity. In future work, we will examine dynamic reassignment of subteams at each new state within MCTS simulations to assess its impact on optimised behaviour.

4.4 Applying Spatial Teamwork to Football

In this section, we apply our model to real-world attack-defence scenarios in football. Specifically, we highlight the relationship between football and our model, and give context to agent actions and state transitions in the football domain.

4.4.1 Multi-Agent Environment in Football

We define a sequence of on-ball events (period of possession) in football as a defence scenario E, consisting of timesteps t representing events (e.g., a pass or shot), and ending with an outcome, Ω_E, which is either a goal or a possession loss. The football pitch, Λ, is discretised into a grid of $(I \times J)$ zones denoted as Z, where $Z_{i,j}$ represents a specific zone. In this paper, the grid comprises (25×16) zones. The value of all pitch zones, V_t^Z, which takes shape $(I \times J)$, are computed using the popular Expected Threat model[1], indicating the probability of a goal in the next five events from each zone in Z respectively. Additionally, we update the team utility to accommodate zones:

$$U_t(C_\alpha) = -\sum_{i \in I} \sum_{j \in J} V_t^{Z_{i,j}} \cdot (1 - \Pr(\vartheta_t^{Z_{i,j}} | \Phi_t^\alpha, \Phi_t^\beta)) \tag{4.6}$$

Where $\Pr(\vartheta_t^{Z_{i,j}} | \Phi_t^\alpha, \Phi_t^\beta)$ is the probability of defenders C_α winning possession if the ball arrives at zone $Z_{i,j}$. This is calculated using zone centroids and a physics-based ball control model (Spearman, 2018). In football, C_α and C_β are the defending and attacking teams, each with 11 players. For each scenario E, the defending team C_α is determined by ball possession. The locations of players on team C_α at time t represent the team's spatial structure Φ_t^α.

To accommodate the football domain, we define an MMDP state s as a tuple $s = \langle \Phi_t^C, \zeta_t^C, B_t^Z \rangle$ where $B_t^Z \in Z$ is the current zone that the ball B occupies at timestep t. Additionally, for this domain, the player characteristics for a specific player c, ζ_t^c, are a 3-tuple, $(\dot{x}_c, \dot{y}_c, \kappa_c)$, where $\dot{x}_c$ and $\dot{y}_c$ denote the x and y velocity of player c respectively and κ_c represents the player's role in the team (e.g., centre defender). Player velocity at a new state s' is calculated as the average velocity required to move from their location in state s to their location in state s' within the time between timesteps. The immediate reward R^c represents the contribution of a player c towards control of valuable pitch space at time t. In this work, we set the discount factor to 1.

[1]Expected Threat: `https://karun.in/blog/expected-threat.html`. Last accessed July 12, 2024.

4.4.2 Transition Function for Football

The transition function in an MMDP predicts state changes probabilistically. In our football framework, it models ball and player movements from state to state. We learn transition probabilities from real-world data. Our transition function, $P = \Pr(s'|s, \{a^c\}_{c \in C})$ is simplified using domain knowledge by making player positions at the next state conditional on ball movement. We define this as follows:

$$P(s', s) = \Pr(B^Z_{t+1}|s) \cdot \prod_{c \in C} \Pr(\phi^c_{t+1}|s, \{a^c\}_{c \in C}, B^Z_{t+1}) \tag{4.7}$$

Where $\Pr(B^Z_{t+1}|s)$ is the probability of the ball moving to each zone, or possession being lost, given the current state s and $\Pr(\phi^c_{t+1}|s, \{a^c\}_{c \in C}, B^Z_{t+1})$ is the probability of player c's next location given the current state, actions and ball movement. At each state, we use a ball transition probability model (Spearman, 2018) to compute $\Pr(B^Z_{t+1}|s)$ and sample from these probabilities during state transitions. For successful ball transitions, given the ball's predicted movement, we then predict future player locations using an adapted version of the *Agent Imputer* model introduced in Chapter 3. This adapted model is trained on 34 games of real-world football data and differs from its predecessor in several key aspects. Firstly, in contrast to the *Agent Imputer* model, which uses only on-ball data, this model incorporates a richer set of features from tracking data, including player locations and velocities, ball locations, player roles, and time elapsed between states. The time passed is estimated using a physics-based model for ball travel times (Spearman, 2018). Additionally, the model is designed to predict future player movement, rather than imputing missing trajectories. Further details on the input data, model structure, and training methodology, in comparison to the original *Agent Imputer* model, are provided in Appendix B.2.

Compared to the same baseline models used in Chapter 3, including XGBoost (2.47m), graph neural network (2.56m) and a simple spline (5.03m), our model achieves the lowest mean Euclidean error (2.31m). Since player movement is conditional on ball movement, from a state s and actions $\{a^c\}_{c \in C}$, the number of possible new non-terminal states equals the total number of pitch zones. The specific details of player actions and their subsequent impact on future positions are elaborated in the following subsection.

4.4.3 Player and Subteam Actions

In football, player positioning is influenced by ball movement and the positions of other players. We model player actions as directional biases that increase the likelihood of players moving towards certain areas. These biases adjust player velocity data in our transition model, increasing the probability of movement towards the desired direction while accounting for ball movement, teammates, and opponents. Adversaries receive no velocity alteration, imitating the average real-world behaviour. We discretise actions to reduce the action space size, speeding up MCTS convergence. The action space is described as follows:

Player Actions Each individual player will choose their action from the following possible directions:

Up	Top Right	Right	Bottom Right
Down	Bottom Left	Left	Top Left

These actions change the player's velocity by 2 m/s in the chosen direction. We set a maximum speed of 5 m/s, extracted from (Spearman, 2018).

Subteam Actions Subteam actions ensure that all players execute the same action with the understanding that their sub-teammates are doing likewise, modelling player communication in football. For instance, if the action is 'Up', all subteam members gain additional velocity in that direction. To enhance flexibility, we use domain-based knowledge to expand the action space beyond a single directional bias:

Avoid ball	Towards opponent	Avoid opponent
Spread out	Close in	

To focus subteam actions on coordination rather than flexibility, we align each agent's velocity with the individual action that best matches the chosen subteam action.

4.5 Empirical Evaluation

In this section, we empirically evaluate our model. We first describe the dataset used in our experiments (Section 4.5.1), then introduce the baselines against which we compare our approach (Section 4.5.2). We subsequently present a series of experiments, including long-term simulations and analyses of varying subteam sizes. Details of the computational resources employed to run the optimisation algorithm and conduct the simulated experiments are provided in Appendix B.

4.5.1 Datasets

We evaluate our approach on all pass and shot events from the same 34-game real-world football events and tracking dataset from the K League 1 introduced in Chapter 3. In total, this amounts to approximately 34,000 events. Player positional and velocity information is obtained from the corresponding tracking data associated with these matches. As detailed in Chapter 3, these constitute gold-standard industry datasets that enable a rigorous evaluation of our approach. Further information on the extracted features and data preprocessing steps is provided in Appendix B.

4.5.2 Baselines

Our model is compared against real-world actions and numerous baselines, listed as follows:

Baselines:

- **Random -** Agents choose their movements randomly.
- **Real-world -** Threat accumulated in the real-world data.
- **Simulated Real-world -** Agents follow their real-world actions within our MMDP simulations, bound to the closest action in the action set.

Optimiser strategies:

- **Individual optimiser -** All agents optimise their contribution individually without any subteams.
- **Pair-Based optimiser -** Agents choose actions in pairs based on the Team-Centred model in (Beal et al., 2020b).
- **(Spatial) Teamwork optimiser -** Agents have the option to coordinate in subteams.

Inspired by Beal et al. (2020b), our Pair-Based strategy segments teams into agent pairs to maximise pass success frequency within pairs from past games. Unlike (Beal et al., 2020b), who focus on lineup selection and in-possession attacking link-ups, we modify this method to in-game defensive decision-making by splitting the 10 outfield defenders into five pairs (excluding the goalkeeper). This pair formation does not explicitly consider spatial context and proximity. Our MCTS model selects effective actions for these pairs. In the following subsections, we present our experimental results.

4.5.3 Experiment 1: Model Performance in Dynamic Defence Simulations

In this experiment, we evaluate the long-term effectiveness of our optimiser against real-world decisions and baseline strategies. For each event in the dataset, we simulate 1,000 independent rollouts of the next five events using our football MMDP environment. This allows us to capture the longer-term impact of an agent's immediate action, since even small positional adjustments can propagate across multiple timesteps and significantly alter match dynamics. Performance is quantified using the reward, which relates to the accumulated threat (goal concession probability) over the five events. Table 4.1 reports the results for each strategy, providing a direct comparison between real-world behaviour, baseline policies, and our proposed optimisation approaches.

TABLE 4.1: Average reward comparison for simulated play sequences using various decision-making strategies. The ± values represent 95% confidence intervals.

Strategy	**Average Reward**	**Reduction**
Random	-0.058 ± 0.001	-9.4%
Real-World	-0.053 ± 0.001	0.0%
Simulated Real-World	-0.051 ± 0.001	3.8%
Individual Optimiser	-0.045 ± 0.001	15.1%
Pair-Based Optimiser	-0.047 ± 0.001	11.3%
Teamwork Optimiser	**-0.042 ± 0.001**	**20.8%**

The teamwork optimiser reduces threat by 20.8%, which is 5.7% more than the individual optimiser, highlighting the benefits of spatial collaboration. Similar threats in real-world and simulated real-world scenarios support the MMDP model's accuracy in the real world. Optimising agent behaviour is critical in high-threat situations; therefore, we focus on these scenarios in Table 4.2. The results in these high-threat scenarios draw similar conclusions.

TABLE 4.2: Average reward comparison for simulated play sequences in high threat situations (≥1% goal probability). The ± values represent 95% confidence intervals.

Strategy	**Average Reward**	**Reduction**
Random	-0.084 ± 0.001	-18.3%
Real-World	-0.071 ± 0.001	0.0%
Simulated Real-World	-0.074 ± 0.001	-4.2%
Individual Optimiser	-0.065 ± 0.001	8.5%
Pair-Based Optimiser	-0.068 ± 0.001	4.2%
Teamwork Optimiser	**-0.061 ± 0.001**	**14.1%**

We analysed how the difference between the average rewards (teamwork optimiser vs. real-world reward) correlates with the number of goals each team conceded in our dataset. Despite noise from low scoring rates across 34 games, a significant Pearson correlation coefficient of 0.71 ($p < 0.05$) confirms the expected relationship between these variables.

Additionally, for every event, we compared threat levels at the next real-world event with those from an augmented next event where defenders chose the optimised actions. This myopic approach assesses performance beyond the MMDP simulation. Results showed that teams reduce threat at just the real next event by 1.5% $\pm$ 0.6 and 2.3% $\pm$ 1.2 by using the individual and teamwork optimisers respectively.

4.5.4 Experiment 2: Attack Location Prediction Using Subteam Collaboration

Having demonstrated in Experiment 1 that our optimiser reduces overall team threat, we now evaluate whether subteam performance provides additional value in predicting the spatial location of the next opponent attack. Specifically, we compare the subteam contribution metric (Θ_t^{Ψ}) against the individual contribution metric (Θ_t^{α}) as input features to a probabilistic classifier that predicts the zone of the next attacking action (pass or shot). For each event, the zone that each player or subteam centroid resides in, and the performance metric of that player or subteam, are used to form a grid of features matching the shape of the grid of pitch zones. For this experiment, the pitch is split into a grid of $Z = (I \times J)$ zones where $I = 6$ and $J = 4$, and a simple convolutional neural network (CNN) model utilises these spatial features to probabilistically predict the next attack zone.

Table 4.3 reports the predictive performance when using the subteam versus individual contribution metrics as input features to the CNN classifier. Subteam features are computed from subteam centroids, while individual features are based on player-level contributions. Predictive accuracy is evaluated using log-loss over five-fold cross-validation. Further implementation details for this experiment, including model architecture and training procedure, are provided in Appendix B.3.2.

TABLE 4.3: Log-loss scores of contribution metrics as features for predicting future attack zones. The $\pm$ values represent 95% confidence intervals.

Metric	Individual	Subteam
Passes	0.156 ± 0.003	$\mathbf{0.147 \pm 0.001}$
Shots	0.0100 ± 0.0004	$\mathbf{0.0093 \pm 0.0002}$

The results show that subteam contributions consistently outperform individual contributions in predicting attack locations, yielding lower log-loss scores for both passes and shots. This emphasises the value of analysing agent coordination, as it uncovers the impact of spatial teamwork on the predictability of future attacks in dynamic settings such as team sports.

4.5.5 Experiment 3: Impact of Subteam Configurations on Threat Reduction

In this experiment, we assess how varying numbers of subteams impact threat reduction in real-world scenarios, aiming to identify the optimal number of players for effective spatial teamwork. Figure 4.2 compares threat reduction using the teamwork optimiser for various numbers of subteams (i.e., the number of subteam clusters generated by the k-means algorithm).

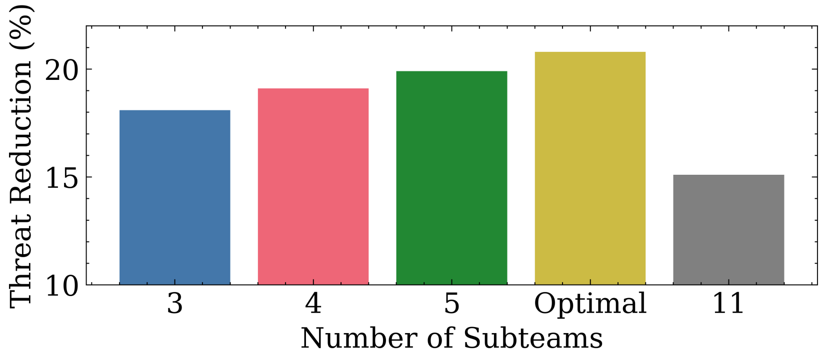

FIGURE 4.2: Effects of number of subteams on threat reduction.

These results show that clustering into five subteams, or using the optimal k-means silhouette score, yields the greatest threat reduction. This likely reflects a balance between coordination and flexibility, where dividing players into smaller subteams allows for more targeted action against specific threats than having just three subteams, while having too many subteams leads to overly fragmented groups where individual players act largely independently, reducing opportunities for effective coordination. We drive deeper insight into the internal structure of these subteams by looking at the distribution of subteam sizes for each number of subteams in Figure 4.3.

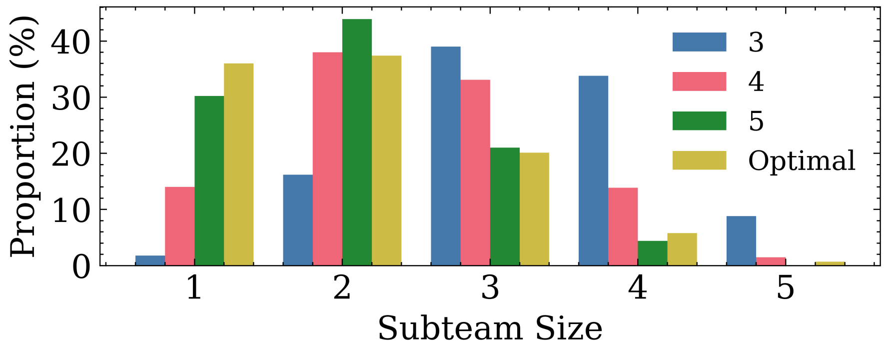

FIGURE 4.3: Distribution of subteam sizes for varying cluster numbers. Bar colours represent the number of subteams.

Figure 4.3 shows the distribution of subteam sizes when dividing the team into a particular number of subteams (denoted by the bar colours). For example, when the algorithm divides the team into three subteams, on average, two percent of the subteams only consist of one member. The results indicate that when clustering into five subteams or using the optimal k-means silhouette score, these clusters often form pairs. This suggests that spatial communication is optimal in pairs, aligning with (Beal, 2022), which found that teamwork in football passing is best represented by pairs.

In many domains, agents have specific roles, such as player positions in football. Our model most commonly suggests subteam pairs as: (`Centre Defender`, `Centre Defender`), (`Centre Forward`, `Centre Forward`), (`Right Defender`, `Right Midfield`) and (`Left Defender`, `Left Midfield`). This aligns with common expert views on important role-based pairs.

4.5.6 Experiment 4: Case Study Analysis of Defensive Optimisations

In this experiment, we compare real-world defensive scenarios with model-optimised decisions to show their impact on game outcomes. These visualisations offer strategic analysis to inform future team preparation. Figure 4.4 highlights a specific defensive situation, demonstrating a 0.4% decrease in goal probability ($U_t(C_\alpha)$) due to the model's recommendations. This case study shows how suggested adjustments in positioning and teamwork can significantly reduce goal likelihood and improve team preparation in a visually compelling way.

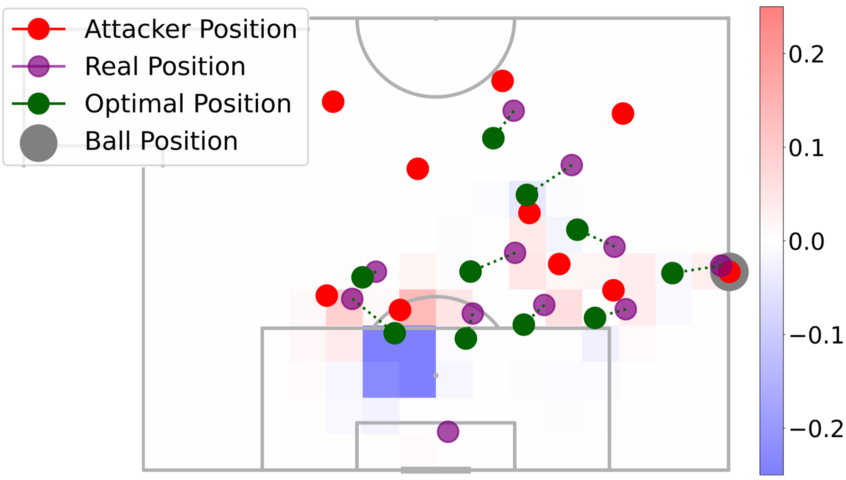

FIGURE 4.4: Optimised positioning to reduce threat. There is a 3.0% goal probability compared to 3.4% with real-world positions. The heatmap shows the change in goal probability for each zone when shifting to optimised positioning.

4.6 Discussion

We validate our spatial teamwork model with real-world football datasets, showing its effectiveness in improving spatial dominance in dynamic settings. Football is a suitable domain due to its rich spatiotemporal data and the clear objective of preventing goals and winning matches. Our model also has potential in other spatiotemporal defence domains, including patrol, security, and other team sports. The model may also be applicable beyond defence, such as for attacking situations in football or emergency response. Future research will extend and validate the model in these other domains as we obtain more high-quality datasets.

The Pair-Based Optimiser improves on real-world performance but is less effective at reducing threats compared to other optimisers. This suggests that, while optimising actions in pairs is valuable (as shown in Section 4.5.5), it is crucial to account for spatial context and proximity when optimising for specific scenarios to avoid suboptimal communication, where individual actions may be preferable. Our findings in Section 4.5.5 may further explain the effectiveness of the teamwork model in (Beal et al., 2020b) due to the implicit link between spatial proximity and passing in football. We also show that integrating spatial context into pair formation is effective for in-game player decision-making.

Our approach optimises agent behaviour by simulating outcomes and minimising average threat. Future work could explore strategies such as a minimax approach to minimise the threat of the worst-case scenario. We plan to evaluate various approaches and the impact of game state, such as maintaining a winning position, on optimised strategies.

The model's primary use is likely to be post-scenario analysis to identify poor teamwork patterns and suggest improvements. It can build databases of similar scenarios to assess typical team responses. In football, this may inform team training.

4.7 Summary

In this chapter, we propose a novel approach to agent coordination in dynamic environments, setting a benchmark for optimising spatial teamwork in adversarial settings. We model defence against adversaries as an MMDP and use MCTS to compute effective decisions. We apply this model to football defence, learning play sequences from real-world spatiotemporal data. Our MCTS approach reduces opponent threat by up to 21% compared to real-world outcomes, with an additional 6% reduction achieved using our spatial teamwork model over individual actions. We also show that our model better predicts future attack locations compared to an individual-based benchmark. Finally, we explain how our model results may offer deeper insight into previous teamwork models that do not explicitly consider the relationship between spatial proximity and teamwork.

Our framework has potential applications beyond football, offering a means to analyse teamwork and spatial coordination in other team-based domains by examining both the types and frequencies of player or agent interactions. It is also noted that player location data is required to model player teamwork and spatial coordination in this way; however, this data is expensive and difficult to access for clubs and researchers with limited resources. As a more affordable alternative, player locations can be estimated using the *Agent Imputer* model introduced in Chapter 3. While this provides a practical solution, it is important to acknowledge that imputed data introduces estimation uncertainty and will not fully capture the precision of ground-truth tracking data.

In the next chapter, we shift from optimising team-level spatial dominance to evaluating individual off-ball defensive impact based on their indirect influence on attacking passing outcomes. We develop a graph attention network that predicts pass receptions and uses attention weights to quantify each defender's influence on attackers. Building on this mechanism, we introduce two metrics for player-level defensive credit assignment, providing new insights to coaches and scouts on players' off-ball defensive performance.

Chapter 5

Evaluating Off-Ball Player Contributions using Graph Attention Networks

In this chapter, we introduce *GAPP*, a graph attention network (GAT) model that predicts football pass reception probabilities and provides interpretable insights into off-ball defending. Evaluating individual contributions from team members is a critical challenge across many domains, such as security and team sports. While progress has been made in valuing contributions, such as target defence in security or on-ball performance in football, many aspects of performance, such as off-ball football actions, remain difficult to quantify. Using attention mechanisms, *GAPP* captures off-ball player interactions and introduces two new metrics to quantify defender contributions. We tested *GAPP* on 306 English Premier League (EPL) matches, and showed it reduces binary cross-entropy (BCE) loss by 6.4% $\pm$ 1.5% compared to multiple baselines for pass reception prediction, while offering unique insights for off-ball defender evaluation for coaches, scouts and teams.

5.1 Introduction

Evaluating individual agent performance within a team is critical in many real-world domains, as it can provide actionable insights for improvement and optimise team strategies. In fields like security (Shieh et al., 2012) and sports analytics (Decroos et al., 2019; Merhej et al., 2021), performance is often assessed through measurable actions, such as stopping a security breach or completing a high-value pass. However, many contributions are indirect and harder to quantify, such as an agent's positioning influencing an attacker's target choice in security or a defender's off-ball positioning affecting pass reception in team sports. Traditional metrics often overlook these subtle impacts, highlighting the need for methods that evaluate indirect contributions to gain deeper insights into performance and improve decision-making.

In this work, we focus on evaluating off-ball performance in football, which has rich availability of real-world spatiotemporal datasets. In football, players position themselves to prevent dangerous opposition attackers from receiving the ball. Yet, these off-ball contributions are challenging to quantify due to the lack of direct links to outcomes like goals, unlike on-ball events such as passes or shots, which are easily measurable. Players spend the majority of a match (∼95%) off the ball, making off-ball performance a critical yet underexplored aspect of the game.

The availability of spatiotemporal datasets, such as event data (e.g., passes, shots) and tracking data (e.g., player locations), has advanced football analytics, enabling machine and deep learning models to evaluate player performance, primarily focusing on on-ball actions like attacking plays (Decroos et al., 2019; Fernández et al., 2021) or defensive actions (Merhej et al., 2021). However, assessing the impact of off-ball defensive positioning remains challenging due to its indirect influence on play outcomes in a dynamic environment. While some models predict future ball locations based on player positioning (Fernández et al., 2021), they do not quantify the specific influence of individual defenders on these outcomes.

Against this background, we introduce *GAPP* (Graph Attention for Pass Probabilities), a novel GAT-based model that predicts the probability of attacking players receiving the ball while providing interpretable insights into off-ball defensive contributions. By representing players and the ball as a dynamic graph, *GAPP* leverages GATs' attention mechanism to identify key relationships between players. The *GAPP* model not only outperforms baselines in pass reception prediction but is also used to provide two novel, attention-based metrics for evaluating off-ball individual defensive contributions.

Thus, this chapter presents the following novel contributions:

- We introduce *GAPP*, a novel graph attention-based model for predicting pass reception probabilities in football.
- The *GAPP* model achieves state-of-the-art performance for pass reception prediction, outperforming multiple baselines, including a traditional graph neural network (GNN), with a $\sim$6.4% $\pm$ 1.5% reduction in BCE loss.
- We introduce two novel attention-based metrics to quantify defender influence (DI) on attackers and overall defensive performance (DP) in a real-world multi-agent system.

- Using real-world event and tracking data from 306 EPL matches, we show how *GAPP* provides explainable insights into defender performance, including their impact on stopping dangerous attackers.

The rest of this chapter is structured as follows. Section 5.2 formalises the ball reception prediction problem, and Section 5.3 introduces our *GAPP* model. Section 5.4 introduces our defensive metrics. Section 5.5 presents empirical evaluations, followed by a case study on the EPL in Section 5.6. Section 5.7 discusses the results, and Section 5.8 summarises the work.

5.2 Predicting Ball Reception

In this section, we model a game scenario and formalise the problem of predicting football on-ball event (e.g., pass, dribble) outcomes. A game is represented as a time series of T events, $E = \{e_1, ..., e_T\}$, where each event $e_t \in E$ corresponds to an on-ball action. This time series is non-uniform, with varying intervals between events.

For each event $e_t \in E$, there is a set of N players, $C = \{c_1, ..., c_N\}$, who are involved in the game at that event. Each player belongs to one of two teams, determined by the function team : $C \to \{1, 2\}$. While players remain on the same team throughout the game, the roles of the teams (attacking vs. defending) can switch dynamically over the sequence of events E depending on which team has ball possession. At each event e_t, the player performing the on-ball action is denoted as c_t^n, where $c_t^n \in C$. The team of c_t^n, denoted as $\text{team}(c_t^n)$, forms the attacking team C_β at that event, i.e., $C_\beta = \{c \in C : \text{team}(c) = \text{team}(c_t^n)\}$. The remaining players, belonging to the opposing team, form the defending team $C_\alpha = C \setminus C_\beta$.

Our goal is to model the probability of a player $c \in C$ receiving the ball at the next event. Let e_t represent the current event and e_{t+1} the next event in the sequence E. At each event, we aim to predict the probability of each player $c_n \in C$ being the on-ball player at e_{t+1}, defined as $\Pr(c_{t+1}^n | e_t)$, where c_{t+1}^n indicates that player c_n possesses the ball at the next event. Accurately estimating this probability helps evaluate the likelihood of each player scoring and the value of the defending team's positioning (Section 5.4).

To model the probability of a player receiving the ball at the next event, we frame this as a graph learning problem. Each game event e_t is represented as a graph $\mathcal{G}_t = (\mathcal{N}_t, \xi_t)$, where $\mathcal{N}_t$ is the set of nodes and ξ_t the set of edges. The nodes $\mathcal{N}_t$ consist of the players $C = \{c_1, ..., c_N\}$ and the ball, denoted as a special node B, such that $\mathcal{N}_t = C \cup \{B\}$. Edges capture relationships between players. Each node $u \in \mathcal{N}_t$ has a feature vector $X_u \in \mathbb{R}^{d_\mathcal{N}}$, and each edge $(u, v) \in \xi_t$, where u and v are nodes, has a feature vector $Y_{u,v} \in \mathbb{R}^{d_\xi}$, where $d_\mathcal{N}$ and d_ξ are the dimensions of the node and edge features.

At each event e_t, predicting the next on-ball player c^n_{t+1} is framed as a node prediction task. We learn a function $f : \mathcal{G}_t \rightarrow [0, 1]$ that maps the graph to the probability of a specific player c_n receiving the ball at the next event: $\Pr(c^n_{t+1}|e_t) = f(\mathcal{G}_t)$. In this chapter, the *GAPP* model represents function f and is introduced in the next section.

5.3 GAPP Model For Ball Receiver Prediction

In this section, we outline the feature engineering (Section 5.3.1), model architecture (Section 5.3.2), and training process (Section 5.3.3) used for our *GAPP* model.

5.3.1 Feature Engineering

The node and edge features of our graph $\mathcal{G}_t$ are derived from football tracking data, which provides the coordinate locations, velocities and accelerations of each player and the ball on the pitch at a given timestep. However, these raw data lack important contextual information, such as the relationships between players and their positions relative to the goals. To address this, we apply feature engineering to extract features for our model. Specifically, for each graph $\mathcal{G}_t$, we compute the following node and edge features:

- **Node Features:** `Location (x, y), velocity (x, y), acceleration (x, y), x-distance to defending goal, y-distance to defending goal, x-distance to attacking goal, y-distance to attacking goal, Euclidean distance to the ball, angle (radians) to the ball, binary indicators for whether the node is on the home team, attacking, and on the ball.`
- **Edge Features:** `X-distance between nodes, y-distance between nodes, Euclidean distance between nodes, edge angle (radians), difference in node angles to the ball, binary indicator for nodes being on the same team.`

In these features, angles are defined as the bearing (in radians) relative to the pitch x-axis, which is aligned with the attacking direction. These features are selected to provide a comprehensive tactical representation of each event. For example, the velocity and acceleration features contribute to capturing aspects of the immediate temporal context and player intention. For the ball node, numeric ball-related features are set to 0, and binary features are set to 1. For the same-team binary indicator edge feature, where one node represents the ball, the value is 1 if the other node's team is attacking. Ball-related features are included to provide the model with easy access to this information during message passing. The orientation of play is consistent across all timesteps, with the attacking team always moving from left to right. Details on the preparation of these features, such as their standardisation for model input, are given in Appendix C.1. For this work, the graph is fully connected and directed, with edges in both directions between every pair of nodes. A fully connected, directed graph enables the model to assess the tactical significance of interactions between all players, allowing the attention mechanism to learn and highlight the most important relationships in a data-driven manner.

5.3.2 GAPP Model Architecture

Our *GAPP* model represents the function f, which predicts ball reception probability $\Pr(c_{t+1}^{n}|e_t)$ using a graph representation $\mathcal{G}_t$ and GATs, which dynamically learn the importance of node connections in graph-structured data. In team sports, interactions between players and the ball are complex and vary in significance. Unlike traditional GNNs, which assign equal attention to all node edges, GATs assign attention weights to edges, capturing the varying importance of player relationships for event prediction. The model's step-by-step architecture is given below and visualised in Figure 5.1.

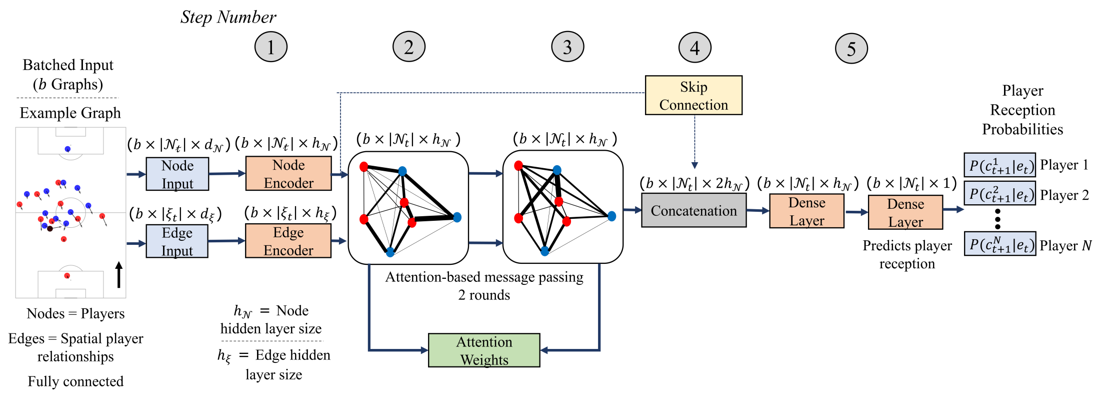

FIGURE 5.1: *GAPP* model architecture. All predictions are made using player feature information described in Section 5.3.1, with batch size b, number of nodes $|\mathcal{N}_t|$, number of edges $|\xi_t|$, number of node features $d_{\mathcal{N}}$, number of edge features d_{ξ}.

Step 1: Encoding Layer The input to the *GAPP* model consists of node features and edge features, represented as tensors $\mathbf{X_t} \in \mathbb{R}^{b \times |\mathcal{N}_t| \times d_{\mathcal{N}}}$ and $\mathbf{Y_t} \in \mathbb{R}^{b \times |\xi_t| \times d_{\xi}}$, where b is the batch size. The *GAPP* model uses dense encoding layers to project node and edge features into latent spaces, allowing the model to learn meaningful representations of these nodes and edges. The encoding outputs are latent representations with shapes of $(b \times |\mathcal{N}_t| \times h_{\mathcal{N}})$ and $(b \times |\xi_t| \times h_{\xi})$ for the nodes and edges respectively where $|\mathcal{N}_t|$ and $|\xi_t|$ are the number of nodes and edges and h_{ξ} are the hidden layer sizes of the node and edge encodings. In this chapter, $b = 64$, $h_{\mathcal{N}} = 32$ and $h_{\xi} = 16$.

Step 2: GAT Layer 1 The second step of the *GAPP* model is the first GAT layer, which performs attention-based message passing (Veličković et al., 2018). This layer takes the node and edge encodings as input. The attention mechanism in the GAT layer computes weights δ_{uv} for each edge $(u, v) \in \xi_t$, reflecting the relevance of node v and the edge connecting it to u. These attention weights are derived by calculating raw attention scores using a learned linear projection of node and edge features, followed by a LeakyReLU activation function, and then normalising them across all neighbours of node u using a softmax function so that they sum to 1. Each node in the graph updates its representation by aggregating the features of its neighbours, weighted by their respective attention scores. This allows the model to focus on the most relevant neighbours during message passing. Note that both GAT layers include self-loops and use 16 attention heads in this work, meaning that δ_{uv} is a vector of 16 values.

The output of this layer is a new set of node embeddings with shape $(b \times |\mathcal{N}_t| \times h_{\mathcal{N}})$, where the attention mechanism adaptively aggregates information from the graph. In football, this corresponds to modelling how players influence each other in ball reception. A rectified linear unit (ReLU) activation function is then applied to the output embeddings, followed by a dropout layer with a zeroing probability of 0.1, resulting in updated node embeddings.

Step 3: GAT Layer 2 The third step of the *GAPP* model is a second GAT layer, which captures higher-order relationships between nodes in the graph through another round of attention-based message passing. This layer takes the updated node embeddings from the first GAT layer and the edge encodings from Step 1 as input. Message passing is performed as in the previous layer, and the output is a new set of node embeddings of shape $(b \times |\mathcal{N}_t| \times h_{\mathcal{N}})$. As in step 2, a ReLU activation function and dropout layer are then applied to the new node embeddings. In the context of football, the second GAT layer enables the model to capture more complex relationships between players in the graph.

Step 4: Skip Connection and Concatenation The original encoded node features from Step 1 are concatenated with the updated node embeddings from Step 3 through a skip connection. This concatenation enables the model to retain original node information while including higher-order node relationships learned by the GAT layers. The resulting concatenated embeddings are of shape $(b \times |\mathcal{N}_t| \times 2h_{\mathcal{N}})$ and are used as input for the final dense layers of the model.

Step 5: Output Dense Layers The concatenated embeddings from Step 4 are passed through a dense layer with a ReLU activation function, resulting in a final latent representation for each node, with a shape of $(b \times |\mathcal{N}_t| \times h_{\mathcal{N}})$. This representation is then passed through another dense layer followed by a softmax activation function, resulting in the final node predictions with shape $(b \times |\mathcal{N}_t| \times 1)$. These predictions represent the probability, ranging from 0 to 1, of each player receiving the ball at the next timestep. While the model produces predictions for all nodes in the graph, including the ball node, in practice, the ball cannot be its own receiver. During training, these probabilities are masked out, and since no instances occur where the ball receives itself, they are effectively disregarded in both training and evaluation.

5.3.3 Model Training

We use football tracking data, which records all player locations during every on-ball event from 306 EPL matches in the 2023/24 season (see Section 5.5.1), to extract our target labels for the *GAPP* model, where each player is assigned a label of 1 if they are the next on-ball player, and 0 otherwise, resulting in one positive label at each timestep. For model training, we focus on attacking player receptions by applying an attacking player mask to the loss function, ensuring that updates are made only for attacking players. The model is trained using a BCE loss function for 200 epochs with a batch size of 64. The Adam optimiser (Kingma, 2014) is used with an initial learning rate of 0.003. The *GAPP* model hyperparameters and architecture were determined empirically through trial and adjustment with no formal hyperparameter tuning. Further details on the compute resources used to train the *GAPP* model are given in Appendix C.1.1.

5.4 Valuing Defensive Positioning

The key goal of the *GAPP* model is to provide insights into how defender positioning impacts attacker reception probabilities, using the GAT attention mechanisms to highlight the most relevant nodes and edges towards predictions. This section explains how attention weights are used to derive a novel defensive influence (Section 5.4.1) and defensive performance (Section 5.4.2) metric for evaluating off-ball defensive contributions.

5.4.1 Extracting Defender Influence

A defender's attention weight shows their influence on an attacker's reception probability. However, in its raw form, the attention weight only indicates the relative importance of the defender without specifying whether their impact is positive or negative. To address this, we compute a more interpretable DI metric that quantifies the effect of a defender on an attacker's reception probability.

Let a node v represent a defender $\alpha \in C_\alpha$ and a node u represent an attacker $\beta \in C_\beta$. The attention weights δ_{uv} quantify the attention assigned to defender v for attacker u, computed separately for each GAT layer and across all attention heads. The probability of attacker β receiving the ball, $\Pr(c^u_{t+1}|e_t)$, is predicted using the trained *GAPP* model f. To measure DI, we recompute the reception probability with δ_{uv} masked to 0 in both GAT layers: $\Pr(c^u_{t+1} \mid e_t, \delta_{uv} = 0)$, i.e. the prediction without defender v's attention. This masking allows the *GAPP* model to predict the reception probability of the attacker without considering the influence of defender v. This approach is inspired by explanation methods in the literature that assess prediction changes by limiting node features (Ying et al., 2019), but we uniquely apply it to quantify defensive contributions in MAS.

Defenders with higher attention weights are expected to have greater impacts on recomputed probabilities than defenders with lower weights (see Section 5.5.5). Using these probabilities, we define a new DI metric for defender v on attacker u:

$$\mathrm{DI}_{uv} = \Pr(c^u_{t+1}|e_t, \delta_{uv} = 0) - \Pr(c^u_{t+1}|e_t) \tag{5.1}$$

The DI metric measures how a defender changes an attacker's ball reception probability. A positive DI_{uv} value indicates that the defender reduces the attacker's reception probability by that amount. We note that Table 5.1 shows that the probabilities in Equation 5.1 are well calibrated, with the GAPP model outperforming state-of-the-art baselines in pass receiver prediction.

5.4.2 Computing Defender Performance

To evaluate the overall performance of a defender v, we first model the attacking threat of an attacker u. To measure attacking threat, we use the expected threat (xT) metric introduced by Singh (Singh, 2019), which divides the football pitch into a grid of zones (16×12 for this work) and estimates the probability of scoring within the next five on-ball events from each zone based on historical data. The threat of an attacker u, denoted as xT_u, is determined by identifying the zone they occupy and assigning the corresponding xT value.

We combined each attacker's threat (xT_u) with the defender's influence on them (DI_{uv}) to calculate defender v's overall positional performance. Our novel DP metric captures this for a defender v defined as:

$$\mathrm{DP}_v = \sum_{u \in C_\beta} \mathrm{DI}_{uv} \cdot \mathrm{xT}_u \tag{5.2}$$

This metric provides a novel evaluation of defender v's positioning by aggregating their influence on all attackers weighted by each attacker's scoring threat, offering a measure of a player's off-ball defensive contribution at an event e_t.

5.5 Empirical Evaluation

This section presents our empirical evaluation of the *GAPP* model and DP metric. We outline our dataset (Section 5.5.1), baseline methods (Section 5.5.2), and four experiments: assessing *GAPP*'s predictive accuracy (Section 5.5.3), evaluating its attention mechanism's fidelity (Section 5.5.4) and faithfulness (Section 5.5.5), and examining the DP metric's correlation with defensive actions (Section 5.5.6).

5.5.1 Datasets

Our *GAPP* model is trained and evaluated on 306 EPL games from the 2023/24 season. This data was supplied to us by Gradient Sports. The dataset includes event data (e.g., passes, shots) and tracking data (player positions and velocities). This is a gold standard industry dataset, enabling a rigorous evaluation of our model. We use tracking frames aligned with on-ball events to construct our feature set and target variables, resulting in 359,040 events. We evaluate our model using five-fold cross-validation, splitting the games into $\sim$80% for training and $\sim$20% for testing (245/61) in each fold, ensuring that each game appears in the test set exactly once. Further details on this dataset are provided in Appendix C.2.

5.5.2 Baselines

We evaluate the performance of *GAPP* against several baselines:

Mathematical Baselines:

- **Distance** - a naïve baseline assigning reception probabilities inversely proportional to player distance from the ball.
- **Spearman** - a physics-based model adapted from (Spearman, 2018) to predict reception probabilities instead of pass locations.

Machine Learning Baselines:

- **Dauxais** - a random forest model using distance-based features (Dauxais and Gautrais, 2018), using the majority of features from the original paper and extending them with our node and edge features.
- **XGBoost** - an XGBoost model (Chen and Guestrin, 2016) using the same node and edge features as the *GAPP* model. XGBoost is shown to be an effective predictor of pass reception in previous studies (Robberechts et al., 2023).
- **GNN** - a GNN with the same architecture as *GAPP* but replacing GAT with SAGEConv (Hamilton et al., 2017) layers without attention.

These baselines provide a diverse range of approaches, from simple distance-based heuristics to advanced machine learning (ML) and graph-based methods, enabling a rigorous evaluation of our model against current pass reception prediction models. Further implementation details of the ML baselines, including the libraries and hyperparameters used, are provided in Appendix C.1.

5.5.3 Experiment 1: GAPP Model Accuracy

We evaluate the predictive performance of the *GAPP* model and baselines towards pass reception in Table 5.1. We use various metrics, averaged across all attacker nodes, to assess the calibration of the model's predicted pass reception probabilities.

TABLE 5.1: Predictive performance of models for pass reception averaged across five folds, along with 95% confidence intervals. Bold results indicate the best performance.

Models	Test BCE	AUC Score	F1 Score
Distance	0.304 ± 0.005	0.655 ± 0.005	0.866 ± 0.001
Spearman	0.298 ± 0.001	0.690 ± 0.003	0.866 ± 0.000
Dauxais	0.286 ± 0.001	0.698 ± 0.003	0.867 ± 0.001
XGBoost	0.264 ± 0.001	0.760 ± 0.004	0.869 ± 0.001
GNN	0.233 ± 0.002	0.830 ± 0.005	0.897 ± 0.001
GAPP	$\mathbf{0.218 \pm 0.003}$	$\mathbf{0.855 \pm 0.006}$	$\mathbf{0.910 \pm 0.002}$

We find that *GAPP* achieves the best performance across all metrics, with a $\sim$6.4% reduction in BCE loss compared to the GNN model with the same architecture. This highlights the benefit of the attention mechanism for learning the importance of player relationships and interactions towards pass reception probability. Additionally, *GAPP* outperforms the XGBoost model with a $\sim$17.4% reduction in BCE loss, highlighting the advantages of using a graph-based approach to model the spatial and tactical context of game situations.

5.5.4 Experiment 2: Testing Model Fidelity

In this experiment, we evaluate the effectiveness of the attention mechanism in our model by testing its impact on pass reception predictions. We test this using various metrics introduced in the literature for evaluating attention-based models. Specifically, we use the fidelity metric (Zhao et al., 2023), a well-established metric for evaluating GNN models, which measures the deterioration in model performance (BCE Loss) when specific defender edge attentions are masked to 0. We compare three masking strategies:

- Top-K - Masking the K defenders receiving the highest average attention from the attacker (averaged across all heads for the attacker-defender edge).

- Random-K - Masking a random selection of K defenders.
- Bottom-K - Masking the K defenders receiving the lowest average attention from the attacker.

We conduct this analysis for K = 1, 2, and 3, in Figure 5.2.

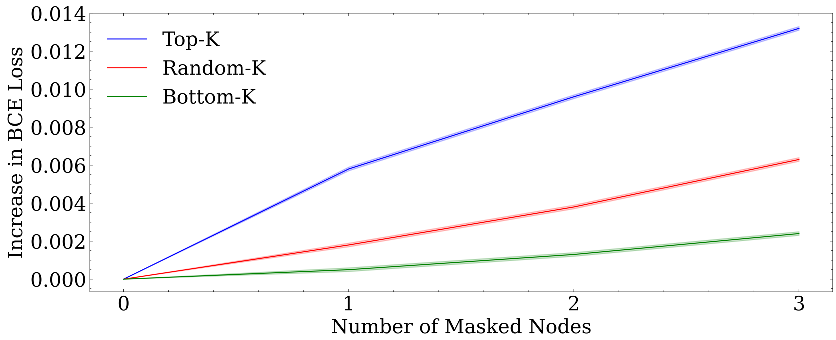

FIGURE 5.2: Fidelity evaluation of the *GAPP* model, illustrating the increase in BCE loss when masking various numbers of nodes for each masking strategy. Shaded 95% confidence boundaries are included, but are minimal.

The results show that masking the Top-K defenders leads to the most significant increase in BCE Loss, with a 5.8×10^{-3} increase for K=1 compared to 1.8×10^{-3} for Random-K and 5.0×10^{-4} for Bottom-K. This trend continues for K=2 and K=3, demonstrating that high-attention defenders contribute significantly to the model's predictions, while low-attention defenders have minimal impact.

5.5.5 Experiment 3: Testing Model Faithfulness

We perform additional tests on the model's attention mechanism by testing *GAPP*'s faithfulness, which assesses how well the attention mechanism aligns with the model's reasoning process (Shin et al., 2025; Liu et al., 2022c). Faithfulness is typically tested by manipulating inputs and observing changes in predictions. We test faithfulness by setting a defender's attention to 0 and measuring the change in an attacker's reception probability. Linear regression shows that a 0.1-unit increase in defender attention corresponds to a ∼6.9% higher relative percentage change of an attacker's reception probability when the defender's attention is removed (coefficient = 69.3, $p < 0.01$). Figure 5.3 shows a scatter plot from a stratified sample illustrating this relationship.

This finding suggests that the model's attention mechanism is closely aligned with changes in its predictions. This alignment further supports the idea that the attention scores capture key features influencing the model's pass reception decisions.

5.5.6 Experiment 4: Impact of Defender Performance on Defensive Actions

In this experiment, we evaluate the relationship between the DP metric and the probability of a defender performing a defensive action. Using logistic regression, we model

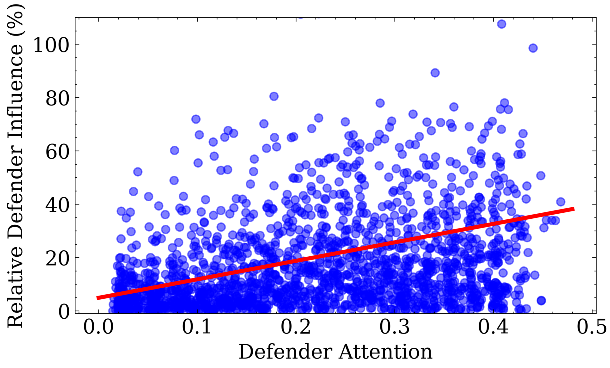

FIGURE 5.3: Assessing *GAPP* faithfulness using the correlation between defender attention and relative defender influence.

the likelihood of a defender performing a defensive action within the following three on-ball events given their relative performance. Relative performance is calculated as the difference between a defender's DP and their team's average DP at the same event.

The analysis shows a significant positive relationship (coefficient = 0.762, $p < 0.01$), where a 0.1-unit increase in relative performance corresponds to a 7.9% increase in the probability of performing an on-ball action. This suggests that higher relative performance improves the likelihood of timely defensive actions. These results may vary for specific teams and opponents. The findings also indicate that defenders actively disrupting attacker receptions are more likely to perform an action, which may suggest that proximity is influencing their impact. The standard deviation of relative performance is $\sim$0.1 (mean $\approx$ 0) in our data, showing the potential for players to improve their likelihood of defensive action with improved off-ball positioning.

5.6 Model Applications

In this section, we apply the *GAPP* model to 2023/24 EPL data, analysing its performance and demonstrating the utility of the DP metric across various contexts. We explore the model's accuracy across different pitch zones, visualise the novel DI and DP metrics, and use DP to rank individual players and defensive partnerships. Furthermore, we analyse how the DP metric changes under varying in-game conditions, including ball location, opponents, player roles, and match time. Further details on the compute resources used to extract the DI and DP metrics across the entire dataset are given in Appendix C.1.1.

5.6.1 Model Accuracy Across Pitch Zones

We analyse how context affects *GAPP*'s prediction accuracy. Specifically, for our EPL dataset, we test the impact of ball location on model predictive accuracy in Figure 5.4.

The model predicts passes more accurately in midfield, suggesting that the structure of play in these areas makes actions easier to predict. Performance declines near the

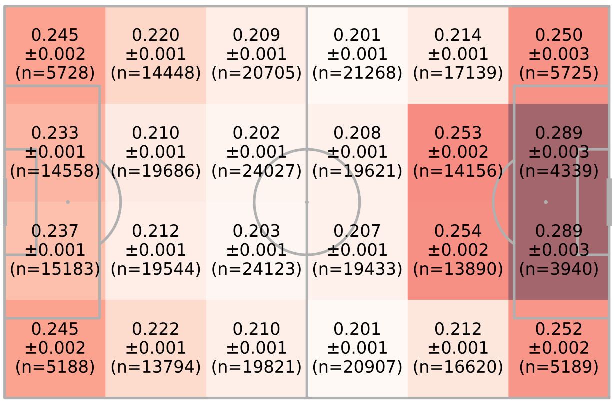

FIGURE 5.4: BCE Loss of the *GAPP* model across our dataset for varying ball zones. Teams attack from left to right. The ± values represent 95% confidence intervals.

goal, where player positions are more congested, matching analysis in Chapter 3 that build-up is more predictable than attacking phases.

5.6.2 Visualising Off-Ball Defensive Contribution

We present visualisations to analyse the DI and DP metrics. These visuals translate our deep learning results and novel defensive metrics into an explainable framework for scenario analysis. Figure 5.5 shows a real-world example, highlighting Rico Lewis's influence (DI) on opposing attackers as predicted by the *GAPP* model.

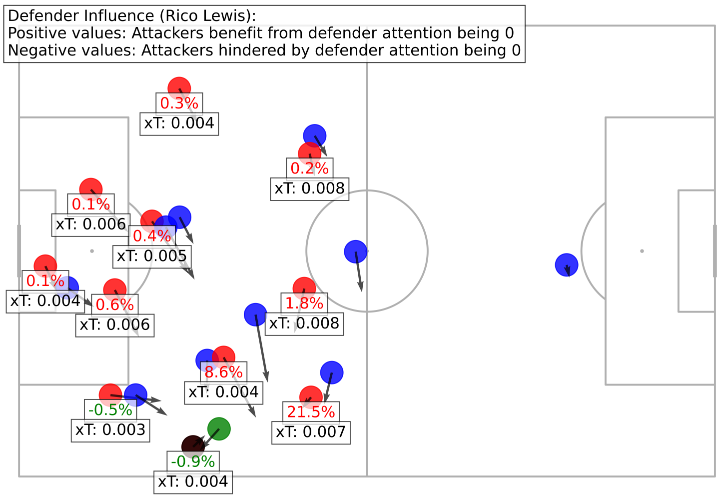

FIGURE 5.5: Example visualisation showing Rico Lewis's (green defender) DI on red attackers, with xT values representing their threat. The red team attacks left to right.

This figure shows Rico Lewis's highest DI on attackers whose passing lanes are being blocked, likely increasing the chance of a backwards pass. Rico Lewis shows minimal influence for attackers positioned further away, where small values may reflect model noise. This visualisation illustrates how *GAPP* captures defensive spatial dynamics. Figure 5.6 further visualises the DP metric for each defender in a specific scenario.

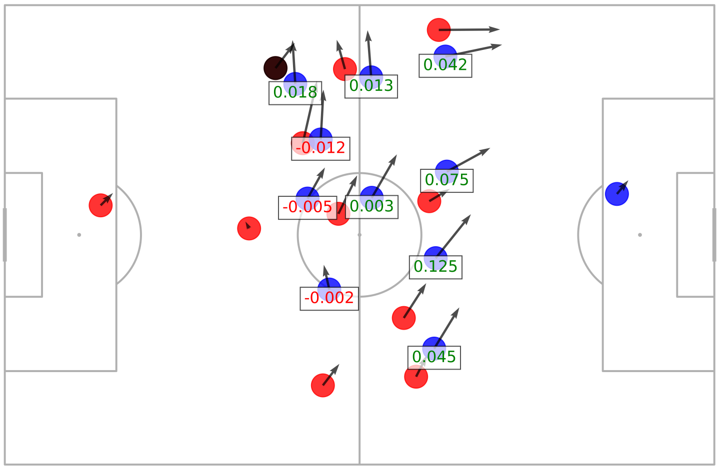

FIGURE 5.6: An example visualisation showing the defender performance (DP) of each defender (blue). The shaded attacker is the ball carrier. Red team attacks left to right.

The DP metric shows that defenders marking high-threat attackers tend to have higher scores, reflecting the intuition that proximity to dangerous attackers is a vital part of defending. These visualisations can analyse full matches or focus on key moments, offering a tool to coaches and teams for evaluating defensive performance.

5.6.3 Rating Players Using Defender Performance

We compared the league's top centre backs using their total DP metric for the season, analysing 99 centre backs from our EPL dataset. To address potential biases, such as each team's number of defensive events or varying match counts in our dataset (which excludes some EPL games), we normalise the DP metric by the number of defensive events by the defender's team in the dataset. The results are shown in Table 5.2.

TABLE 5.2: Ranking EPL centre-backs by defender performance (DP). IR and TR are rankings among EPL centre-backs for most interceptions and tackles, respectively. Each team's 2023/24 EPL position is also shown to indicate overall team performance.

Defender	Team	DP	Rank	IR	TR	Team Pos. (/20)
F. Schär	Newcastle	0.0397	1	11	26	7th
T. Gomes	Wolves	0.0395	2	27	21	14th
J. Andersen	C. Palace	0.0394	3	7	6	10th
M. Kilman	Wolves	0.0393	4	8	17	14th
J. Tarkowski	Everton	0.0387	5	2	4	15th
I. Zabarnyi	Bournemouth	0.0386	6	17	5	12th
J. Branthwaite	Everton	0.0384	7	3	1	15th

Interestingly, top defenders based on the DP metric mostly come from EPL clubs with around average league performance, likely because DP is valued more when influencing high-threat attackers in dangerous areas. Top teams defend less in these areas, leading to lower DP rankings. Similar trends appear in metrics like tackles and interceptions, which are often led by players from non-top teams. Among the 99 centre-backs in our dataset, those with high DP scores also typically rank highly in tackles and interceptions. Analysing defensive partnerships also provides valuable insights (Beal et al., 2020b), as it allows managers to assess how well players work together. Table 5.3 ranks top centre back pairings by combined DP (i.e., the sum of their individual DPs), which is calculated only when both players are on the pitch together as the sole centre backs. The DP is normalised by the team's total defensive events in our EPL dataset.

TABLE 5.3: Ranking EPL centre back partnerships based on combined Defender Performance (DP). GC is the team's goals conceded in the 2023/24 EPL season.

Centre Back Partnership	Team	DP	Team Pos. (/20)	GC
C. Romero M. van de Ven	Tottenham	0.0496	5th	61
J. Tarkowski J. Branthwaite	Everton	0.0444	15th	51
F. Schär S. Botman	Newcastle	0.0431	7th	62
Gabriel W. Saliba	Arsenal	0.0420	2nd	29

These results highlight defensive partnerships that coordinate well off the ball. We also include each partnership's team and average goals conceded, noting that for all teams, the number of goals conceded is below the league average (62.3).

5.6.4 Performance Across Various Opponents

Assessing DP against different attacker types is also valuable for pre-game preparation. We categorise attackers by their teams' final 2023/24 EPL positions: top teams with higher quality players typically dominate possession, while lower-ranked teams usually play deeper and counterattack. Table 5.4 compares the DP of top-rated centre backs against each opposition type.

TABLE 5.4: Ranking EPL defenders and how well they perform (DP) against each type of opposition. Opposition type is separated based on the final team league ranking.

Defender	DP	1-5	6-10	11-15	16-20
F. Schär	0.0397	0.0436	0.0304	**0.0444**	0.0347
T. Gomes	0.0395	0.0433	0.0309	0.0348	**0.0490**
J. Andersen	0.0394	**0.0445**	0.0350	0.0359	0.0415
M. Kilman	0.0393	0.0398	0.0356	0.0389	**0.0431**
J. Tarkowski	0.0387	**0.0439**	0.0358	0.0380	0.0364
I. Zabarnyi	0.0386	**0.0431**	0.0418	0.0341	0.0319
J. Branthwaite	0.0384	**0.0441**	0.0374	0.0362	0.0336

Most defenders record their highest DP against the top 1-5 teams, as DP tends to be higher when the ball is in high-threat areas. However, players like Max Kilman had the highest off-ball performance against lower-ranked teams. This analysis can help coaches assess player strengths and weaknesses. We extend this by evaluating Fabian Schär's influence on various centre forwards by examining his mean DI when he has the highest attention score for an attacker out of all defenders. Table 5.5 lists the top and bottom three centre forwards by mean DI against Schär, along with their total xT/90 (xT per 90 minutes), computed using the centre forward's locations for all their on-ball attacking events in matches against Fabian Schär.

TABLE 5.5: Top 3 and bottom 3 EPL centre forwards ranked by F. Schär's mean DI in limiting their ball receptions, based on at least 200 events where Schär has the highest attention.

Attacker	**Team**	**Mean DI**	**Minutes**	**xT/90**
		Top 3		
E. Adebayo	Luton	**0.0104**	166	0.48
R. Muniz	Fulham	**0.0072**	114	0.40
N. Jackson	Chelsea	**0.0068**	157	0.45
		Bottom 3		
D. Solanke	Bournemouth	**0.0057**	180	0.56
C. Archer	Sheffield United	**0.0054**	156	0.78
E. Haaland	Man City	**0.0042**	90	0.49

Interestingly, Fabian Schär recorded his lowest DI score against Erling Haaland, the league's top scorer during the 2023/24 EPL season. This analysis provides valuable insight into the specific attackers whom Schär demonstrated the greatest effectiveness in limiting ball reception.

5.6.5 Exploring Defender Performance

In this section, we analyse the DP metric of players under varying contexts. Firstly, we compare the DP of players when the ball is in different locations on the pitch, as shown in Figure 5.7.

The results show that the average DP is higher in dangerous pitch areas because attackers pose greater threats, which gives defenders with the same influence a higher DP. Thus, teams that defend frequently in high-risk areas tend to have higher DP scores. However, the link between DI and defensive actions (Section 5.5.6) indicates that DP increases in situations where the defending team is more likely to win the ball, giving insight into team defence beyond just the attacking threat. Interestingly, DP is notably better on the left wing compared to the right. This bias is similar to findings in (Robberechts et al.,

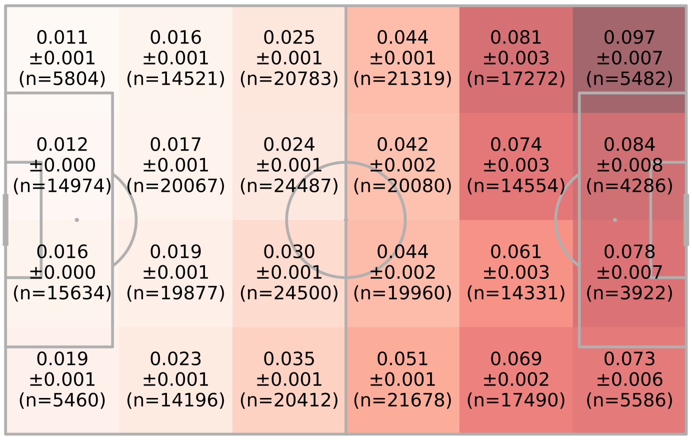

FIGURE 5.7: Comparison of the average DP for each zone the ball occupies on the pitch. Teams attack from left to right. The ± values represent 95% confidence intervals.

2023), where players had higher on-ball decision ratings on the right wing for 2021/22 EPL data. Further analysis could explore if this trend holds for other seasons or leagues.

When comparing the average DP across different team roles, goalkeepers had the lowest mean DP (0.008), while wide defenders (0.053) and centre midfielders (0.045) had the highest. This aligns with trends for tackles in the EPL, where these positions typically make the most tackles. We also analyse the average DP for teams over different match times for all games in our dataset, comparing the league average with Newcastle and Bournemouth in Figure 5.8.

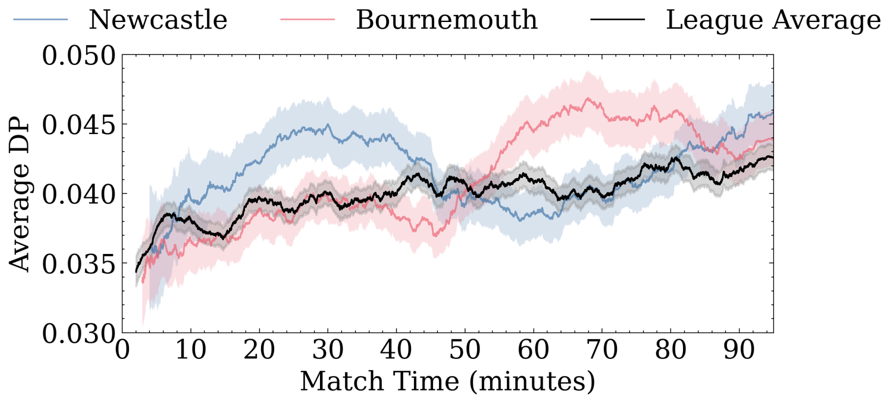

FIGURE 5.8: Comparison of the mean DP in teams (rolling average) for varying match times for all games in our dataset, shown for two example teams and the overall league average, with shaded 95% confidence boundaries.

These results show that Newcastle's mean DP is above the league average in the first half, while Bournemouth's is higher in the second. T-tests found a statistically significant difference between the rolling averages of Bournemouth and Newcastle ($p<0.01$), indicating that they vary in their defensive patterns over time. This analysis helps clubs identify periods of weaker performance or increased defending for teams or players.

5.7 Discussion

Our framework introduces novel metrics to evaluate off-ball defensive contributions in football. While most research focuses on events, top teams use their positions to prevent events from happening, a task that is harder to measure. Merhej et al. (2021) were the first to analyse on-ball defensive actions by predicting what was prevented. Our approach quantifies individual off-ball defensive performance for the first time, which players spend the vast majority of the match. Future work could extend the model to assess off-ball attacking contributions, such as runs or creating space, by analysing attacker attention and its influence on pass recipient probabilities. While *GAPP* uses velocity and acceleration features to represent player intentions and direction of play, future work could explore whether using sequences of past events as input enhances the model's understanding of temporal context further.

The *GAPP* model could also cluster similar play situations for players. Recent research compares past team situations (Stöckl et al., 2021), focusing on entire team structures. Using *GAPP*'s node representations from the GAT layer outputs, we could generate latent representations of a player and their surroundings, enabling clustering of similar scenarios a player has faced, improving the efficiency of tactical analysis for clubs.

In future work, we plan to investigate alternative graph construction methods beyond the fully connected approach, such as proximity-based graphs, and conduct a computational analysis to compare the efficiency and effectiveness of these different graph structures for model performance and understanding true influence between players. Additionally, we plan to perform ablation studies to systematically evaluate the impact of various input features, hyperparameters, and architectural choices on model performance.

Football is a suitable testbed for our model as it is a data-rich, real-world environment, where tracking data provides many data points across games. In future work, we plan to test our approach across more leagues to capture a broader football context, as well as in other real-world MAS, such as security scenarios or other team sports. For instance, in security games, the spatial positioning of defenders can influence attack locations, drawing similarities to the dynamics observed in football. Adapting the model to new domains would require thorough validation and potential adjustments, such as modifying the number of attention heads, layers or input data to reflect differences in relationships between agents. Additionally, the target variable in the *GAPP* model may need to be tailored to each domain, for example, by predicting the probability of an attack on each target in security settings.

While *GAPP* achieves state-of-the-art performance in football pass receiver prediction, its key value lies in interpreting off-ball defensive contributions. This interpretability could support post-match analysis, highlight team weaknesses, and introduce new metrics for player evaluation and recruitment. The model's explainable insights could make a substantial real-world impact by bridging the gap between domain expert football coaches and advanced ML, thereby fostering trust and usability in data-driven decision-making.

5.8 Summary

This chapter introduces *GAPP*, a novel GAT-based model for predicting pass reception probabilities in football. The model achieves a $\sim$6.4% $\pm$ 1.5% reduction in BCE loss compared to the best performing baseline while offering insights into off-ball defensive contributions, a critical yet underexplored area in football analytics. Derived from the *GAPP* attention mechanism, we introduce novel DI and DP metrics, providing new ways to evaluate off-ball defensive performance. A case study on the EPL highlights *GAPP*'s ability to assess players and defender-attacker dynamics. We also show how explainable visualisations can improve trust in data-driven decision-making in football. Future work will explore the application of *GAPP* to other dynamic MAS, such as security and team sports.

In the next chapter, we apply focus to team formation (TF). More specifically, we present a novel sequential TF model that considers the risks of long-term agent unavailability and connects pre-event team selection with in-event team adaptation. We demonstrate its application in football to inform pre-match team selection and in-game substitution decisions, improving long-term performance by reasoning over player welfare. This TF framework utilises the *Agent Imputer* model from Chapter 3 to estimate tracking data and derive player workload from the estimated data, making the approach accessible to a wider range of clubs.

Chapter 6

Optimising Short- and Long-Term Team Selection

In this chapter, we present a sequential decision-making algorithm for optimising team formation (TF) and agent replacement in dynamic, adversarial environments, applying this problem to football. Effective team selection and adjustment are essential for maximising rewards, such as match points, while minimising risks, such as agent fatigue and failure, yet existing models often overlook real-time adjustments to the team and the importance of player welfare. Our model integrates ex-ante (pre-game) decision-making with real-time (in-game) substitutions, both modelled as stochastic Markov decision processes (MDPs) that incorporate predictive models of event outcomes and agent fault risk. We use Monte Carlo tree search (MCTS), guided by human expert policies, to approximate optimal actions across these decision stages. Validated on two seasons of English Premier League (EPL) data, our algorithm improves season-average points by approximately 1% compared with baseline models, while reducing squad injuries by 3%, first-team injuries by 15%, and wages inefficiently spent on injured players by 3%. By explicitly reasoning over player skill, fatigue, and injury risk, our approach improves long-term team performance and player availability, showing potential to inform managerial decision-making and presenting broader applicability to other dynamic multi-agent domains.

6.1 Introduction

In many real-world domains, forming and managing teams in MAS is critical for achieving collective objectives such as maximising rewards or minimising risks. These domains include disaster response, where teams of agents are assigned to newly emerging tasks to maximise overall task completion rates (Ramchurn et al., 2010; Wu and Ramchurn, 2020; Capezzuto et al., 2020), as well as high-stakes adversarial environments such as port security and wildlife protection, where teams of agents are deployed to monitor and protect vulnerable areas with the aim of minimising the risk of successful adversary attacks (Shieh et al., 2012; Wang et al., 2019b). In such settings, effective team management often requires not only selecting the best team of agents for the current situation by matching agent skillsets to the task, but also dynamically adapting the team by selecting or replacing agents in response to changes in the environment, which may include agent fatigue, injury, or adversary decisions.

While many existing TF approaches for MAS address spatio-temporal constraints, including travel time and task completion deadlines (Ramchurn et al., 2010; Capezzuto et al., 2020), and skill-based factors such as agent compatibility to tasks (Gaston and DesJardins, 2005), they often overlook the challenge of managing agent fatigue and injury risk. For instance, in real-world domains, emergency responders have their shifts managed to prevent exhaustion and injury, and robot swarms are scheduled to avoid battery depletion or damage. Some frameworks are designed to be robust to faults by selecting teams that can tolerate a certain number of agent failures and ensure task completion (Okimoto et al., 2015). However, these approaches do not fully address the importance of proactive team management strategies, which can improve the long-term welfare of human or robotic agents. Anticipating and mitigating fatigue and injury not only enhances agents' well-being but also maintains team effectiveness by ensuring availability for future tasks. Therefore, effective team selection and adaptation are needed to balance the immediate and long-term objectives, to maximise outcomes and minimise risk in dynamic, complex settings.

A key aspect of real-world multi-agent team management is the temporal structure of decision-making. In many domains, teams must operate over extended periods and across multiple events or locations. For example, in disaster response, teams are repeatedly deployed to different sites or incidents, requiring managers to allocate resources and schedule shifts over time to maintain operational effectiveness. Similarly, in football, teams compete in a sequence of league or cup matches, where player selection and management decisions must account for both the demands of the current game and the cumulative effects across a season.

Within this temporal context, there are critical decision points at which managers can influence both short-term and long-term outcomes. Pre-event decisions involve selecting the initial team configuration or lineup, considering factors such as agent fatigue, injury risk, and anticipated challenges. In contrast, in-event decisions involve dynamic adaptation, such as substitutions or reassignments, in response to unexpected events or environmental changes. By explicitly modelling and optimising both pre-game and in-game decision-making, it is possible to balance immediate performance with safeguarding agent welfare and sustaining team effectiveness over time.

Football provides a suitable real-world testbed for these challenges, as teams of human agents must handle congested match schedules and frequent player injuries over the course of a typical eight-month league season. These factors, combined with the need to respond to evolving opponent strategies, make effective team selection and in-game substitutions essential for maintaining long-term performance and improving player welfare. The abundance of detailed data on player performance, injuries, and match outcomes enables football to be effectively modelled as a real-world team selection problem. Moreover, team performance and player injuries have a pivotal impact on the financial stability of football clubs. For instance, the 2022/23 season saw the 'Big Five' European football leagues incur €705 million in injury-related costs.[1] Injury rates during football matches can range from 10 to 60 incidents per 1000 playing hours (Owoeye et al., 2020), and research shows that squad injuries significantly impact both league and cup outcomes (Hägglund et al., 2013). League performance also influences club social standing and financial health, as teams receive prize money based on their final ranking. These factors highlight the importance of prioritising player welfare, not only for ethical reasons but also to enhance the strategic and financial performance of teams.

Recent advances in sports science, wearable technology, and tracking systems have made it possible to collect detailed data on player workload, fatigue, and injury risk (Hulin et al., 2016). Furthermore, developments in sports analytics have produced models that guide team selection and substitutions, taking into account player skills, teamwork, and tactical requirements (Beal et al., 2020b,a). However, research has yet to fully address the integration of long-term planning with dynamic in-game adaptation to manage player fatigue and performance, leaving this challenge underexplored. Most existing approaches also focus particularly on pre-game team selection and overlook the importance of real-time adjustments during matches in response to changing opponent strategies, such as tactical changes or substitutions.

This chapter addresses these gaps by proposing a sequential decision-making algorithm that integrates pre-game and in-game team management. The pre-game team management model determines the starting lineup by considering factors such as player fatigue, injury risk, and expected opponent strategies. This decision process is modelled as a stochastic finite-horizon MDP and optimised using MCTS. The MCTS algorithm uses human-expert policies, which are derived from historical managerial decisions, to guide the model towards more realistic and effective team selections. Expanding on the pre-game team management model, our approach enables in-game player substitutions to adapt to changes in match conditions, such as shifts in opponent tactics or a player's physical state. The in-game decision process is also modelled as a stochastic finite-horizon MDP, where MCTS is used to approximate optimal substitution strategies in response to real-time developments. The two team management models are coupled

[1] https://www.howdengroupholdings.com/news/howden-2022-23-mens-european-football-injury-index

in the algorithm, with the team chosen by the pre-game model serving as the initial starting team for the in-game model.

We evaluate our algorithm using a comprehensive dataset of 760 EPL matches, extracting information on player skills, injury probabilities, team strategies, and match transition probabilities. This enables our model to accurately simulate and adapt to real-world match conditions. Our results demonstrate that this unified TF approach not only enhances long-term team performance but also reduces player injury rates and minimises financial losses associated with injured players.

Thus, this chapter presents the following novel contributions:

- We present a unified approach for both pre-game team selection and in-game player replacement, modelling these as two connected MDPs, enabling dynamic adaptation to evolving adversarial environments.
- We develop a multi-stage MCTS algorithm that sequentially optimises pre-game lineups and in-game substitutions, approximating optimal decision policies to maximise long-term team performance by reasoning over agent welfare.
- We demonstrate how player abilities, injury risks, replacement strategies and match transition probabilities can be derived from real-world data to simulate and optimise future scenarios.
- We empirically evaluate our approach using real-world football data, benchmarking against a real-world manager model that simulates human decision-making with $\sim$80% accuracy. Our model not only achieves performance gains on human-level performance but also reduces first-team player injuries by up to $\sim$15% and decreases inefficient spending on injured players by approximately 3%.

By explicitly modelling agent reliability and risk, such as player fatigue and injury, our work advances the state-of-the-art in dynamic TF for real-world, adversarial settings like football. Central to our approach is a two-step decision-making process that integrates both pre-game team selection and in-game adaptation, allowing managers to respond proactively to evolving match conditions. Furthermore, by embedding these decisions within the context of repeated interactions, across a season or series of events, our method enables teams to optimise long-term performance while considering agent welfare. Although football serves as the primary testbed for our study, the proposed framework is broadly applicable to other high-risk, team-based domains, such as emergency response and other team sports. We aim to validate its effectiveness in these settings in future work.

The remainder of this chapter is structured as follows. Section 6.2 defines the sequential team selection model. Section 6.3 models team selection as an MDP, and Section 6.4

introduces the injury probability model. Section 6.5 details the team dominance and match transition model. Section 6.6 presents the application of MCTS for optimising team selections and substitutions. Section 6.7 empirically evaluates our approach, and Section 6.8 discusses the findings. Finally, Section 6.9 summarises the work.

6.2 Football Team Selection Model

We derive the fundamentals of our sequential team selection model from observed properties of real-world football team selection. We present a framework diagram of the decision-making process in Figure 6.1.

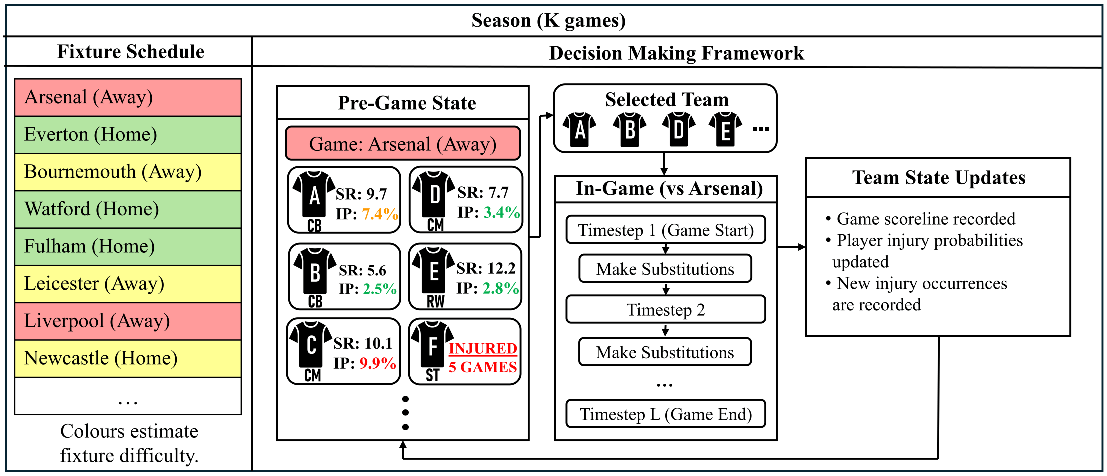

FIGURE 6.1: Decision-making framework across a season for football clubs. This framework illustrates the interconnected pre-game and in-game decision phases across a season (right side), emphasising how clubs navigate a condensed fixture list of varying difficulties (left side) while managing stochastic player injuries and injury risks. (SR = skill rating, IP = injury probability).

A high-level overview of the decision-making process outlined in the figure is given as follows:

1. **Pre-Game Phase: Starting Team Selection**
   (a) **Current Team State:** Before a game, our model evaluates the squad, analysing player injury risks and skill ratings (see Sections 6.4 and 6.5), as well as the upcoming opponent and future fixtures.
   (b) **Manager Decision-Making:** The model identifies the starting team that maximises long-term performance, balancing immediate rewards (e.g., winning the current game) and future risks (e.g., player injuries). This process is aided by a player-based match prediction model.
   (c) **Team Selection:** The manager selects the suggested starting team based on the output of the model.

2. **In-Game Phase: Substitution Decisions**
   (a) **Dynamic Game State:** During the game, the model monitors real-time factors such as player fatigue, injuries, and in-game performance metrics.
   (b) **Substitution Decision-Making:** The manager dynamically selects substitutions to optimise in-game performance, considering the current match scoreline and minute, player fatigue, player performance and the opponent's tactics.
3. **Team State Updates:** Following the game, player injury probabilities are updated, and any new injuries are recorded, reflecting changes in the team's condition. The scoreline of the game is also recorded. This new team state is applied to the next game in the sequence (fixture schedule).

The team selection problem involves many factors, such as team skill, game importance, and player injury risk. We divide the team selection problem into pre-game and in-game phases to address their distinct objectives: strategic, long-term planning for the pre-game phase and dynamic, real-time adjustments during the in-game phase. The methodology used to optimise pre-game and in-game decision-making is given in Section 6.6. We formally define the sequential team selection model in the following subsections.

6.2.1 Basic Definitions

We define a sequential set of K games in a season, indexed by $k \in \{1, ..., K\}$. The set of games is denoted as $\mathbf{G} = \{G_1, ..., G_K\}$ where G_k is the k'th game of the season for a given team. Each game $G \in \mathbf{G}$ involves two adversarial teams, denoted as $C = \{C_\alpha, C_\beta\}$, where C_α is the team our model aims to optimise and C_β is the opposing team (adversary). Each team consists of a squad of players. Specifically, the squad for team C_α is defined as $C_\alpha = \{\alpha_1, ..., \alpha_{N_\alpha}\}$ where N_α is the total number of players in the squad for team C_α. For each game G_k, the team selects a subset of players to start the game, while the remaining players are designated as reserves. We define the set of starting players as $C_\alpha^+ \subset C_\alpha$ and the set of reserve players as $C_\alpha^- \subset C_\alpha$ such that $C_\alpha^+ \cup C_\alpha^- = C_\alpha$ and $C_\alpha^+ \cap C_\alpha^-$ is the empty set. For our football application, $|C_\alpha| = 18$ and $|C_\alpha^+| = 11$. Each game G_k is further divided into a sequence of L discrete timesteps, indexed by $\tau \in \{1, ..., L\}$ where each game timestep τ corresponds to a uniform time interval and L is the number of game timesteps (for this work, we model a 90-minute football game as 10-minute intervals). Each timestep represents a game state, and player substitutions can occur at any game timestep τ during game G_k.

Every game G_k results in an outcome Ω_k, which includes the final match result (i.e., the score) and updated player states, such as skill, fatigue and injury status. Additionally, within game G_k, each game timestep τ has an associated outcome $\Omega_{k,\tau}$ representing the progress towards the final game outcome. This includes information such as the current match score and the injury status of players at that specific timestep. In the next subsection, we formally define the characteristics of players within a team.

6.2.2 Player Characteristics

Each player $\alpha \in C_\alpha$ is characterised by a set of attributes that influence both individual and team performance. We formalise the set of attributes as a 6-tuple of player characteristics at timestep τ of game k:

$$\zeta^{\alpha}_{k,\tau} = \langle \kappa_\alpha, \psi_\alpha, \psi^{\alpha}_{k,\tau}, \mathcal{I}^{\alpha}_{k,\tau}, \Pr(\mathcal{I}^{\alpha}_{k,\tau}), l^{\alpha}_{k} \rangle$$

Firstly, κ_α represents the role of player α within the team, such as goalkeeper, defender or attacker. Secondly, ψ_α denotes the historical performance of player α across all previous games they participated in, reflecting the player's contribution towards the team scoring and conceding goals in prior matches. Additionally, $\psi^{\alpha}_{k,\tau}$ represents the performance of player α in the current game G up to game timestep τ. Before the game begins, and for reserve players during the game, $\psi^{\alpha}_{k,\tau}$ is initialised to 0. A player's in-game performance only updates when they are actively participating in the game (i.e., when they are part of C^{+}_{α}).

The binary variable $\mathcal{I}^{\alpha}_{k,\tau}$ indicates the injury status of player α at timestep τ of game k. Specifically, $\mathcal{I}^{\alpha}_{k,\tau} = 1$ if the player is injured at game timestep τ, and $\mathcal{I}^{\alpha}_{k,\tau} = 0$ otherwise. An injured player cannot be included in the playing team C^{+}_{α}, (i.e., $\forall \alpha \in C^{+}_{\alpha}, \mathcal{I}^{\alpha}_{k,\tau} = 0$). If a player in C^{+}_{α} becomes injured during the game (i.e., $\mathcal{I}^{\alpha}_{k,\tau} = 1$ at some timestep τ), they must be substituted out at the next timestep. The variable $\Pr(\mathcal{I}^{\alpha}_{k,\tau})$ denotes the probability of player α sustaining an injury at timestep τ. This probability is influenced by several factors, including the player's fatigue level and their injury history (see Section 6.4). Finally, l^{α}_{k} represents the number of future games (starting from game k) that player α will miss due to injury. If a player is not injured, $l^{\alpha}_{k} = 0$. For injured players, l^{α}_{k} is updated after each game. Specifically, if the player was already injured before game G_k, l^{α}_{k} decreases by 1 after the game. If the player sustains an injury during game G_k, the absence duration l^{α}_{k} is sampled from a real-world injury distribution $\mathcal{F}$. The sampled duration is converted into a game count based on the number of games in $\mathbf{G}$ that fall within the injury period (see Section 6.4). For simplicity, we assume that all injury lengths are drawn from the same overall distribution, based on our entire injury dataset, and do not account for differences due to injury type or individual player characteristics.

We review the impact of fatigue on player performance (Rampinini et al., 2009; Dambroz et al., 2022) and introduce a fatigue decay on player ability, where ψ^{*}_{α} is player ability scaled by fatigue. We define the decay parameter as: $\Xi^{\alpha}_{k,\tau} = 1 - \lambda \cdot \frac{\Pr(\mathcal{I}^{\alpha}_{k,\tau})}{\Pr(\mathcal{I}^{\alpha}_{k,\tau})_{\max}}$, where $\Xi^{\alpha}_{k,\tau}$ is the decay parameter, λ is a scaling factor and $\Pr(\mathcal{I}^{\alpha}_{k,\tau})_{\max}$ is the maximum possible injury probability of that player (according to real-world data and our injury model). The decay parameter updates the player's performance value to be $\psi^{*}_{\alpha} = \psi_{\alpha} \cdot \Xi^{\alpha}_{k,\tau}$. We set the fatigue scaling factor, λ, to 0.25. This value signifies that at maximum injury probability, a player's ability is reduced by 25%. A higher λ, such as 0.99, would mean that the ability is nearly completely reduced at maximum fatigue, while a lower value would suggest a more negligible impact. We selected 0.25 as a subjective choice informed by findings in the literature, which highlight the significant effect of fatigue on technical and physical performance in football players (Rampinini et al., 2009; Dambroz et al., 2022).

6.2.3 Team Dominance

Each team has a dominance value, which quantifies its ability to create and concede goal-scoring opportunities. The dominance of team C_{α} is generally denoted as $D_{\alpha} \in \mathbb{R}$. For a specific game G_k, the dominance of the selected starting players C^{+}_{α} is denoted as $D^{\text{pre}}_{\alpha,k}$. This value is computed based on the individual contributions towards goal scoring of the players in C^{+}_{α} across all prior games, scaled by their fatigue at the start of game k, formally defined as $\sum_{\alpha \in C^{+}_{\alpha}} \psi^{*}_{\alpha}$. The pre-game dominance of the opposing team C_{β} is similarly defined as $D^{\text{pre}}_{\beta,k}$. The relationship between a team's pre-game dominance and the probability of achieving a specific final game outcome Ω_k is modelled as:

$$\Pr(\Omega_k) = f^{\text{pre}}_{\Omega}(D^{\text{pre}}_{\alpha,k}, D^{\text{pre}}_{\beta,k}) \tag{6.1}$$

Where f^{pre}_{Ω} is a function that captures the influence of the dominance values of both teams on the game outcome. At the level of individual timesteps, the dominance of team C_{α} in game G_k up to timestep τ is denoted as $D^{\text{in}}_{\alpha,k}(\tau) \in \mathbb{R}$. Similarly, the probability of achieving a specific timestep outcome $\Omega_{k,\tau}$ is modelled as:

$$\Pr(\Omega_{k,\tau}) = f^{\text{in}}_{\Omega}(D^{\text{in}}_{\alpha,k}(\tau), D^{\text{in}}_{\beta,k}(\tau), D^{\text{pre}}_{\alpha,k}, D^{\text{pre}}_{\beta,k}) \tag{6.2}$$

Where $D^{\text{in}}_{\alpha,k}(\tau) = \sum_{\alpha \in C^{+}_{\alpha}} \psi^{\alpha}_{k,\tau}$ and $D^{\text{in}}_{\beta,k}(\tau) = \sum_{\beta \in C^{+}_{\beta}} \psi^{\beta}_{k,\tau}$ represent the in-game dominance values of the two teams up to game timestep τ. In this formulation, the dominance values $D^{\text{pre}}_{\alpha,k}$ and $D^{\text{pre}}_{\beta,k}$ serve as key predictors of game outcomes in the pre-game phase, while the timestep-level dominance values $D^{\text{in}}_{\alpha,k}(\tau)$ and $D^{\text{in}}_{\beta,k}(\tau)$ provide a finer-grained view of the dynamics within each game. The functions f^{pre}_{Ω} and f^{in}_{Ω} model the probabilistic relationships between dominance and outcomes at both the game-level and timestep-level. This offers a foundation towards optimising pre-game team selection and in-game substitutions. While these functions focus on match outcome (e.g., win/draw/loss), the updated injury status of each player at each game timestep is governed by their characteristics in $\zeta^{\alpha}_{k,\tau}$. In the next section, we model the pre-game phase and the in-game phase as two MDPs.

6.3 Modelling Team Selection as an MDP

We model football team selection as a two-phase decision process: pre-game and in-game. Each phase is represented by a separate MDP with distinct objectives, states, and outcomes. The pre-game MDP focuses on selecting the starting lineup before each match. This MDP considers factors such as player availability and long-term performance across the season. Critically, the pre-game MDP incorporates a match prediction model to estimate game outcomes, allowing it to optimise for overall season performance without tracking detailed in-game states. This makes the pre-game MDP self-contained, which we use to optimise pre-game team selections. The in-game MDP uses the starting lineup selected by the pre-game MDP to form the initial state. It then governs substitution decisions during the match, focusing on real-time performance, player fatigue, and tactical adjustments. The in-game MDP operates independently of the pre-game MDP, allowing it to react to game dynamics.

This separation into two MDPs allows each phase to be optimised independently: the pre-game MDP for strategic season-long planning, and the in-game MDP for tactical adaptation during a match. To simulate a full season, we link the two MDPs. The starting lineup (action) from the pre-game MDP initialises the in-game MDP, and the final outcome of the in-game MDP updates the state of the pre-game MDP for the subsequent game.

In both the pre-game and in-game MDPs, we assume the processes are approximately Markovian, which means that the current state contains all relevant information for making optimal decisions. For the pre-game MDP, whose states include player availability, prior player performance, and the remaining games in the season, we expect this representation to be largely sufficient since line-up decisions primarily depend on these factors rather than the full history of past matches. For the in-game MDP, the state encodes the current scoreline and player characteristics such as in-game performance so far, which similarly captures the essential information for substitution decisions, however, short-term momentum or tactical patterns may introduce mild dependencies on prior timesteps. Overall, we assume these processes to be sufficiently Markovian to allow us to tractably optimise decision-making.

MDPs provide a robust framework for capturing how a stochastic environment evolves as a decision-maker interacts with it. This makes them well-suited for modelling team selection decisions, where the decision-maker (e.g., the manager) interacts with the environment (e.g., the game) by dynamically updating the team's selected players. We present the pre-game MDP in the following subsection.

6.3.1 Pre-Game MDP

We formulate our pre-game team selection model as a finite-horizon MDP where a season of football games is defined as a sequence indexed by $k \in \{1, \ldots, K\}$. The MDP is defined as $\mathcal{M}_{\text{pre}} = \langle S_{\text{pre}}, A_{\text{pre}}, P_{\text{pre}}, R_{\text{pre}}, \gamma_{\text{pre}} \rangle$ with a set of states S_{pre} and set of actions A_{pre}. The horizon of the MDP is finite and corresponds to the number of games in the season K. We index the MDP steps by the game index k, so at each step the state and action correspond to a particular game G_k.

At each game G_k, the state is $s_k \in S_{\text{pre}}$, the action is $a_k \in A_{\text{pre}}$, and the next state is s_{k+1}. For $s_k, s_{k+1} \in S_{\text{pre}}$ and $a_k \in A_{\text{pre}}$, the transition function is $P_{\text{pre}} = \Pr(s_{k+1}|s_k, a_k)$, and the reward function is $R_{\text{pre}}(s_k, a_k)$. Finally, we have a discount factor $\gamma_{\text{pre}} \in [0, 1]$, which is set to 1 to focus on the overall season performance. Each MDP state is a tuple consisting of two components: $s_k = \langle k, \{\zeta_k^\alpha \forall \alpha \in C_\alpha\} \rangle$. Here, k is the index of the current game, implicitly defining the set of remaining games in the season. Additionally, the second component, $\zeta_k^\alpha \forall \alpha \in C_\alpha$, represents the set of characteristics ζ_k^α for each player α in the set of players C_α prior to game G_k. These game-level characteristics differ slightly from the timestep-level characteristics defined in Section 6.2.2. In particular, in-game performance $\psi_{k,\tau}^\alpha$ is ignored in the pre-game MDP, and injury status and probability are defined at the game level ($\mathcal{I}_k^\alpha$ and $\Pr(\mathcal{I}_k^\alpha)$) rather than at the timestep level ($\mathcal{I}_{k,\tau}^\alpha$ and $\Pr(\mathcal{I}_{k,\tau}^\alpha)$).

The set of actions A_{pre} is the set of all selectable teams such that all players in the team must be available to play. The action set updates given the player availability encapsulated in the current state. We define this action set as:

$$A_{\text{pre}} = \{C_\alpha^+ \subset C_\alpha \,||C_\alpha^+| = 11, \forall \alpha \in C_\alpha^+ (\mathcal{I}_k^\alpha = 0)\} \tag{6.3}$$

The reward function $R_{\text{pre}}(s_k, a_k)$ assigns a reward to team C_α based on the selected lineup determined by a_k and the current state s_k. This reward represents the expected points from the match, considering the chosen lineup and the opponent's predicted lineup. Details on the calculation of expected points and the methodology for predicting the opponent's lineup are provided in Sections 6.5 and 6.6.3.1, respectively.

Our MDP transition function P_{pre} stochastically models the probabilistic nature of player injuries. For a state s_k and a selected action a_k, a successor state s_{k+1} is sampled using the following process:

1. Sample $\mathcal{I}_k^\alpha \sim \Pr(\mathcal{I}_k^\alpha) \; \forall \; \alpha \in C_\alpha^+$.

2. Update absence times:
   (a) *New Injuries*: Sample injury durations $l_k^\alpha \sim \mathcal{F} \; \forall \; \alpha \in C_\alpha^+ \wedge \mathcal{I}_k^\alpha = 1$.
   (b) *Existing Injuries*: Reduce injury duration $l_k^\alpha = l_k^\alpha - 1 \; \forall \; l_k^\alpha > 0$.
3. Increment game index to $k+1$ and update $\{\zeta_{k+1}^\alpha \forall \alpha \in C_\alpha\}$.

Here, $\Pr(\mathcal{I}_k^\alpha)$ is the injury probability of a player α in game G_k. At each game k, we begin by sampling potential new injuries for every player α according to their injury probability $\Pr(\mathcal{I}_k^\alpha)$. For new injuries ($\mathcal{I}_k^\alpha = 1$), the injury duration l_k^α (in games) is sampled from the injury length distribution $\mathcal{F}$. For players with existing injuries, their absence time is reduced by one as we transition to the next match. After updating injury statuses, we progress to the next game G_{k+1} and update the state variables and player characteristics $\{\zeta_{k+1}^\alpha \forall \alpha \in C_\alpha\}$. Injury probabilities are updated given the injury model introduced in Section 6.4. For the pre-game MDP, prior performance ψ_α is extracted from real-world historical data. All states where the game index $k = K$ are terminal states. Thus, a trajectory from the initial state ($k = 1$) to a terminal state represents an entire football season. In the next subsection, we describe the in-game MDP.

6.3.2 In-Game MDP

We formulate a game G_k as a finite-horizon MDP defined as $\mathcal{M}_{\text{in}} = \langle S_{\text{in}}, A_{\text{in}}, P_{\text{in}}, R_{\text{in}}, \gamma_{\text{in}} \rangle$. The horizon corresponds to the number of timesteps L in the game. We index the MDP steps by $\tau \in \{1, \ldots, L\}$, where each step corresponds to a uniform time interval. At each step τ, the state is $s_\tau \in S_{\text{in}}$ and the action is $a_\tau \in A_{\text{in}}$. The MDP reaches a terminal state when the game ends at timestep $\tau = L$.

Each MDP state $s_\tau \in S_{\text{in}}$ is a tuple $s_\tau = \langle \text{score}_\tau, \zeta_{k,\tau}^C \rangle$ where score_τ is the current match scoreline at timestep τ and $\zeta_{k,\tau}^C$ represents the characteristics of all players on both teams at timestep τ of game k. The action set for a team at timestep τ is the set of all possible player substitutions. We define this as:

$$A_{\text{in}} = \{(\mathcal{O}_\alpha^+, \mathcal{O}_\alpha^-) \; | \mathcal{O}_\alpha^+ \subseteq C_\alpha^+, \mathcal{O}_\alpha^- \subseteq C_\alpha^-, |\mathcal{O}_\alpha^+| = |\mathcal{O}_\alpha^-|\} \tag{6.4}$$

Where $\mathcal{O}_\alpha^+$ is the set of players being replaced and $\mathcal{O}_\alpha^-$ is the set of players replacing the players in $\mathcal{O}_\alpha^+$. Players being substituted off must be from the deployed team (C_α^+), those being substituted on must be from the reserves (C_α^-), and the number of players substituted on and off must be equal. Note that $\mathcal{O}_\alpha^+$ can be the empty set, meaning that no players will be replaced. A terminal state, corresponding to the final game timestep $\tau = L$, signifies the end of the match, at which point no further substitutions can be made. The decision-maker aims to identify the optimal set of replacements to maximise the reward by considering the characteristics of both playing and backup players.

The transition function $P_{\text{in}} = \Pr(s_{\tau+1}|s_\tau, a_\tau)$ represents the probability distribution over possible next states given the current state and action. The transition dynamics are influenced by the decision-maker's action a_τ and the opponent's selected substitutions, which are treated as part of the environment's dynamics and modelled in Section 6.6.3.2. For transitions involving the updated injury status of players, we utilise the injury model described in Section 6.4. For this work, we update the players' in-game performance in our in-game model by sampling from a distribution of their in-game performance in historical matches. Additionally, the scoreline, $\text{score}_{\tau+1}$, is sampled using the outcome transition model f_Ω^{in} we defined previously. We discuss this outcome transition model in more depth in Section 6.5. At a terminal state, no state transition can occur, and the MDP terminates.

The reward function $R_{\text{in}}(s_\tau, a_\tau)$ computes the immediate utility for team C_α upon reaching state s_τ. Since s_τ encodes the current match scoreline, we define the reward as the expected points for the team from the match, given both the current scoreline and the expected scoring rates of each team between the current and next game timestep, divided by the total number of timesteps to normalise the reward to a per-timestep value. As detailed in Section 6.5, the expected points for a match outcome, $\mathbb{E}[\text{Points} \mid s_\tau, a_\tau]$, are calculated by modelling the probability of each final scoreline using a Poisson distribution. To account for the impact of player injuries, we adjust this reward by subtracting a penalty that reflects the expected points lost in future matches due to injuries that may occur at the next state transition. Formally, we define the in-game reward as:

$$R_{\text{in}}(s_\tau, a_\tau) = \frac{\mathbb{E}[\text{Points} \mid s_\tau, a_\tau]}{L} - \mathbb{E}[\text{Points Lost} \mid s_\tau, a_\tau] \tag{6.5}$$

Where $\mathbb{E}[\text{Points} \mid s_\tau, a_\tau]$ is the expected number of points (3 for a win, 1 for a draw, 0 for a loss) based on the scoreline and expected scoring rates in state s_τ, and $\mathbb{E}[\text{Points Lost} \mid s_\tau, a_\tau]$ is the expected number of points the team may lose in future matches due to injuries that may be incurred at the next game timestep. The detailed computation of $\mathbb{E}[\text{Points Lost} \mid s_\tau, a_\tau]$ is provided in Section 6.5. At a terminal state (game timestep $\tau = L$), the reward is simply the points from the game, calculated as the final scoreline divided by L. In the next section, we introduce the injury risk model.

6.4 Injury Risk Model

In this section, we present an injury risk model for football. The features used to train our model are curated from a two-season EPL on-ball events dataset and a historic injury dataset (see data sources in Section 6.7.1). We use XGBoost (Chen and Guestrin, 2016) due to its strong performance on tabular data and train the model using binary cross-entropy loss. The feature set is listed in Table 6.1. Prior work has identified relationships between these features and injury risk for athletes in sports (Hulin et al., 2016; Orchard and Powell, 2003; Kucera et al., 2005).

TABLE 6.1: Injury risk model features. Rolling av. refers to the weighted rolling average in past games (decay factor of 0.8).

Player Physical Features	Player Workload and Fatigue	Wider Environmental Features
Number of past injuries	Acute workload (dist. covered in past week)	Player's current team
Age	Chronic workload (dist. covered in past month)	Opposing team
Career total days injured	Distance Covered (rolling av.)	Travel distance to game (km)
Career total games missed	Total dribbles (rolling av.)	Opposition tackles (rolling av.)
Most recent injury length (days)		Opposition fouls (rolling av.)
Most recent injury length (days)		Game day temperature (C)
Total days injured with modal injury		Game day precipitation (mm)
Number of modal injury occurrences		
Duration of longest injury (days)		
Days since last injury		
Number of injuries in the past year		

We assess each feature's contribution to our model using Shapley additive explanations (SHAP) values (Lundberg and Lee, 2017) - an application of cooperative game theory to fairly distribute feature contributions to predictions in an machine learning (ML) model. The most contributory feature was acute workload with a mean SHAP value of 0.0033, followed by the number of past injuries (0.0029), career total days injured (0.0023) and distance covered (0.0022). Features such as team form, days since the previous game, and other physical features were also curated. However, these features had minimal performance impact, with some already incorporated in other interrelated features. Measuring feature contributions to explain injury risks enhances the model's interpretability and accessibility by providing coaches and players with insights into the underlying factors that drive a high injury risk or influence suggested team selections. This added transparency supports more informed and nuanced decision-making. A full analysis of all feature contributions to the injury risk model is provided in Appendix D.4. We use our feature set as input to our model to predict player injury probabilities for a game k, denoted $\Pr(\mathcal{I}_k^\alpha)$. The injury probability for a given timestep τ within game k, denoted $\Pr(\mathcal{I}_{k,\tau}^\alpha)$, is adjusted based on the proportion of time elapsed between timesteps. Specifically, the model assumes that the probability $\Pr(\mathcal{I}_k^\alpha)$ represents the likelihood of a player sustaining an injury if they were to play the entire game. This probability is scaled uniformly during in-game transitions to reflect the shorter duration between timesteps. A player's injury risk is updated using the model at each timestep, given the additional workload since the previous timestep.

As outlined in Chapters 2 and 3, while many professional teams now use wearable technologies such as Catapult[2] for load monitoring during training and matches, access to comprehensive physical data across all competitions and contexts remains challenging as granular tracking data is often unavailable for historical matches or lower-tier leagues. To ensure our team selection model is broadly applicable and accessible, we estimate player physical metrics using widely available on-ball event data. This approach allows

[2] https://www.catapult.com/

us to derive fatigue-related features in settings where advanced tracking or wearable data are inaccessible. Specifically, our *Agent Imputer* model (introduced in Chapter 3) uses on-ball data to predict player locations for every match event, from which we estimate physical outputs such as total distance covered. When simulating fatigue-related features in our model, we extrapolate the mean distance covered and dribbles per game for individual players from real-world data to the minutes they have played in the MDP. Injury-related features are also updated when new injuries occur during our simulated seasons. Environmental features update deterministically using our real-world data.

Finally, we predict the injury length distribution $\mathcal{F}$ by applying Gaussian kernel density estimation to our entire real-world injury dataset (described in Section 6.7). For simplicity, all injury durations in our model are sampled from this single distribution, rather than accounting for differences between injury types. When an injury occurs during MDP simulation, the injury duration l_k^α for player α is updated with a sample from $\mathcal{F}$, converted from days to games. In the next section, we introduce the model that underpins team dominance and match scoreline transitions.

6.5 Team Dominance and Match Transitions

In football, we define team domination at a specific game timestep τ of game k as the degree of control a team has in terms of their likelihood of scoring versus conceding. This is quantified using the Valuing Actions by Estimating Probabilities (VAEP) framework (Decroos et al., 2019), which assigns a value to each on-ball event based on its impact on the team's probabilities of scoring and conceding. Formally, the cumulative in-game domination metric up to timestep τ, denoted as $D_{\alpha,k}^{\text{in}}(\tau)$, is given by:

$$D_{\alpha,k}^{\text{in}}(\tau) = \sum_{e \in \mathcal{E}_{k,\leq\tau}} \left[\left(\Pr_{\alpha}^{\text{scores}}(e) - \Pr_{\alpha}^{\text{concedes}}(e) \right) - \left(\Pr_{\alpha}^{\text{scores}}(e') - \Pr_{\alpha}^{\text{concedes}}(e') \right) \right] \tag{6.6}$$

Here, $\mathcal{E}_{k,\leq\tau}$ is the set of VAEP events (i.e., on-ball actions such as passes, shots, or dribbles) that have occurred in game k up to timestep τ. For each on-ball event e, $\Pr_\alpha^{\text{scores}}(e)$ and $\Pr_\alpha^{\text{concedes}}(e)$ denote the estimated probabilities that team C_α will score or concede, respectively, as a result of event e, and e' denotes the event immediately preceding e. The VAEP equation calculates the net change in scoring probability (scoring probability minus conceding probability) attributed to an on-ball action. Thus, $D_{\alpha,k}^{\text{in}}(\tau)$ aggregates the net value of all on-ball actions completed by team C_α for the game so far up to timestep τ, providing a granular, event-based measure of team dominance.

The skill value of a player α, denoted ψ_α, is defined as the cumulative contribution of their on-ball actions to team domination in historical matches. Specifically, ψ_α is the sum of the player's individual VAEP values, reflecting their overall impact on goal-scoring probabilities. These player skill values are subsequently used as features in our transition model to simulate goal events.

Traditional goal prediction models, such as the Maher model (Maher, 1982), rely primarily on historical club-level results and do not account for individual player contributions. We therefore adapt the Maher model to incorporate player-level features, and denote this model as f_{Ω}^{in} (and f_{Ω}^{pre}). The original Maher model employs Poisson regression to estimate the expected goal-scoring rate (xG) for each team, with the regression coefficients for the two teams coupled through team attack and defence parameters and a home advantage term. In our adaptation, focussing first on f_{Ω}^{in}, we retain this coupling structure but replace the feature set with the following predictors: pre-game dominance values for both teams ($D_{\alpha,k}^{\text{pre}}$ and $D_{\beta,k}^{\text{pre}}$, representing the summed historical dominance of the starting lineups scaled by fatigue), cumulative in-game dominance up to timestep τ ($D_{\alpha,k}^{\text{in}}(\tau)$ and $D_{\beta,k}^{\text{in}}(\tau)$), the interval between timesteps (10 minutes in this work), the number of defenders and attackers for each team, and a binary indicator for home advantage.

At each game timestep, the goals scored by teams C_α and C_β are assumed to follow independent Poisson distributions with expected rates μ_α and μ_β respectively, where μ_α and μ_β are obtained using regression on the features mentioned above. This formulation allows us to compute the probability of all possible goal outcomes at each game timestep, which we use in two ways: (i) to simulate match score transitions in the MDP transition function P_{in}, updating the game scoreline for the next state $s_{\tau+1}$ given the current state s_τ, and (ii) to compute expected points for the reward function R_{in} (Equation 6.5). For the latter, we use the current scoreline together with the estimated scoring rates to derive the probabilities of each possible match outcome (win, draw, or loss) under the Poisson model. The expected points, $\mathbb{E}[\text{Points} \mid s_\tau, a_\tau]$, are then calculated as the weighted sum of these outcome probabilities.

To account for rare situations where a team has fewer than 11 players due to injuries, we set a team dominance decay variable to $\varphi = 0.46$, or $\varphi = 1.67$ when facing teams with fewer players, based on the literature (Badiella et al., 2023). In most situations where team numbers are equal, $\varphi = 1$. This factor is applied multiplicatively to the relevant team's scoring rate.

For the pre-game MDP, the reward function R_{pre} requires a model that predicts the outcome of the entire match, rather than just a single match time interval. We adapt the in-game prediction model f_{Ω}^{in} to create a pre-match model, denoted f_{Ω}^{pre}. This model retains the Poisson regression framework but uses only pre-game features. Specifically, it uses the pre-game dominance values for both teams ($D_{\alpha,k}^{\text{pre}}$, $D_{\beta,k}^{\text{pre}}$) and home advantage. The model f_{Ω}^{pre} is used to predict the full-time scoreline (Ω_k) for each match, from which we derive the probabilities of all possible match outcomes by aggregating over scorelines.

The expected match points, calculated from these outcome probabilities, serve as the reward in the pre-game MDP. This approach enables analysis of the expected points of all season matches during the pre-game team selection phase. The match outcome model f_{Ω}^{pre} is also used to compute the expected points lost parameter in the in-game reward function at game timestep τ, $\mathbb{E}[\text{Points Lost} \mid s_\tau, a_\tau]$. This is achieved using the mean injury duration from the injury length distribution $\mathcal{F}$, identifying the set of all future games $\{G_{k'} | k' > k\}$ within the injury duration, and, for each player $\alpha \in C_\alpha^+$, calculating the expected points difference for each game with and without the player (by subtracting the player's skill value from the overall team dominance). This difference is weighted by the player's injury probability at timestep τ, and summed across all players to derive $\mathbb{E}[\text{Points Lost} \mid s_\tau, a_\tau]$. In the next section, we discuss solving the team selection model.

6.6 Solving the Team Selection Model

With the pre-game and in-game models established, we now turn to the problem of optimising decision-making for team C_α in both pre-match team selection and in-game substitutions. To address this, we use MCTS and team strategy modelling. Figure 6.2 illustrates the overall framework, showing how MCTS is applied in both the pre-game and in-game phases to optimise team selection and substitution decisions.

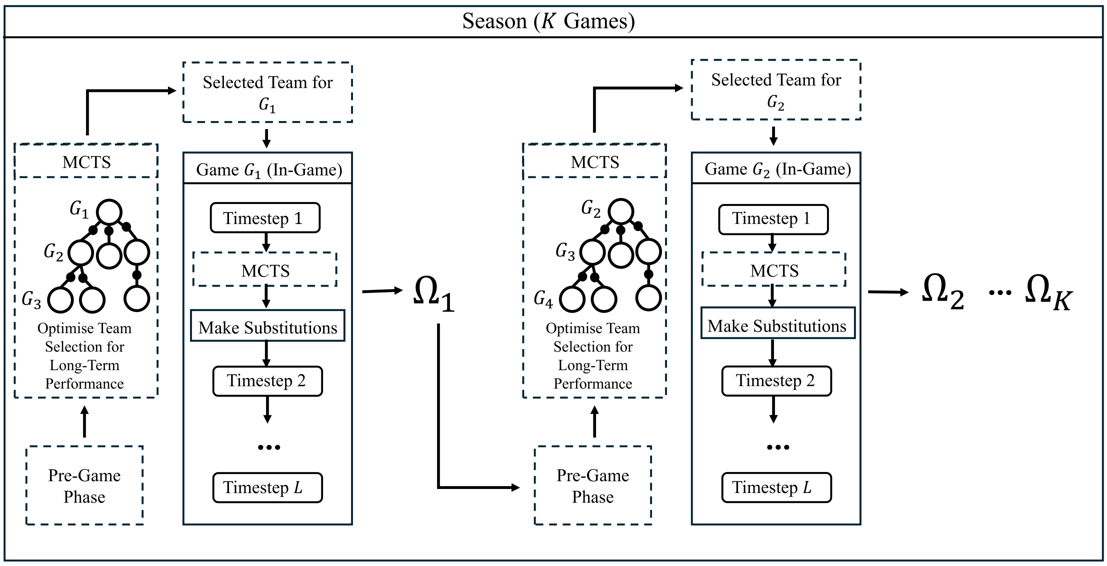

FIGURE 6.2: Framework for optimising pre- and in-game decision-making using MCTS.

For each game in the season, the framework alternates between a pre-game phase and an in-game phase. In the pre-game phase, MCTS is applied to the pre-game MDP to select a starting lineup by simulating long-term season outcomes influenced by match results and injury risks. Once the lineup is chosen, the in-game phase begins, where MCTS is applied again to the in-game MDP at each timestep to determine whether substitutions should be made. These substitution decisions balance the immediate objective of maximising match performance with the longer-term objective of managing fatigue and injury risks. When a game finishes, the match result, along with updates to player states such as fatigue and injury, feeds back into the next pre-game phase. This iterative process forms the decision-making of our model across a set of games. The remainder of this section formulates the optimisation problem, along with the algorithms and modelling approaches that support this framework.

6.6.1 Optimising Long-Term Performance

The decision-making agent for team C_α is designed to select actions that maximise the team's expected long-term reward. During the pre-game team selection phase, the agent's objective is to optimise performance over the entire season $\mathbf{G}$, indexed by $k \in K$. In contrast, the in-game selection model focuses on maximising the reward within a single game G_k, while also accounting for the impact of in-game injuries through the expected points lost parameter. The long-term reward of being in state s_k is formalised using a value function, defined as follows:

$$V_{\text{pre}}(s_k) = \max_{a_k} \left(\sum_{s_{k+1}} \Pr(s_{k+1} \mid s_k, a_k) \cdot \left[R_{\text{pre}}(s_k, a_k) + \gamma_{\text{pre}} V_{\text{pre}}(s_{k+1}) \right] \right) \tag{6.7}$$

Where $V_{\text{pre}}(s_k)$ denotes the expected value of the optimal action for state $s_k \in S_{\text{pre}}$. The recursion terminates at terminal states, i.e., when the final game of the season has been played ($k = K$). At these terminal states, the value function is defined as: $V_{\text{pre}}(s_K) = \max_{a_K} R_{\text{pre}}(s_K, a_K)$.

Analogously, the in-game value function of being in state s_τ at game timestep τ is defined as:

$$V_{\text{in}}(s_\tau) = \max_{a_\tau} \left(\sum_{s_{\tau+1}} \Pr(s_{\tau+1} \mid s_\tau, a_\tau) \cdot [R_{\text{in}}(s_\tau, a_\tau) + \gamma_{\text{in}} V_{\text{in}}(s_{\tau+1})] \right) \tag{6.8}$$

Where $V_{\text{in}}(s_\tau)$ captures the optimal expected long-term reward for in-game decision-making. Here, the recursion terminates at the end of the match, i.e., when the final timestep $\tau = L$ is reached. At this terminal state, the value function is set to: $V_{\text{in}}(s_L) = \max_{a_L} R_{\text{in}}(s_L, a_L)$, where this reward is the points from the match divided by the number of timesteps and the only action is to make no substitutions.

The objective of the decision-maker is to select, at each state, the action that maximises this value function, thereby optimising both immediate and future rewards by accounting for factors such as fatigue and injury risk. However, accurately estimating the value of actions and states is challenging, particularly due to uncertainty about the opponent's future actions. Since the decision-maker cannot directly observe the opponent's strategy, they must instead form a belief to approximate the opponent's behaviour.

In addition to uncertainty about the opponent, the decision-maker faces the complexity of a large action space and numerous sources of randomness such as injuries, injury durations and match outcomes. These stochastic elements, combined with the wide range of possible actions, greatly expand the state space and result in an extremely large search tree when evaluating long-term rewards. Specifically, the size of the in-game action space is given by $|A_{\text{in}}| = \sum_{i=0}^{\min(|C^+|,|C^-|)} \binom{|C^+|}{i}\binom{|C^-|}{i} i!$ and the size of the pre-game action space is $|A_{\text{pre}}| = \binom{|C|}{|C^+|}$. Although we apply constraints to the action set to maintain football team formations and in-game substitution rules (as discussed in Section 6.6.2), these action spaces remain large. Consequently, computing the exact optimal action from each state becomes computationally infeasible. To address this, we use MCTS to estimate the long-term values of actions through simulation, allowing us to approximate the optimal action in a tractable manner. To improve efficiency, we form beliefs about the team's strategy π_α to guide MCTS toward promising actions. Additionally, we divide MCTS into multiple steps to reduce the branching factor and make the search process more computationally tractable. In the next subsection, we describe the constraints on the action sets that ensure the decision process adheres to the rules and tactical structure of football.

6.6.2 Team Selection Constraints

In football, teams typically select a team of players whose roles fit a chosen formation. Teams also face several player substitution constraints during a match G. In this section, we impose constraints on our action set to model the aspects the manager must consider during team selection.

6.6.2.1 Pre-Game Team Selection Constraints

Each player $\alpha \in C_\alpha$ has a role κ_α which can be goalkeeper (GK), defender (DEF), midfielder (MID), or attacker (ATT). To reduce the dimensionality of the pre-game action space A_{pre}, we restrict team selections to only those lineups that satisfy a set of positional constraints, ensuring that all chosen formations are plausible. All formations observed in our real-world dataset adhere to these constraints. For simplicity, we assume that managers select the most favoured formation within these constraints (learned from real-world data) and use that formation consistently throughout the season. The action space is therefore defined as:

$$\begin{aligned}
A_{\text{pre}} = \{C_\alpha^+ \subset C_\alpha \,||C_\alpha^+| = 11, \forall \alpha \in C_\alpha^+ (\mathcal{I}_\alpha = 0),\\
\sum_{\alpha \in C_\alpha^+} \mathbb{I}(\kappa_\alpha = \text{GK}) = 1,\\
3 \leq \sum_{\alpha \in C_\alpha^+} \mathbb{I}(\kappa_\alpha = \text{DEF}) \leq 5,\\
3 \leq \sum_{\alpha \in C_\alpha^+} \mathbb{I}(\kappa_\alpha = \text{MID}) \leq 5,\\
1 \leq \sum_{\alpha \in C_\alpha^+} \mathbb{I}(\kappa_\alpha = \text{ATT}) \leq 3\}
\end{aligned} \tag{6.9}$$

Where $\mathbb{I}$ is an indicator function, equal to one if the condition is satisfied and zero otherwise. In cases where the available player pool lacks sufficient players to fill a specific role (e.g., requiring three midfielders), we model the use of a reserve player α with a baseline skill value, ψ_α, set to 0. Similarly, reserve players are used as substitutes (C_α^-) when injuries prevent a full squad from being fielded. To maintain model tractability and reflect real-world scenarios, we simplify the injury probability of such reserve players by assigning it the mean injury probability observed across all players in our real-world dataset. Similarly, we model the in-game performance of reserve players by sampling from the empirical distribution of in-game performances observed across all players in our real-world dataset. This models the common real-world solution to player injuries, and the reduced skill value serves as an implicit penalty within the model, discouraging team compositions that overly depend on reserve players due to excessive squad injuries.

6.6.2.2 In-Game Team Selection Constraints

In football, teams are subject to several in-game player selection constraints. Firstly, each team is allowed a maximum of three substitutions.[3] Secondly, if a player is injured at any game timestep τ of game k, denoted $\mathcal{I}_{k,\tau}^\alpha$, they must be substituted. If no substitutions remain, the player is substituted without replacement, forcing the team to play with one fewer player. Finally, once a player has been substituted, they cannot re-enter the match. To account for these constraints, the logic governing the set of available in-game actions for a team C_α is updated accordingly and defined as follows:

[3]For our dataset, teams were limited to 3 substitutions, though some leagues allow 5.

$$
\begin{aligned}
A_{\text{in}} = \{(\mathcal{O}^+, \mathcal{O}^-) \mid & \mathcal{O}^+ \subseteq C_\alpha^+, \mathcal{O}^- \subseteq C_\alpha^-, |\mathcal{O}^+| \geq |\mathcal{O}^-|, \\
& \sum_\alpha \alpha_{\text{sub}} + |\mathcal{O}^-| \leq 3, \forall \alpha \in C_\alpha^+ (\mathcal{I}_{k,\tau}^\alpha = 1 \implies \alpha \in \mathcal{O}^+), \\
& \forall \alpha \in C_\alpha^- (\alpha_{\text{sub}} = 1 \implies \alpha \notin \mathcal{O}^-)\}
\end{aligned} \tag{6.10}
$$

Where α_{sub} is a binary variable indicating if an agent α has been substituted. When $\sum_\alpha \alpha_{\text{sub}} = 3$, the substitution limit is reached and the action set is restricted to a single action: a tuple consisting of empty sets for both $\mathcal{O}^+$ and $\mathcal{O}^-$. However, if an injured player must be substituted while $\alpha_{\text{sub}} = 3$, the action will contain an empty set for $\mathcal{O}^-$ and the set of injured players at timestep t for $\mathcal{O}^+$. A player who has already been substituted off cannot be substituted back on.

6.6.3 Modelling Team Strategies

In this subsection, we describe how we model team selection and substitution strategies using insights derived from human expert data. These models are used to simulate opponent decision-making during MDP state transitions and to guide our MCTS algorithm towards more realistic and effective actions.

6.6.3.1 Pre-Game Strategy

One of the primary challenges in optimising team selection is the uncertainty regarding the opponent's strategy and lineup. Incorporating a model of expected opponent behaviour enables more proactive decision-making by simulating games with greater accuracy. We represent the opponent's pre-game strategy as π_β^{pre}, and the probability of the opponent taking action (selecting a lineup) a_k in state s_k as $\pi_\beta^{\text{pre}}(a_k \mid s_k)$. The decision-maker does not know what action the opponent will select and can only approximate it through the estimated policy π_β^{pre}.

The opponent's pre-game policy, π_β^{pre}, is learned from historical data using a logistic regression model, which predicts the probability of each player in the squad C_β starting the game. The model uses several features for each player: their total VAEP in past games (ψ_β), the difference in total VAEP between the player and other teammates in the same position, the number of players expected to play in that position based on the team's chosen formation, the absolute difference in total VAEP between the team and its opponent ($|D_{\beta,k}^{\text{pre}} - D_{\alpha,k}^{\text{pre}}|$), and the player's injury probability ($\Pr(\mathcal{I}_k^\beta)$). Given a vector of player selection probabilities and a set of possible team selections A_{pre}, we aggregate the individual probabilities to estimate the overall likelihood of each team selection. By systematically combining and normalising these probabilities across all possible actions, we obtain a probability distribution that reflects the relative plausibility of each team selection, grounded in the underlying player probabilities. These probabilistic models capture team selection decisions at each state, assuming that while teams do know their opponent's strategy π_β^{pre}, they remain uncertain about the exact actions their opponent will take. The same framework is applied to model the strategy of team C_α, represented as π_α^{pre}.

To train the logistic regression model, we use real-world data from the 2017/18 EPL season and evaluate its performance by predicting lineups made by human managers during the 2018/19 EPL season. The model achieves an F1 score of 0.65 and a log-loss of 0.56 for predicting which players will start in a given game G. While these results focus on players, we also compare the similarity of the entire starting teams chosen using this model to those selected by human managers in Section 6.7.5. These results establish an initial benchmark for modelling human manager player selection.

6.6.3.2 In-Game Strategy

As with the pre-game model, optimising player substitution decisions is challenging partly due to uncertainty regarding the opponent's future actions. To address this, we incorporate a model of expected in-game opponent behaviour, which enables more proactive and informed decision-making by simulating the remainder of the game, G_k, with greater accuracy when estimating action values. We represent our belief about the opponent's in-game strategy as π_β^{in}, and define the probability of the opponent taking the in-game action a_τ at state s_τ as $\pi_\beta^{\text{in}}(a_\tau \mid s_\tau)$. Although the opponent possesses a true underlying in-game strategy and action distribution, the decision-maker can only approximate these through the belief π_β^{in}.

To model the structure of in-game substitution decisions, we decompose the opponent's policy π_β^{in} into three components:

1. The probability of the number of substitutions being made at state s_τ, denoted as $\pi_\beta^{\text{in}}(|\mathcal{O}_\beta^+| \mid s_\tau)$.
2. The probability of specific players being substituted off, given the number of substitutions, expressed as $\pi_\beta^{\text{in}}(\mathcal{O}_\beta^+ \mid s_\tau, |\mathcal{O}_\beta^+|)$.
3. The probability distribution over the players being substituted on, given the players substituted off, represented as $\pi_\beta^{\text{in}}(\mathcal{O}_\beta^- \mid s_\tau, \mathcal{O}_\beta^+)$.

Combining these components, the probability of the opponent taking action a_τ in state s_τ is given by:

$$\pi_\beta^{\text{in}}(a_\tau \mid s_\tau) = \pi_\beta^{\text{in}}(|\mathcal{O}_\beta^+| \mid s_\tau) \cdot \pi_\beta^{\text{in}}(\mathcal{O}_\beta^+ \mid s_\tau, |\mathcal{O}_\beta^+|) \cdot \pi_\beta^{\text{in}}(\mathcal{O}_\beta^- \mid s_\tau, \mathcal{O}_\beta^+) \tag{6.11}$$

These policy components are learned from historical data on past substitution behaviour. For each component, we train a multi-class logistic regression model to estimate the relevant probabilities. The probability distribution for the number of substitutions, $\pi_\beta^{\text{in}}(|\mathcal{O}_\beta^+| \mid s_\tau)$, is modelled using the following features: the current game timestep τ, team dominance metrics for both teams ($D_{\alpha,k}^{\text{in}}(\tau)$ and $D_{\beta,k}^{\text{in}}(\tau)$), and player characteristics $\zeta_{k,\tau}^{C}$. The same approach and feature set are used to model the probability of players being substituted off, $\pi_\beta^{\text{in}}(\mathcal{O}_\beta^+ \mid s_\tau, |\mathcal{O}_\beta^+|)$, and substituted on, $\pi_\beta^{\text{in}}(\mathcal{O}_\beta^- \mid s_\tau, \mathcal{O}_\beta^+)$, with the latter also incorporating the characteristics of the players being substituted off. The same modelling approach is applied to team C_α, with its in-game strategy represented by π_α^{in}.

To train these models, we use real-world data from the 2017/18 EPL season and evaluate their performance by predicting substitutions made by human managers during the 2018/19 EPL season. The model achieves an F1 score of 0.80 and a log-loss of 0.45 for predicting the number of substitutions at a given game timestep τ. It further achieves an F1 score of 0.83 and a log-loss of 0.31 for predicting which players will be substituted off, and an F1 score of 0.77 and a log-loss of 0.41 for predicting players substituted on. These results establish a strong initial benchmark for modelling human manager substitution strategies, and more detailed data, such as team morale and fatigue metrics, may be incorporated in future work. This initial model of team strategies enhances the MCTS algorithm by providing probabilistic estimates of expected human manager behaviour for both teams during a match, thereby improving the realism and accuracy of decision-making simulations. During state transitions within our MDPs, opponent actions are sampled from the policies defined in this section. Furthermore, the policy for the decision-making team, C_α, is used to guide the MCTS algorithm, as detailed in the following subsection.

6.6.4 Monte Carlo Tree Search

MCTS is an anytime algorithm for sequential decision-making, which incrementally constructs an asymmetric search tree using heuristic-guided exploration (Browne et al., 2012). For stochastic MDPs, MCTS employs stochastic simulation to estimate action values. MCTS has demonstrated success in complex sequential domains characterised by large decision spaces, such as the game of Go (Silver et al., 2016) and Coalition Structure Generation (Wu and Ramchurn, 2020). In this work, we leverage MCTS to address both pre-game and in-game team selection, formulating each as a stochastic MDP. For each game G_k and each timestep τ within that game, MCTS iteratively expands the search tree, prioritising actions with high estimated value while balancing exploration and exploitation.

Within our stochastic MDP framework, the search tree alternates between decision nodes, where the agent selects an action, and chance nodes, where the environment samples an outcome according to the underlying transition probabilities. We employ distinct MCTS processes for pre-game and in-game decision-making: the pre-game MCTS optimises team selection with respect to long-term, season-level objectives, while the in-game MCTS dynamically adapts to evolving match conditions, optimising tactical decisions conditional on the initial team selection and observed in-game events.

For conciseness, we present a general description of the MCTS process that applies to both pre-game and in-game settings in this section. Here, we use generic notation for states, actions, and policies, denoted as s, a and π respectively. Specific adaptations for each algorithm are detailed in the following subsections. Both MCTS approaches have the following key characteristics:

- The root decision node refers to the initial state of the MDP.
- Each decision node encapsulates an MDP state $s \in S$.
- Each tree branch refers to an action chosen by the manager at the current tree decision node.
- Chance nodes are outcomes chosen by the environment given the connected decision node and branch.
- Terminal decision nodes refer to terminal MDP states.

For the pre-game MCTS, the search tree has maximum depth K, with each node at depth k representing the state for game G_k in the season. Thus, a path from the root to a terminal node represents all team selections across the season. In contrast, the in-game MCTS has a maximum tree depth of L, where each node contains a state at a given game timestep τ within game G_k where $\tau = L$ is the end of the game. Consequently, a path from the root to a terminal node in the in-game tree captures the sequence of tactical decisions and stochastic events throughout a single match.

While we follow the fundamental structure of the MCTS algorithm, our approach incorporates adaptations tailored to the pre-game and in-game models, detailed in the following subsections. Each MCTS iteration involves four key steps: selection, expansion, simulation, and backpropagation. There are also minor differences in how we apply the MCTS algorithm for this problem compared to our approach to spatial teamwork in Chapter 4. We explain the MCTS steps for this work below.

1) **Selection -** The most promising action from the root is selected using the PUCT algorithm (Rosin, 2011; Silver et al., 2016), which balances exploration of less-visited nodes with exploitation of high-value nodes identified in previous simulations. Additionally, it integrates prior knowledge from human manager strategies introduced in Section 6.6.3, initially biasing the tree search towards more promising actions. The algorithm is defined as:

$$a_s^\alpha = \text{argmax}_a(Q(s,a) + \epsilon_{\text{puct}}\pi_\alpha(a \mid s)\frac{|(a,s)|}{|s|}) \tag{6.12}$$

Where a_s^α is the selected action, $Q(s,a)$ is the action-value function (Q-value), representing the expected cumulative reward obtained by taking action a in state s according to our MCTS simulations, ϵ_{puct} is a constant determining the weight of exploration, and $\pi_\alpha(a \mid s)$ is the probability of selecting action a under the team's policy in state s. Additionally, $|(a,s)|$ denotes the number of times action a has been selected from the current node at state s, while $|s|$ indicates how many times the current node has been visited. The selection phase continues until an unexplored or terminal node is reached. Each node η_s leads to a child node in state s', determined by action a_s^α and the transition function P.

2) **Expansion -** Expand a node η_s by selecting an unexplored action. To prioritise promising actions and improve convergence speed, actions are sorted according to the policy $\pi_\alpha(a \mid s)$, and the most probable action is selected without replacement for each node.

3) **Simulation -** At a leaf node η_s, we approximate the value of state s,$V(s)$ (see Equation 6.7 for pre-game and Equation 6.8 for in-game), by simulating the remainder of the game until a terminal state is reached, returning the cumulative reward. During simulation, the decision-making team selects its pre-match lineup according to its policy, while the opponent's starting lineup is fixed to match the actual lineup from real-world data. For in-game decisions, both teams select actions at each timestep according to their respective policies, π_α and π_β, using the corresponding action probabilities $\pi_\alpha(a \mid s)$ and $\pi_\beta(a \mid s)$.

4) **Backpropagation -** Backpropagate the value of the new child node η_s up the tree to the root node.

In the following subsection, we give details on adaptations to this core MCTS structure for pre-game team selection.

6.6.4.1 Pre-Game MCTS

To address the large action space arising from the many possible combinations of pre-game team selections, we employ a progressive widening approach (Chaslot et al.,

2008) for the pre-game MCTS. The set of candidate actions is initially artificially limited and gradually expanded over time. This approach enables more efficient allocation of computational resources within a limited time budget by focusing exploration on the most promising actions. Empirically, we found an effective progressive branching factor at a node η_s to be $W = 1 + \sqrt{|\eta_s|}$ where W is the size of the explorable action space (rounded to the nearest integer) and $|\eta_s|$ is the number of times node η_s has been visited.

To determine the order in which actions are added to the candidate set, we rank them according to the policy $\pi_\alpha(a_k \mid s_k)$, which gives the probability of selecting action a_k under the team's strategy in state s_k. This ensures that actions most likely to be chosen by a human manager, according to our team strategy model, are explored first. As the search proceeds and W increases, the action set is widened to include less probable, but potentially valuable, alternatives. For the pre-game MCTS model, $\epsilon_{\text{puct}} = 2$.

6.6.4.2 In-Game MCTS

In contrast to the pre-game model, which employs progressive widening to manage the large action space, we address the large in-game action space by using a multi-step MCTS approach. In this approach, action selection is decomposed into two sequential stages: first, the model estimates the optimal number of player substitutions; second, it selects the players to substitute based on this number. This process enhances computational efficiency by ensuring that the model first commits to a substitution before expending resources to explore the space of all possible player substitutions.

In the first stage, the MCTS algorithm for team C_α at state s_τ approximates the optimal number of player substitutions. Here, the action space comprises all possible substitution counts at the current state. During simulations in this phase, the specific player substitutions are sampled according to the team's strategy π_α, allowing the model to estimate the value of different substitution counts while focusing on substitutions which would be favourable according to the human manager model.

After determining the number of substitutions, the algorithm evaluates individual player substitutions. At this point, the action space consists of all possible single-player replacements. Each potential substitution is simulated to estimate its Q-value, thereby ranking substitutions according to their projected long-term impact on in-game reward. In practice, our MCTS algorithm at this stage does not compute the value function as in Equation 6.8. Instead, after simulating the remainder of the match, the algorithm backpropagates the accumulated scoring rates and expected points lost values up the search tree. At each node, the value function represents the expected points for the team in the game, based on the current scoreline and the backpropagated scoring rates, minus the backpropagated expected points lost. Thus, rather than relying on the specific sequence of goals in each simulation, the value function at each node reflects the overall scoring potential and injury risk, averaged across many simulated outcomes. This method yields a more stable estimate of expected points for the remainder of the match at each node, reducing the noise from individual simulated goals and improving convergence speed.

For the selection phase of the multi-step in-game MCTS, $\epsilon_{\text{puct}} = 2$ and $\epsilon_{\text{puct}} = 200$ for selecting the number of substitutions and players to substitute, respectively. These hyperparameter values were selected empirically. After both stages are complete, the final substitution decision is made by selecting the set of player substitutions with the highest Q-values, based on the number of substitutions recommended in the first stage. For example, if two substitutions are recommended, the two unique player substitutions with the highest Q-values are chosen, ensuring that no player is substituted on or off more than once within the set. Figure 6.3 illustrates the overall structure of the in-game substitution optimisation model based on the multi-step MCTS approach.

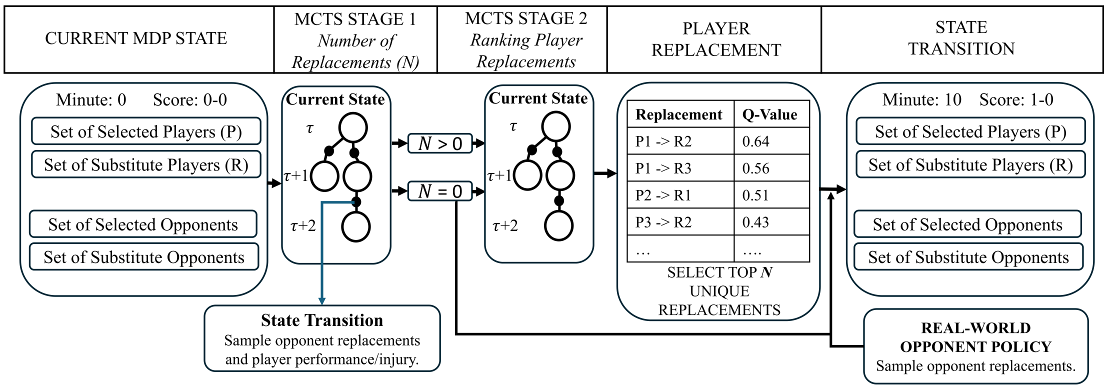

FIGURE 6.3: Decision-making framework for the multi-step MCTS algorithm to optimise substitutions at each in-game timestep.

As shown in the figure, the first stage of the multi-step MCTS algorithm estimates the optimal number of substitutions to make in the current in-game state. If the recommendation is zero, this action is taken, an opponent move is sampled using the real-world policy, and the game proceeds to the next state. If the recommendation is to make N substitutions, MCTS is then used to rank all possible single-player substitutions, from which the top N unique replacements are selected. An opponent action is again sampled, and the game transitions to the next state. This process repeats sequentially until the game ends. In the following section, we present an empirical evaluation of our team selection algorithm.

6.7 Empirical Evaluation

In this section, we assess the performance of our team selection algorithm compared to alternative decision-making strategies. We refer to our algorithm as MCTS throughout this section. The computing resources used for these experiments, along with the hyperparameters and their selection process for all models, are detailed in Appendix sections D.1 and D.2. We begin our evaluation by introducing the dataset used for our evaluation.

6.7.1 Datasets

To evaluate our approach, we use on-ball events data for the 2017/18 and 2018/19 EPL seasons, along with a player injury dataset covering these seasons and historic career injuries. The on-ball event data includes every ball-related action, such as passes, shots and tackles, for each game in these EPL seasons. For each event, the dataset records the event type, the player involved, the event location, and the time it occurred. This data is used to curate injury risk features introduced in Section 6.4 and player-based features for the team dominance and match transition model in Section 6.5. The dataset also provides team lineups and match schedules, which we use to model match times and to inform analyses of human decision-making in team selection and substitutions. This data was provided by StatsBomb.[4] The player injury dataset includes all injuries throughout the career of all players that participated in the 2017/18 and 2018/19 EPL seasons. Injuries that occurred after the 2018/19 EPL season are not included. Injury data includes the type of injury, the date the injury occurred, and the number of days the player was injured. This data was accessed from TransferMarkt.[5] These datasets allow us to assess the real-world value of our models. We evaluate our models on data from the 2017-2019 EPL seasons, as these are the only freely available league datasets containing the required data for this study. Focusing on these datasets ensures the reproducibility of our methods. More details on the datasets used in this work are provided in Appendix D.3.

We evaluated our injury risk and match transition models by training them on the 2017-18 EPL season and validating them on the 2018-19 season. A five-fold cross-validation (CV) is used on the 2018-19 season data. In each fold, the testing set consists of a chronologically ordered fifth of the 2018–19 season, while the training set includes all data from the 2017–18 season as well as all preceding matches from the 2018–19 season up to the start of the current test fold.

6.7.2 Baselines

To evaluate the performance of our proposed algorithm, we compare it against several baselines:

[4] https://statsbomb.com/.
[5] https://www.transfermarkt.co.uk/

- **Greedy:** This baseline selects the team to maximise the total player skill (ψ_α). This approach prioritises immediate team strength without considering long-term injury, fatigue, or opponent strategies. Consequently, it always fields the highest-skilled team at the start of the match and makes no substitutions.
- **Real-World Policy (RWP):** Team selection and substitutions are determined by choosing the most probable actions according to the real-world pre-game and in-game strategy models described in Section 6.6.3. This policy uses historical data and is designed to emulate the pre-game and in-game decision-making behaviour of human managers.

Directly comparing to actual real-world manager decisions is challenging due to the inherent noise and the fact that only a single sequence of injury events and match outcomes occurred in the real world. By instead evaluating against the real-world policy baseline, which is derived from aggregated historical managerial behaviour, we are able to benchmark our algorithm against human-like strategies across a wide range of simulated scenarios. This approach provides a more robust and generalisable assessment of our model's performance.

6.7.3 Experiment 1: Performance Comparison of Injury Risk Model

In this experiment, we compare the performance of our XGBoost injury model to several other ML models, all trained using the same set of features, as well as a heuristic baseline that uses the historical injury rate of players in past games as the predicted probabilities. Optimal hyperparameters are selected for each model using grid search CV (see Appendix D.2 for more details). The XGBoost model achieved a mean log loss of 0.1617 across the five CV folds. The Logistic Regression, Random Forest and Neural Network models achieved mean log loss scores of 0.1670, 0.1664 and 0.1682, respectively. The heuristic baseline achieved a mean log loss score of 0.1688. The XGBoost model therefore has lower log-loss compared to the Logistic Regression, Random Forest, Neural Network, and heuristic baseline by 0.0053 ± 0.0021, 0.0047 ± 0.0019, 0.0065 ± 0.0023 and 0.0071 ± 0.0025, respectively. We find the predictive performance improvement in using XGBoost with player-based features compared to all of the presented baselines to be statistically significant using a corrected paired t-test ($p < 0.01$) (Nadeau and Bengio, 1999).

Despite the challenging low injury rate in football ($\sim$4% in our dataset), XGBoost consistently outperforms the other methods across all five folds. We also conduct a five-fold CV with random shuffling across both seasons in our dataset to test non-season-specific patterns. For each team, we compare the predicted and actual number of injuries per season, averaged across both seasons. Figure 6.4 illustrates the relationship between the average predicted and actual injury counts for each team.

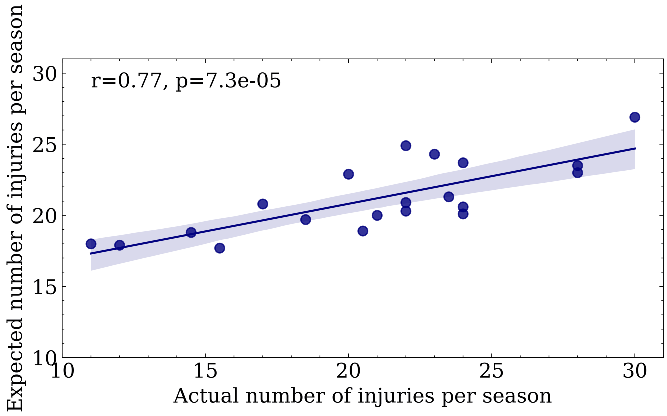

FIGURE 6.4: Predicted injuries in a season (using the XGBoost injury model) for each team compared to actual injuries.

The mean percentage difference between the expected and actual number of injuries for each team in the league is $\sim$14%. The Pearson correlation coefficient is 0.77, significant at $p<0.01$, showing a relationship between model predictions and real-world injury occurrence. The probabilistic nature of injuries and the consequential noise for injury frequency for a team in a single season contribute to the percentage difference in injuries. When assessing league-wide predictive performance, where the model predicts the total injuries in the league in a single season, the mean percentage difference between the expected and actual number of injuries is 0.8%. Specifically, in the 2018/19 season, the predicted and actual numbers of injuries in the league were 438 and 441, respectively.

6.7.4 Experiment 2: Match Transition Model Performance Evaluation

In this experiment, we compare the predictive performance of our player-level match transition models, f_{Ω}^{pre} and f_{Ω}^{in}, to the original team-level Maher model using the EPL dataset. We assess model performance using both log-loss, which measures the calibration of predicted match outcome probabilities to actual match outcomes, and the 'Compute Area Under the Receiver Operating Characteristic Curve' (AUC) score, which quantifies the overall ability of the model to discriminate between classes, which in this case is match outcomes (win/draw/loss). For the three-class match outcome setting, AUC is computed using a one-vs-rest formulation, where each outcome is treated as the positive class and the other outcomes as the negative class, and the resulting AUC values are averaged across classes. For the pre-game setting, the player-level and team-level models achieve log-loss scores of 0.94 ± 0.02 and 0.91 ± 0.02, respectively, and corresponding AUC scores of 0.69 and 0.71.

For the in-game setting, match outcome probabilities are computed by aggregating predicted goal rates at each timestep for each strategy and then deriving scoreline probabilities from these rates. This process and the features of the models were described in Section 6.5. The in-game player-level and team-level models yield log-loss scores of

0.89 ± 0.02 and 0.91 ± 0.02, and AUC scores of 0.72 and 0.71, respectively. The ROC-AUC curves for all pre-game and in-game models are presented in Figure 6.5.

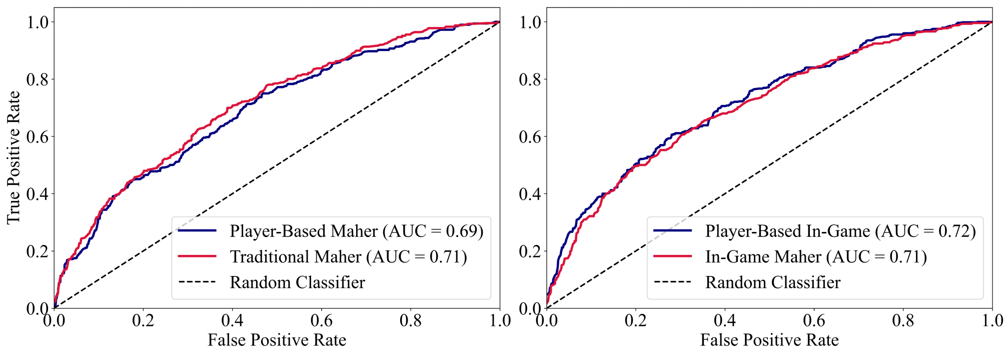

FIGURE 6.5: ROC-AUC curves for the pre-game (left) and in-game models (right) for match outcome prediction.

Although the player-level models do not yield statistically significant improvements in log-loss compared to the team-level Maher model, they achieve comparable performance within the reported error margins (95% confidence interval) to this well-established baseline. Importantly, the player-level models provide greater flexibility for measuring outcome probabilities based on the selected players for a match. This granularity enables our simulation framework to capture the impact of different team compositions on match outcomes, which is highly valuable when performing a prescriptive analysis using MCTS and evaluating the performance of our dynamic team selection and substitution model.

6.7.5 Experiment 3: Similarity to Human Team Selections

In this experiment, we compare MCTS team selections to teams chosen by human managers. We use real-world team states, including player injuries and workloads, to ensure that the comparison reflects the actual decision context faced by managers. Team similarity is quantified as the average percentage of shared players between the real-world and MCTS-selected lineups. Table 6.2 presents the similarities for MCTS and several baselines, compared to human manager team selections extracted from the EPL dataset. Among the baselines, we include a 'Last Team' strategy, which repeats the lineup chosen by the manager in the previous match from the real-world data.

TABLE 6.2: Similarity comparisons between a manager's team selections and our team selection strategies. The $\pm$ values show the 95% confidence interval of the mean value across all teams.

Strategies	% Similarity to Real-World Selection
MCTS	76.08 ± 0.66
Greedy	79.85 ± 0.68
RWP	79.91 ± 0.69
Last Team	80.92 ± 0.90

Our results indicate that real-world team selections are most similar (∼80%) to the Greedy strategy and the real-world policy baseline, rather than to the MCTS approach. This finding supports the view that managers are often more driven by short-term objectives, such as maximising immediate match performance, rather than by longer-term considerations like player fatigue management. The 'Last Team' strategy achieves the highest similarity, reflecting the common managerial tendency to minimise changes between matches. However, the performance of the other models falls within the error margins of this baseline. Unlike the 'Last Team' strategy, these models are better suited for simulated seasons because they do not rely on knowledge of the actual manager's previous team selection, which is only available in real-world data.

Given the inherent challenges of directly comparing MCTS to human decisions due to limited and variable real-world data in our two-season dataset, evaluating our model against a real-world policy baseline over numerous simulated seasons, constructed from real-world data, enables a more comprehensive assessment of its robustness across a wider range of scenarios.

We also assess the similarity of in-game substitution timings between strategies. Figure 6.6 shows the distribution of substitution timings across match minutes, highlighting the typical substitution patterns for each approach. The results show that the MCTS strategy tends to make substitutions in very similar timing patterns to those of human managers.

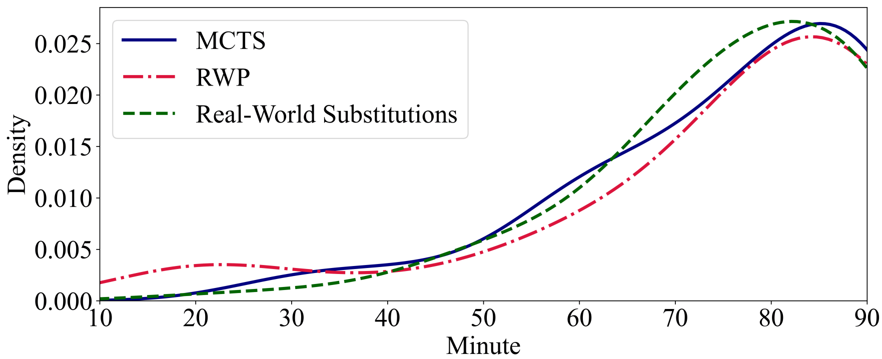

FIGURE 6.6: Distribution of the substitution timings for the MCTS and RWP strategies compared to substitution timings in real-world matches for the 2018/19 EPL dataset.

This convergence in substitution strategies is likely due, in part, to the incorporation of the real-world policy as a guide during the selection phase of MCTS, which encourages alignment with established managerial practices. Additionally, it may suggest that, when connected with the pre-game strategy, the optimal timing for substitutions identified by MCTS often coincides with the decisions made by expert human managers. In real-world settings, managers frequently delay substitutions for reasons such as tactical flexibility, a preference for making changes at half-time, or concerns about the psychological impact on players since early substitutions can be perceived as a sign of poor performance and may attract media scrutiny. While the MCTS algorithm is designed to optimise team performance without regard for these social factors, the integration of the real-world policy in its selection phase and action thresholds may implicitly capture some of these practical considerations. Nevertheless, explicitly modelling such social and psychological factors within the MCTS framework could further enhance its realism and applicability. The real-world policy baseline continues to closely mirror actual managerial behaviour, reinforcing its validity as a comparative benchmark. Given its alignment with human decision-making in both pre-game and in-game contexts, we adopt the real-world policy as the baseline for our season simulation experiments.

6.7.6 Experiment 4: Model Season Performance

In this experiment, we compare the performance of our MCTS-based approach against several baseline strategies through many season simulations. For each team in the 2018/19 EPL season, we conduct 500 simulated seasons using MCTS and 500 simulated seasons for each baseline strategy. The performance of each strategy is evaluated by the average points accumulated over the season, as summarised in Table 6.3. While Table 6.3 includes the actual points achieved by each team in the 2018/19 EPL season for reference, these values should be interpreted as a single set of outcomes for a season, conditioned on a specific sequence of matches, injuries, refereeing decisions, and other stochastic factors that can affect team performance. In contrast, our simulated season results show outcomes over many possible match and availability trajectories. As a result, discrepancies between simulated and actual points are expected, as shown by the high confidence intervals for the mean actual points across all teams. This motivates the use of repeated simulations and comparisons to the RWP baseline to assess performance.

Due to the computational demands of MCTS, these simulations were executed in parallel on a compute cluster, with each node allocated 4GB of RAM. Each simulated season follows a consistent procedure. For every match, the pre-game MCTS algorithm is run for three minutes to select the lineup with the highest estimated value. To mitigate the influence of noise, we require that an action must be selected at least 200 times within the MCTS search to be considered. This prevents actions with very few selections from being chosen solely due to high but unreliable Q-values. During the match, at each timestep, MCTS is employed with a 30-second search to determine the number of substitutions, subject to the constraint that an action must have been selected more than the total number of selections divided by the number of possible actions plus one. Subsequently, a two-minute MCTS search is used to select the specific player to substitute, if the first stage suggested making a substitution. This stage only considers the top 30 most frequently selected actions. These time constraints and action selection thresholds were determined empirically to balance computational efficiency with solution quality, given the extensive number of simulations required. In each simulated game, the opponent's

TABLE 6.3: Season performance of our MCTS strategy compared to the baselines and actual Premier League points (2018–19). Bold values indicate the superior strategy for each team, with ± values showing the 95% confidence interval of the mean value across all simulations for that team.

	Average Points			
Team	MCTS	Greedy	RWP	Actual Points
Arsenal	**58.6 ± 0.6**	56.7 ± 0.6	56.7 ± 0.6	70
Bournemouth	**53.0 ± 0.6**	52.8 ± 0.7	52.7 ± 0.6	45
Brighton	**41.5 ± 0.6**	40.9 ± 0.6	41.2 ± 0.6	36
Burnley	**48.8 ± 0.6**	48.2 ± 0.6	47.9 ± 0.6	40
Cardiff	**33.1 ± 0.5**	32.9 ± 0.5	32.9 ± 0.5	34
Chelsea	**69.2 ± 0.6**	68.6 ± 0.7	68.8 ± 0.6	72
Crystal Palace	55.5 ± 0.6	**55.7 ± 0.7**	55.6 ± 0.6	49
Everton	**49.3 ± 0.6**	**49.3 ± 0.6**	**49.3 ± 0.5**	54
Fulham	30.0 ± 0.5	**30.1 ± 0.5**	30.0 ± 0.5	26
Huddersfield	**40.1 ± 0.6**	39.9 ± 0.5	39.9 ± 0.6	16
Leicester	**51.0 ± 0.6**	50.7 ± 0.7	50.7 ± 0.7	52
Liverpool	**78.9 ± 0.6**	77.1 ± 0.6	77.6 ± 0.6	97
Man City	**86.7 ± 0.6**	85.9 ± 0.6	85.5 ± 0.6	98
Man United	**61.4 ± 0.6**	59.3 ± 0.6	60.1 ± 0.6	66
Newcastle	**42.2 ± 0.5**	41.7 ± 0.6	41.5 ± 0.6	45
Southampton	**45.9 ± 0.5**	45.3 ± 0.6	45.1 ± 0.6	39
Tottenham	**69.4 ± 0.6**	68.2 ± 0.6	68.5 ± 0.6	71
Watford	54.5 ± 0.6	**54.7 ± 0.7**	54.4 ± 0.6	50
West Ham	**50.3 ± 0.6**	50.0 ± 0.6	50.0 ± 0.6	52
Wolves	**35.5 ± 0.5**	35.0 ± 0.5	35.1 ± 0.5	57
Mean	**52.7 ± 0.3**	52.1 ± 0.3	52.2 ± 0.3	53.5 ± 9.0

starting lineup matches the actual lineup from real-world data, while their in-game substitutions are chosen according to their policy π_{β}^{in}.

Table 6.3 shows that MCTS yields a modest but consistent improvement in average season points, with a mean increase of 1.0% over the best-performing baseline. A paired t-test also supported that the MCTS strategy yielded significantly higher average points compared to Greedy and RWP ($p < 0.01$). This highlights the benefit of accounting for player fatigue and strategically planning team selections and substitutions. Figure 6.7 compares the strategies further by visualising the probability of teams achieving key season objectives in our simulations using the RWP and MCTS strategies. These probabilities are based on each team's simulated points totals relative to the thresholds required for each objective, assuming all other teams follow the real-world policy.

The results in Figure 6.7 show how teams can improve the probabilities of reaching their objectives for the season. For example, within our season simulations, Liverpool has a 4.6% added probability of winning the title using the MCTS strategy compared to if they were to follow the real-world policy. Achieving European football (finishing in the

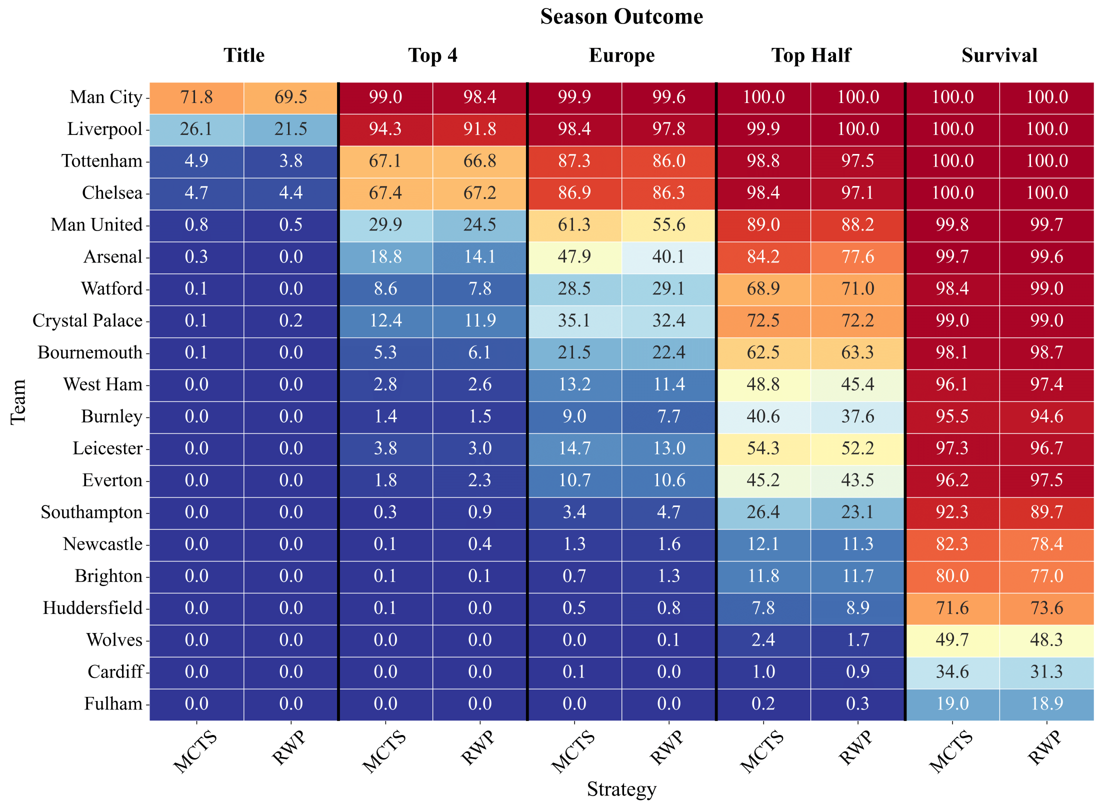

FIGURE 6.7: Percentage probabilities for each team achieving particular season objectives using the MCTS or RWP strategy.

top 6 places) is often an important goal for clubs because of the financial rewards and publicity associated with reaching these European competitions. The results show that, for example, Man United's chances of qualifying for European competition improve by 5.7% when using MCTS compared to the real-world policy. While some clubs do not see significant gains, likely due to stochastic elements such as injuries and match outcomes throughout a season, the overall trend indicates that MCTS enhances the probability of achieving season goals for most teams.

6.7.7 Experiment 5: Model Season Injury Analysis

In this experiment, we evaluate the impact of each strategy on player welfare by comparing the average number of injuries sustained by the squad and by the optimal team (the eleven highest-skilled players who fit the pre-game formation constraints) across season simulations, as shown in Table 6.4. This analysis provides insight into how effectively each approach manages player health during congested fixture periods.

On average, MCTS team selections reduce total squad injuries by approximately 2.9% and optimal team injuries by 14.7% compared to the Greedy baseline. This shows that MCTS more effectively rests key players and manages workloads, not only improving player welfare but also team performance across a season. The MCTS strategy also

TABLE 6.4: Team injuries across season simulations for our MCTS strategy versus baseline models. Values are rounded to one decimal place. Bold values indicate the strategy with the fewest injuries for each team. "Decr (%)" shows the percentage decrease in injuries achieved by MCTS relative to the Greedy strategy (which maximises skill without injury consideration).

	Squad Injuries				Optimal Team Injuries			
Team	MCTS	Greedy	RWP	*Decr (%)*	MCTS	Greedy	RWP	*Decr (%)*
Arsenal	**25.1**	26.6	26.7	*5.6*	**15.9**	19.0	18.9	*16.3*
Bournemouth	**20.6**	21.1	21.1	*2.4*	**13.3**	15.6	15.5	*14.7*
Brighton	**17.6**	18.1	**17.6**	*2.8*	**11.2**	13.7	13.4	*18.2*
Burnley	**19.1**	19.8	19.3	*3.5*	**13.0**	15.1	14.6	*13.9*
Cardiff	**20.2**	20.7	20.3	*2.4*	**13.6**	15.7	15.3	*13.4*
Chelsea	**22.0**	22.7	22.8	*3.1*	**14.4**	16.3	16.4	*11.7*
Crystal Palace	**20.2**	21.2	21.0	*4.7*	**13.1**	15.6	15.5	*16.0*
Everton	**21.3**	22.4	21.7	*4.9*	**14.1**	16.5	15.9	*14.5*
Fulham	**20.0**	20.4	20.4	*2.0*	**13.1**	15.0	14.9	*12.7*
Huddersfield	**19.9**	20.6	20.7	*3.4*	**13.3**	15.5	15.6	*14.2*
Leicester	**18.7**	19.0	18.9	*1.6*	**12.4**	14.5	14.3	*14.5*
Liverpool	**24.1**	24.5	24.7	*1.6*	**14.9**	17.5	17.4	*14.9*
Man City	**23.2**	24.2	23.8	*4.1*	**15.0**	17.9	17.5	*16.2*
Man United	**25.0**	26.0	25.9	*3.8*	**15.7**	18.7	18.6	*16.0*
Newcastle	**20.8**	21.2	21.3	*1.9*	**13.2**	15.6	15.2	*15.4*
Southampton	20.0	19.9	**19.8**	*-0.5*	**12.4**	14.8	14.5	*16.2*
Tottenham	**22.6**	23.1	23.0	*2.2*	**14.7**	16.8	16.1	*12.5*
Watford	**20.0**	20.8	20.7	*3.8*	**13.3**	15.4	15.4	*13.6*
West Ham	**23.7**	23.9	23.8	*0.8*	**14.3**	16.8	16.8	*14.9*
Wolves	**17.1**	17.7	17.7	*3.4*	**11.7**	13.7	13.5	*14.6*
Mean	**21.1**	21.7	21.6	*2.9*	**13.6**	16.0	15.8	*14.7*

reduces the number of squad and optimal team injuries compared to the RWP by 2.3% and 13.9%, respectively.

We also assess the financial implications of injuries for each strategy in the 2018/19 season using estimated EPL wage data from Capology[6]. We assume equivalent competition prize money and other financial rewards for both methods, focusing primarily on the cost incurred due to player injuries, where money is spent inefficiently on player wages whilst they are injured. To compute this value, we multiplied the number of weeks each player was expected to be injured (based on the injury numbers in our simulations) by their weekly wage (obtained from our wage dataset), and then summed these costs across all players on the team. The results of this analysis for each model are provided in Table 6.5.

Our results show that MCTS reduces inefficient wage expenditure on injured players by an average of 3.3% per club compared to the Greedy model, translating to a typical saving of around £300,000 per club per season. These estimates are likely conservative, as they do not account for the long-term costs of repeated injuries, the need for larger, more expensive squads to compensate for frequent player absences, or additional prize money from improved long-term performance.

[6]Estimated wage data available at: `https://www.capology.com/`.

TABLE 6.5: Estimated wages (£M) spent on injured players across season simulations for MCTS versus baseline models. Values rounded to one decimal place. Bold values indicate the minimum wage cost per team. "Decr (%)" shows the percentage decrease in wages inefficiently spent by MCTS relative to the Greedy baseline.

Team	MCTS	Greedy	RWP	*Decr (%)*
Arsenal	**16.3**	17.1	17.1	*4.7*
Bournemouth	**4.1**	4.1	4.1	*0*
Brighton	**3.3**	3.4	**3.3**	*3.0*
Burnley	**3.4**	3.6	3.5	*5.6*
Cardiff	**2.5**	2.6	2.6	*3.8*
Chelsea	**15.0**	15.7	15.8	*4.5*
Crystal Palace	**5.7**	6.1	6.1	*6.6*
Everton	**7.3**	7.9	7.7	*7.6*
Fulham	**4.3**	4.4	4.4	*2.3*
Huddersfield	**2.6**	2.7	2.8	*3.7*
Leicester	**6.1**	6.4	6.4	*4.7*
Liverpool	11.9	**11.7**	12.0	*-1.7*
Man City	**17.8**	19.6	18.4	*9.2*
Man United	**16.8**	17.2	17.4	*2.3*
Newcastle	**3.3**	3.5	3.7	*5.7*
Southampton	5.1	**4.9**	**4.9**	*-4.1*
Tottenham	**9.2**	9.4	9.6	*2.1*
Watford	**3.9**	4.0	**3.9**	*2.5*
West Ham	**7.8**	7.9	8.0	*1.3*
Wolves	**3.5**	3.6	3.6	*2.8*
Mean	**7.5**	7.8	7.8	*3.3*

Figure 6.8 provides a detailed analysis of how the MCTS strategy reduces the average number of optimal team injuries on a gameweek-by-gameweek basis throughout the season. This reduction is particularly significant, as a high number of concurrent injuries among key players can substantially affect match outcomes. The results show that MCTS consistently achieves lower average optimal team injuries across all gameweeks, particularly during the middle portion of the season.

The gap in optimal team injuries between strategies narrows towards the end of the season, likely because MCTS shifts focus to short-term results as the long-term impact of injuries diminishes with fewer matches remaining. Additionally, the average points difference between MCTS and the real-world policy increases as the season progresses. This reflects the cumulative benefit of improved player welfare, where higher-skilled players are available and less fatigued for critical matches later in the season. In summary, our MCTS-based team selection strategy not only outperforms baseline models in terms of season results but also significantly reduces player injuries and financial costs.

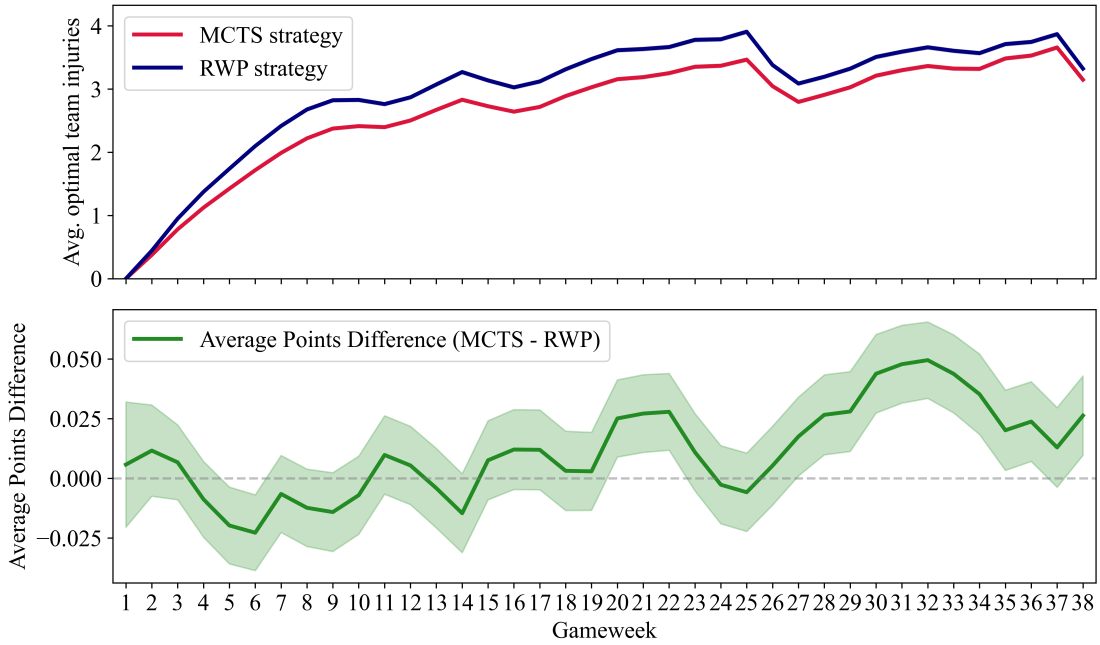

FIGURE 6.8: The average number of optimal team injuries at each gameweek across the season for the MCTS and RWP strategies. The average points difference between MCTS and RWP at each gameweek of the season is also plotted, which is the difference in average points for that gameweek alone between strategies. Shaded regions represent 95% confidence intervals, computed across the 500 independent simulations of each gameweek for both the MCTS and RWP strategies.

6.7.8 Experiment 6: Player Welfare Case Study

In this experiment, we present a case study evaluating the injury risk of an individual player over the course of the 2018/19 EPL season. Specifically, we focus on Trent Alexander-Arnold, a regular Liverpool starter, whose consistent playing time, high performance and importance to the team make him an ideal subject for assessing the impact of different team selection strategies. We compare Alexander-Arnold's participation and injury risk across simulated seasons using both the RWP and MCTS strategies, aiming to highlight the benefits of long-term planning and proactive rest management enabled by MCTS. Figure 6.9 illustrates Alexander-Arnold's injury probability throughout the season under both strategies.

As illustrated in Figure 6.9, the MCTS strategy consistently reduces Alexander-Arnold's average injury probability compared to the RWP approach. This demonstrates the effectiveness of MCTS in identifying suitable rest opportunities and managing workload to mitigate injury risk. These findings reinforce the value of proactive workload management for long-term player welfare, especially when considered alongside earlier results showing that MCTS can improve both team performance and player health. To further understand how these strategies affect player utilisation, Figure 6.10 presents the proportion of simulated matches in which Alexander-Arnold was selected each gameweek for both strategies.

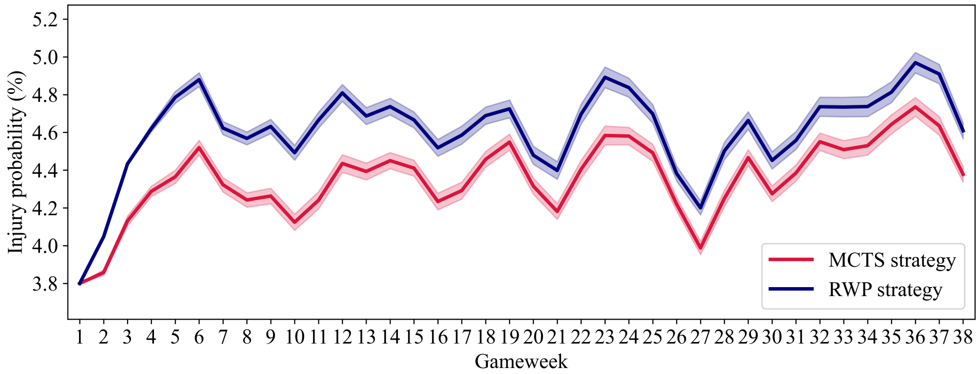

FIGURE 6.9: The average injury probability of Alexander-Arnold at each gameweek across the season for the MCTS and RWP strategies. Shaded regions represent 95% confidence intervals, computed across the 500 independent simulations of each gameweek for both the MCTS and RWP strategies.

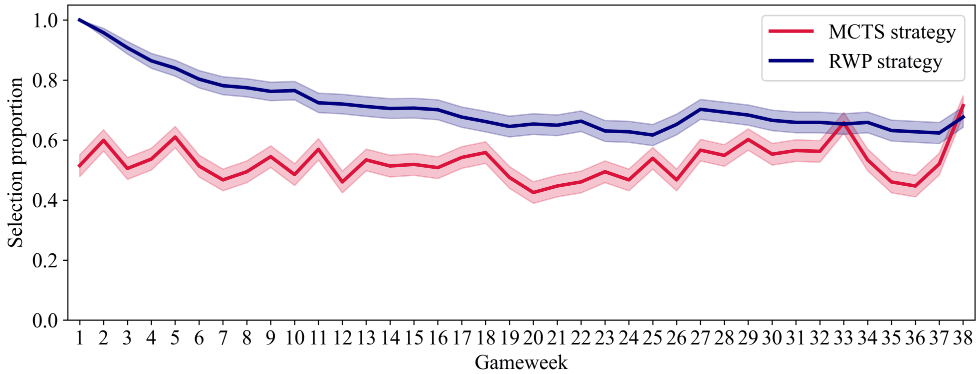

FIGURE 6.10: The average proportion of times Alexander-Arnold was selected to play at each gameweek across the season for the MCTS and RWP strategies. Shaded regions represent 95% confidence intervals, computed across the 500 independent simulations of each gameweek for both the MCTS and RWP strategies.

Figure 6.10 shows that MCTS rests Alexander-Arnold more frequently while maintaining competitive team performance. The greater variance in selection under MCTS suggests that the algorithm strategically prioritises his participation in more critical fixtures. The decrease in selection rate for the RWP strategy as the season progresses will be significantly due to the player's increasing unavailability as a result of injuries. Notably, this decrease is not shown for the MCTS strategy, suggesting that the strategy is less impacted by injuries, which aligns with the increasing points difference shown in Figure 6.8. Additionally, the convergence in selection rates for both strategies towards the end of the season may indicate that MCTS places less emphasis on injury risk when the long-term consequences are minimal. Extending the model to consider multiple seasons could further incentivise the protection of player health for the entire set of matches, including the matches near the end of the season.

6.8 Discussion

We now discuss the main findings of our work, structured into key themes to highlight their implications and limitations.

Comparison with Human Decision-Making Our analysis reveals that real-world team selections are most similar to the Greedy and real-world policy baselines, supporting the idea that managers often prioritise short-term objectives. The divergence between MCTS and human team selection likely reflects the influence of social and psychological factors, such as player morale, media scrutiny, and the immediate pressure to achieve results, that are not explicitly modelled in our approach. In contrast, the MCTS approach focuses on systematically planning rest and substitutions from a long-term perspective, explicitly considering player welfare and fixture congestion. Prioritising actions that are more similar to human managerial decisions in the MCTS selection phase aims to guide the model towards more typical pre- and in-game decision-making, as seen in the distribution of substitution timings for each strategy. These findings suggest that while algorithmic strategies like MCTS can provide valuable insights, their integration into practice should consider the broader context of human decision-making. As such, MCTS-based recommendations are best positioned as decision-support tools that inform, rather than replace, managerial expertise. Furthermore, incorporating social and psychological considerations into future models could further enhance their relevance and effectiveness in real-world settings. This could be achieved using data derived from domain expertise (e.g., expert annotations and contextual features), sentiment analyses of social and mainstream media, and survey data focussing on player morale.

Applicability Beyond Football Our sequential team selection algorithm, while initially applied to football due to the availability of real-world data, has significant potential across various domains, including other team sports and emergency response. Its ability to improve player welfare, reduce injuries, and optimise performance for objectives such as points underscores its wide-ranging applicability. This is particularly pertinent in sports with higher injury risks. Future work could further validate the model's robustness by deploying it in more real-world settings with available data. We are optimistic about its applicability to other team sports, as previous studies have identified many shared challenges across various sports (Radke and Orchard, 2023). However, there are domain-specific components to the model, such as the match reward and injury model, which would require appropriate adaptation to capture the objectives and welfare risks of agents in other domains.

Generalisability Over Time In addition to extending the model to other domains, we aim to assess the model's generalisability over time by applying it to more recent football datasets as they become available. While this study focuses on the 2017-19 EPL seasons to ensure reproducibility, expanding our analysis to include additional seasons and leagues would allow us to evaluate the model's robustness across different contexts and its adaptability to the evolving nature of the sport, including changes in tactics and player conditioning. This broader evaluation would help to support our claims about the practical value of the model in improving team performance and player welfare.

Practical Feasibility Currently, the model is capable of analysing weekly real-world data to offer insights into injury risks and actionable recommendations for team selections. For in-game substitutions, MCTS operates as an online algorithm, running for 2.5 minutes at each decision timestep to provide substitution recommendations. While this is feasible with 10-minute intervals, further reducing computation time would enhance its practicality for real-time use, as major changes to the game state, such as goals scored or new injuries, could occur during the algorithm's runtime, meaning that recommendations may become outdated. Pre-game team selection was limited to three-minute run times to balance computational feasibility with solution quality; however, in real-world settings, teams typically have much longer periods between games to run the algorithm. Since MCTS is an anytime algorithm and the search space in this domain is highly complex, it is challenging to know the action optimality. Therefore, it is likely that longer runtimes could yield further performance improvements for teams. Investigating the impact of longer search times remains an important direction for future work.

Future Research Areas Incorporating richer data sources, such as training load, GPS tracking and other sensor-derived physical metrics, could enhance the accuracy of injury risk predictions and further improve the effectiveness of the MCTS decision-making. We also recognise that players are susceptible to different types of injuries, and that certain injuries may result in longer absences. Although this chapter assumes all injury durations are drawn from a single distribution for simplicity, future work will account for varying injury types and their specific impacts on recovery time. Furthermore, teams often participate in additional cup competitions beyond the EPL, leading to more congested schedules and increasing the importance of proactive player welfare management. We aim to incorporate these additional fixtures into future analyses. We will systematically evaluate how these enhancements affect both the predictive accuracy of the injury risk model and the overall performance of MCTS throughout the season. We also plan to extend our approach to multi-season optimisation, enabling a deeper assessment of long-term player welfare. Finally, an interesting direction for future research is to explore scenarios where both competing teams employ MCTS strategies, providing insight into the dynamics of algorithmic, long-term decision-making when widely adopted across the sport.

Summary of Findings This model offers a baseline algorithm for team selection that focuses on player welfare and injury management to improve team performance. Experimental results demonstrate that this approach can reduce injury incidence, optimise financial investments, and support more effective long-term squad management while improving season performance. Additionally, it encourages more sustainable and transparent team management practices through explainable injury risks. While our evaluation leverages simulated seasons to enable robust statistical analysis, we acknowledge that this may not fully capture the complexities of real-world implementation, such as player psychology and media pressure. To bridge this gap, we will discuss this model with coaches and plan to integrate it into real-world decision-making, combining algorithmic insights with human expertise to assess its practical impact. Developing explainable artificial intelligence (AI) tools to communicate model recommendations to managers and coaches could enhance the connection between algorithmic and human decision-making, leading to more effective adoption in practice.

6.9 Summary

In this chapter, we introduced a sequential algorithm for TF and substitutions in football, a dynamic and adversarial environment. Our approach focuses on long-term team performance by reasoning over player fatigue, injury risk and opponent strategy. By modelling both pre-game team selection and in-game player replacement as an MDP and leveraging a multi-step MCTS approach guided by human-expert policies and an XGBoost-based injury risk model, we developed a system capable of making informed, long-term focused decisions that account for both immediate match outcomes and the cumulative effects of player workload.

Our empirical evaluation on EPL data from the 2017-19 seasons demonstrates that our model can significantly reduce first-team player injuries by $\sim$15% while also showing a $\sim$1% increase in long-term team performance, leading to more efficient financial investment by reducing wage expenditure on injured players by approximately 3%. Our analysis shows that real-world managerial decisions often favour short-term, skill-based strategies, highlighting the value of algorithmic tools for supporting more sustainable, long-term planning. Although our evaluation uses simulated seasons for robust statistical analysis, we recognise that this does not fully capture the psychological and social complexities of real-world football management. However, discussions with coaches and the model's ability to analyse real-world data demonstrate its practical potential.

Several areas of future work have been discussed to advance this work further. Firstly, incorporating richer data, such as training loads, physiological sensor data, and more recent seasons, could improve injury risk predictions and enhance the adaptability of the model to the evolving dynamics of sport. Secondly, integrating social and psychological factors into the model may allow it to better reflect the considerations of human managers and provide deeper insight into how these factors influence team performance. Thirdly, completing real-world trials and developing explainable interfaces would be vital to integrating the algorithms' recommendations into insights for coaches and decision-makers. Finally, we plan to validate our algorithm in other dynamic adversarial settings.

In summary, this research establishes a baseline algorithm for long-term team planning in football, demonstrating clear benefits for both sustained team performance and player welfare. Our results highlight the importance of integrating real-time adaptation and proactive risk management into TF, advancing beyond current approaches that overlook the long-term effects of fatigue and injury or focus on pre-game selection. With suitable domain-specific adjustments, our model holds potential for broader application to other dynamic, adversarial TF challenges.

In the next chapter, we review the outcomes from the previous four chapters and discuss future directions for the use of AI in football.

Chapter 7

Discussion

In this chapter, we summarise the general contributions of our work and its anticipated impact on the field of artificial intelligence (AI) in sports analytics (Section 7.1). We then summarise the emerging challenges identified in this thesis (Section 7.2). Finally, we discuss the future for the application of AI in football (Section 7.3).

7.1 Research Retrospective

We reflect on the key insights and lessons derived from the research presented in this thesis, highlighting both its theoretical contributions and practical implications for the field of AI in sports analytics.

The Advantages of Spatial Imputation This thesis addresses the common challenge of limited data availability in sports analytics by introducing a spatial imputation model capable of inferring player tracking data from sparse on-ball event data. By imputing player positioning, this approach enables a wide range of analyses previously restricted to organisations with access to comprehensive tracking systems. Notably, it facilitates the estimation of player physical metrics, the generation of player heatmaps, and the calculation of pitch control. Furthermore, the model's utility extends to advanced applications, including the simulation of play sequences and the prediction of player injuries, as demonstrated in Chapters 4 and 6. This approach lays the foundation for applying such a model to make existing and future tracking-based research more accessible in the field. Beyond football, the methodological framework established here offers a generalisable solution for domains where observability may be very limited, such as other team sports or disaster response scenarios.

Advancing Off-Ball Analysis This thesis contributes to the emerging field of off-ball analysis in sports analytics, a domain which is relatively new due to recent advances in tracking technology. While prior research has primarily leveraged off-ball data to enhance the contextual understanding of attacking metrics (Fernández et al., 2021; Spearman, 2018), significant gaps were identified, particularly in evaluating defensive performance, which has traditionally relied on event-based data (Merhej et al., 2021). The research presented here addresses these limitations by developing an optimisation framework for defender positioning to mitigate attacking threats and by introducing novel off-ball defensive contribution metrics (Chapter 5). These contributions not only enhance the predictive and prescriptive capabilities of defensive analysis but also lay the methodological foundations for applying deep learning and search algorithms to more comprehensive off-ball evaluation in football and other team sports. The thesis uses practical visualisations to illustrate how real-world clubs and analytics teams could integrate these models. Collectively, these advancements deepen our understanding of the critical role of off-ball actions in football match outcomes and set a new direction for future research on off-ball performance evaluation.

AI for Decision-Making in Football Optimising decision-making is an emerging area in football analytics, with recent studies beginning to optimise corner positioning (Wang et al., 2024), on-ball actions (Rahimian et al., 2021), and team selection using event data (Beal et al., 2020a). This thesis demonstrates how integrating off-ball tracking and injury data can further enhance prescriptive decision-making. We show that optimising off-ball player positioning can reduce opposition threat (Chapter 4), and that team selection strategies accounting for injury risk can improve long-term performance and player welfare (Chapter 6). These methods offer practical tools for club analytics teams in both pre- and post-match analysis. A key challenge in optimising decision-making is accounting for the effect of behavioural changes on the opposition, such as how a defender's repositioning influences attacker behaviour, or how substitutions may lead to tactical changes from opponents. We address this by employing predictive methods that model opponent responses to tactical adjustments. However, we emphasise that this remains an important consideration for future research. Notably, Monte Carlo tree search (MCTS) serves as a core optimisation technique in our prescriptive applications in this thesis, demonstrating its effectiveness for optimisation tasks in football and its potential for wider use in sports analytics. Further exploration of other reinforcement learning (RL) methods would be valuable to assess their performance for similar tasks.

Spatial Teamwork In this thesis, we have highlighted the importance of player teamwork, which refers to the communication and coordination between players, in improving team performance. However, it is challenging to measure these factors, and current research is limited. In football, existing teamwork research has focused on one-to-one interactions through football passing sequences (Beal et al., 2020b). Studies in other team-based domains mainly focus on efficient task allocation, aiming to minimise travel times and assign the most suitable agents to specific roles. This thesis advances the field by focusing on spatial teamwork, where agents coordinate their positions to maximise defensive spatial control. Furthermore, to our knowledge, this work introduces the first approach in football analytics aimed at optimising player teamwork. Interestingly, our work found that it is most favourable to consider player spatial coordination in pairs, which aligned with previous findings related to football passing (Beal et al., 2020b). Overall, this research establishes a theoretical foundation for optimising spatial teamwork in real-world settings, such as sports teams or security patrols. It also demonstrates how club analytics teams could apply these methods to assess and improve player coordination, particularly during post-match analysis.

Football as a Testbed for Team-Based AI Research In this thesis, a commonly discussed theme is the suitability of football as a testbed for developing and validating new models in team-based AI research. This is primarily due to clear objectives in the sport and the wide availability of real-world datasets. This point has been emphasised in previous studies (Beal et al., 2019; Tuyls et al., 2021) and remains true for future football research. Our work further underscores football's relevance by explaining similarities to other team sports, as well as to disaster response, security and patrol, where teams operate in spatiotemporal settings. The methods and findings in this thesis, which are grounded in real-world data, have the potential to inform AI research in these domains. These advancements include spatial imputation in limited observability environments, modelling spatial teamwork and agent interactions, evaluating indirect defensive contributions, and optimising long-term team formation (TF) with consideration for agent welfare and availability. While our models are designed to be generalisable, we note that domain-specific adaptations may be necessary when applying them beyond football, and extensive model validation is also required for these domains.

7.2 Emerging Challenges

Although each chapter of this thesis includes its own discussion of future work, here we summarise several overarching challenges for further research. These challenges arise across multiple areas of this work.

Real-Time Prescriptive Analysis The current prescriptive analytics models developed in this thesis, which use MCTS search algorithms and linear programming, are designed to take between 30 seconds and 3 minutes to run to ensure the models reach satisfactory convergence. However, for real-time applications, such as optimising player spatial positioning in real-time or making in-game substitution decisions, these runtimes may limit practical use, as the game state can change significantly by the time the algorithm produces a solution. It is likely that the primary use of models in this thesis, and particularly the spatial teamwork model, would be for post-match analysis. A key area for future research would be to improve the efficiency of these prescriptive models towards reaching solutions, thereby increasing their potential for informing decision-making in real-time during matches. This process would involve exploring ways to improve algorithm efficiency, such as implementing deep Q-networks within MCTS or experimenting with other RL-based approaches.

Improved Dataset Depth and Quality The datasets used in this thesis include on-ball events and tracking data from a season of the South Korean K League 1 and English Premier League (EPL) data. Additionally, two seasons of EPL events data and historical injury data were used. These datasets provide a strong foundation for off-ball and physical player information in the development of AI models for football. As new datasets become available, such as tracking data that spans a wider range of seasons, it will be important to validate our approaches across more leagues. This would help assess how model performance varies between different competitions and seasons, further improving the robustness of model evaluations. Furthermore, the inclusion of more advanced training and player physical data could improve the predictive accuracy of the injury prediction model and the effectiveness of the long-term team selection model. In the future, as even more detailed datasets emerge, such as body pose data, our models could be updated to incorporate these new factors. For example, integrating information on player balance could provide deeper insight into defensive contributions for the *GAPP* model in Chapter 5.

Off-Ball Attacking Teamwork and Credit Assignment This thesis offers contributions to off-ball credit assignment and the optimisation of the spatial positioning of players. These studies concentrated on defenders, whose purpose is to reduce attacking threats. Although these contributions address important research gaps, especially given the limited exploration of off-ball defending in the field, there are also gaps in optimising attacking positioning to increase threat and in assigning credit for the spatial positioning of attackers in facilitating space for teammates. The models in this thesis set a foundation for these research gaps to be addressed, and future work could implement and modify the spatial teamwork (Chapter 4) and *GAPP* (Chapter 5) models to focus on attacking positioning.

Model Validation in Other Team-Based Domains The models presented in this thesis, as shown in the research retrospective, can be applied beyond football, such as in other team sports and security. Certain models may require domain-specific modifications to enhance their applicability to specific domains. Testing the generalisability of our models in other team-based areas would illustrate their wider relevance and practical significance. This validation would also help us better understand the utility of football as a testbed for developing and validating models in team-based domains where acquiring real-world data is more difficult.

7.3 The Future of AI in Football

Our research has made several contributions to understanding how AI can be effectively applied in football, with a particular focus on the use of spatiotemporal data. However, rapid developments in data collection and AI models will continue to open up new areas of potential research. For example, advances in body pose and eye-tracking data are likely to create many opportunities for more detailed analysis of player behaviour. This data could extend the work presented in this thesis by enabling a finer-grained assessment and optimisation of player communication and coordination in matches. Furthermore, such data would facilitate improved modelling of player intentions, which may enhance the effectiveness of inverse reinforcement learning (IRL) and imitation learning (IL) approaches for capturing and replicating typical player behaviours in specific game situations.

Reinforcement learning is a branch of machine learning (ML) that has seen significant advances in recent years. However, its application in football analytics is relatively limited. Existing work includes Rahimian et al. (2021), who use deep RL to optimise on-ball actions, and Matthews et al. (2012), who use Bayesian RL to optimise fantasy football lineups. In contrast, this thesis employs search algorithms to optimise spatial coordination between defenders, as well as team selection and substitutions. Simulated environments like Google Research Football (Kurach et al., 2020) have also been developed to provide valuable platforms for training RL agents in football scenarios. Despite these developments, there is still considerable potential for RL to be applied more broadly in football, such as optimising attacking runs or adapting tactics dynamically during matches. Importantly, football presents several challenges that make it a complex environment for RL models, including sparse reward signals due to infrequent goals, and a complex, dynamic environment with many interacting decision-makers and a vast state space. Despite these difficulties, exploring deep RL, multi-agent RL, and other advanced approaches on real-world football data offers promising directions for future research (Radke and Orchard, 2023).

Most current approaches in football analytics focus on pre-match and post-match analysis, providing tools and metrics to evaluate players for team selection and recruitment (Beal et al., 2020a; Decroos et al., 2019). However, real-time analysis presents a promising area for future research, offering the potential to deliver actionable insights during matches. Such insights could support decisions related to player positioning, tactical adjustments, and substitutions by leveraging spatiotemporal data that reflects the evolving state of the game. The key challenge faced by real-time models is the computational constraints that limit the speed of data processing and analysis. This makes it difficult to deliver timely insights capable of keeping pace with football's fast and dynamic nature. To overcome this, the development of fast, robust AI models and efficient data pipelines capable of handling large volumes of streaming data is likely to be essential. If these challenges are addressed, AI systems could provide valuable real-time recommendations, enabling coaches and analysts to make more informed decisions under the pressures of a game. A promising direction could be the use of digital twins as virtual replicas of teams, which can simulate match scenarios in real time. These models could be used for scenario planning, helping coaching staff explore the potential outcomes of tactical changes before applying them on the pitch.

Another area likely to be transformed by AI is officiating (Spitz et al., 2021), particularly in offside decisions. Advances in multi-camera tracking and pose estimation can support offside technologies, reducing decision times and improving consistency even in crowded situations. Beyond offside decisions, learned models for foul detection, trained on league-specific officiating rules, could assist virtual assistant referee processes by informing decision-making and providing interpretable visual evidence to the referees (Held et al., 2025). These advancements would also introduce challenges, including fairness and bias, adaptability to varying stadiums, and the need for transparent protocols that preserve referee authority while leveraging the accuracy of well-calibrated AI models.

A significant barrier to the widespread adoption of AI models in real-world football analytics is the issue of trust (Beal et al., 2019). For coaches to rely on these models, interpretability is essential. However, many state-of-the-art AI techniques, such as deep neural networks, are black boxes, making it difficult to understand the reasoning behind their predictions. This challenge is also evident in parts of this thesis, for instance, while our MCTS algorithm in Chapter 6 can identify optimised team selections and the injury prediction model can provide feature importance, it is challenging to fully explain why a particular team configuration is chosen for a match. Consequently, combining human expert knowledge and explainable AI methods tailored to football analytics represents a key area for future research. Approaches such as feature importance analysis, attention mechanisms, and intuitive visualisation tools can enhance the transparency of model outputs, enabling coaches to better understand and trust the recommendations. This area of research is vital to bridge the gap between complex AI systems and domain experts and to ensure that AI tools are trusted within the football community.

The rise of large language models (LLMs) in AI presents new opportunities for enhancing model explainability in football analytics through natural language. In fact, some models are already being used to generate explanations for scouting reports and goal probability predictions (Rahimian et al., 2025). There are many potential applications of LLMs in the field beyond these applications, such as for analysing match reports, generating game commentary, and compiling tactical documents. When combined with structured data models, LLMs can effectively summarise complex datasets and model insights, enabling coaches to interact with and understand sophisticated analytics through natural language, without the need for constant involvement by model developers (Naveed et al., 2025). However, there exists a number of challenges, including ensuring the factual accuracy of LLM outputs and adapting these models to the specialised language and domain knowledge of football. Rigorous verification of these models will be essential to build trust and prevent the model from outputting misleading or incorrect information to coaches and analysts.

Our work in this thesis contributes to the field by providing additional data-driven tools to support decision-making in football. While there are additional human factors that these models do not capture, such as psychology and individual personalities, it is clear that they would be most valuable in informing, rather than replacing, human decision-making (Beal, 2022). This raises questions about the appropriate weight to assign to data-driven insights and how best to integrate these models within football clubs to maximise efficiency in addressing the practical challenges faced by coaches. To maximise the impact of emerging AI tools, close collaboration between data scientists and coaching staff is likely to be essential, and a continuous feedback loop could be formed through regular meetings, workshops and collaborative development of analytics tools. As AI becomes increasingly prevalent across many industries, a key question for football is the broader impact of emerging models and algorithms, such as those presented in this thesis, on the competitive dynamics of the sport, especially if their adoption becomes widespread across many teams. Finally, football is likely to continue to serve as a valuable real-world testbed for emerging AI technologies, offering a platform to rigorously evaluate new models in a complex, dynamic environment and helping to drive progress in the broader field of AI (Tuyls et al., 2021).

Chapter 8

Conclusions

In this thesis, we have investigated the practical application of artificial intelligence (AI) to advance football analysis. In particular, we focus on spatiotemporal domains and challenges such as agent trajectory prediction, agent performance modelling, and team formation (TF). We also review existing AI implementations in football, focussing on approaches that leverage spatiotemporal data. It is also identified that football's well-defined objectives and rich datasets provide a suitable testbed for developing and testing new AI models and algorithms. Our research demonstrates the potential for AI to enhance proactive decision-making and optimise club processes, including training, tactical analysis, and player recruitment. Throughout the thesis, we have identified several open challenges in this field and contributed novel approaches that address these gaps, thereby advancing the state-of-the-art in football analytics.

Our research presents several key contributions and findings. In Chapter 3, we have shown that an long short-term memory and graph neural network-based model, named *Agent Imputer*, can impute player locations from sparse on-ball event data under limited observability and non-uniform timesteps, outperforming a set of naïve imputation and machine learning baselines. By using this model, player locations can be estimated to within $\sim$6.9m, enabling downstream off-ball analysis, such as physical metrics, player coverage, team pitch control, and advanced off-ball models like those presented in this thesis for clubs with limited tracking resources. Next, in Chapter 4, we introduced a novel spatial teamwork model that captures agent coordination and optimises agent defence against adversaries. When evaluated using real-world football data, our spatial teamwork model reduces opponent threat by up to 24% compared to real-world outcomes, with a 7% reduction achieved compared to an Monte Carlo tree search (MCTS) based approach that ignores inter-agent interactions.

In Chapter 5, we introduced a novel framework for quantifying indirect agent contributions in team-based defence and proposed two new defensive metrics. Our graph attention network-based model, named *GAPP*, predicts football pass reception with a 6.4% reduction in binary cross-entropy (BCE) loss compared to the best-performing baseline. We demonstrated how this approach advances off-ball defensive analysis in football and identified a relationship between the proposed defensive metrics and future on-ball defensive actions. Finally, in Chapter 6, we present a TF and dynamic agent replacement model that incorporates agent unavailability risk learned from real-world data, expanding on existing TF frameworks that typically focus on agent skills or travel times. Using MCTS, the model effectively balances short-term team performance with long-term injury risk. When applied to football team selection and substitutions, the model achieves a $\sim$1% improvement in long-term performance compared to simulated human decision-making, while reducing first-team injuries by $\sim$15%. Additionally, the model achieves a $\sim$3% reduction in wages inefficiently spent on injured players, which equates to approximately £300,000 saving per club per season for English Premier League teams, highlighting the competitive and financial advantages of this strategy.

Building on prior work in the field (Beal, 2022), which was limited to event data and exposed tracking data as a key factor for more in-depth analysis, this thesis has evaluated how AI can advance football analytics by leveraging emerging spatiotemporal datasets. Practically, we have made contributions that expand access to player tracking data for clubs and develop new methods using tracking data to evaluate how players coordinate spatially with teammates and contribute indirectly to game outcomes through off-ball positioning. On a prescriptive level, we show how these insights can be used to optimise spatial coordination and inform team selection and substitution strategies. Together, these contributions enhance both the availability and practical use of off-ball spatiotemporal data in football. Theoretically, our contributions include novel approaches for imputing agent locations under limited observability and irregular timesteps, optimising agent coordination and quantifying indirect contributions of agents in team defence, and designing a TF framework that balances long-term player welfare with team performance.

In conclusion, the work presented in this thesis includes several novel models that contribute to both football analytics and the wider AI community. On a practical level, our work allows clubs with limited tracking resources to access and perform off-ball analysis, introduces new metrics and recommendations to enhance pre- and post-match off-ball tactical analysis and scouting, and provides data-driven recommendations for team selection and substitutions both before and during matches. We envision these models complementing domain expertise to support informed human decision-making, and we see building trust in these models as a key objective. Moreover, we contribute to general AI research on team prediction, evaluation, and optimisation, validating our models on real-world data. This opens potential pathways for these models to be validated in other team-based domains, such as other team sports and security. Finally, we highlight key areas for future research, especially as more in-depth datasets become increasingly accessible, setting the stage for further AI applications in football analytics.

Appendix A

Appendix for Multi-Agent Spatial Imputation to Infer Player Movement

In this appendix section, we give additional details on the models and data used in Chapter 3 of this thesis. Specifically, we discuss the implementation of the *Agent Imputer* model (Section A.1) and give additional details on the dataset used (Section A.2).

A.1 Implementation Details

In this section, we list the compute resources used in the work in Chapter 3 (Section A.1.1), and provide further details on the feature embedding process (Section A.1.2).

A.1.1 Compute Resources

This work was implemented in Python 3.8.8, and our machine learning functionality used PyTorch.[1] Model training was carried out on a remote graphics processing unit (GPU) service using a NVIDIA V100 GPU with 32 GB of video random access memory (VRAM) facilitated by Google Colab.[2] This used Compute Unified Device Architecture (CUDA) to enable training on the GPU. Training the *Agent Imputer* model took a maximum of three hours for each fold.

A.1.2 Feature Embedding

We embedded our categorical variables (`agentRole`, `agentSide`, `agentObserved`, `goalDiff`, `eventType`) using random embedding vectors where each vector had size equal to the

[1] `https://pytorch.org/`
[2] `https://colab.research.google.com/`

square root of the number of classes for each variable. This increased the total categorical feature dimensionality from 5 to 14. Future work should consider the use of trainable embeddings. We normalise spatial features using the `MinMaxScaler` for the X and Y directions individually, and for time-based features, we use the `RobustScaler` from the `scikit-learn` package.[3] This appropriately scales time whilst being robust to outliers caused by abnormally long periods between events (e.g., an injury).

A.2 Dataset Details

In this section, we provide further details on the dataset structure (Section A.2.1) and player positions (Section A.2.2) in Chapter 3.

A.2.1 Dataset Structure

The ground truth tracking dataset we use logs player position at 30 Hz. To align player tracking with event data, we find the closest tracking timestamp to the event timestamp (usually within a few milliseconds). For our data, the set of possible event types is as follows:

- **Attacking Events -** `Pass, Pass Received, Foul Won, Shot, Take-On, Ball Received, Control Under Pressure, Carry`
- **Defensive Events -** `Foul, Duel, Tackle, Save, Intervention, Recovery, Interception, Block, Clearance, Error, Goal Conceded, Aerial Clearance, Own Goal`
- **Other Events -** `Offside, Pause, Deflection, Substitution`

The player who receives a pass is not highlighted for a pass event, and a pass received event is instead referenced. This pattern is consistent throughout the data, whereby related events are referenced for each event.

A.2.2 Player Positions

The event data provided gives a label for player role in each game. There are 16 unique roles which we pass into our model. When applying downstream analysis, we group roles for clarity and to increase the amount of data for each role. Below, we present the role groups (bold) and roles within each group:

[3] https://scikit-learn.org/stable/

- **Goalkeeper -** `Goalkeeper`
- **Central Defender -** `Centre Back, Central Defensive Midfielder`
- **Wide Defender -** `Left Back, Right Back, Left Wing Back, Right Wing Back`

- **Central Midfielder -** `Central Midfielder`
- **Wide Midfielder -** `Left Midfielder, Right Midfielder`
- **Central Attacker -** `Central Attacking Midfielder, Centre Forward`
- **Wide Attacker -** `Left Winger, Right Winger, Left Forward, Right Forward`

The event and tracking data used in Chapter 3 were provided by Bepro Group Ltd. Inquiries regarding access to these data can be made at: https://www.bepro.ai/.

Appendix B

Appendix for Learning and Optimising Spatial Teamwork in Multi-Agent Teams

In this appendix, we provide additional information about various aspects of the work presented in Chapter 4 of this thesis. Section B.1 outlines the compute resources employed in the work. Section B.2 presents the model structure, hyperparameters, and inputs for the transition function described in Section 4.4. Finally, additional implementation details for Chapter 4 are shared in Section B.3.

B.1 Compute Resources

This work was implemented using Python 3.8.15. Our Monte Carlo tree search (MCTS) algorithm was evaluated on every on-ball pass and shot event in our dataset of 34 K League 1 games, comprising 34,112 events. For Experiments 1 and 3 in Chapter 4, we tested model performance using 1,000 simulations per event. For each event and defending player in the dataset, we executed the MCTS algorithm for thirty seconds, and the action with the highest value at the end of this period was selected. The linear programming algorithm (*Spatial teamwork*) typically takes between 30 and 90 seconds to execute.

Simulations were executed in parallel on the IRIDIS High Performance Computing Facility at the University of Southampton. Each cluster node was allocated 4GB of memory.

The event and tracking data used in Chapter 4 were provided by Bepro Group Ltd. For inquiries regarding access to these data, the company can be contacted directly at: `https://www.bepro.ai/`.

B.2 Transition Function Model

In Section 4.4, we introduced an agent location prediction model adapted from the *Agent Imputer* described in Chapter 3. This model processes state information to predict future player positions, which are then used within our state transition function. In this appendix subsection, we provide details of the curated input data used for training (Section B.2.1), followed by a description of the model architecture along with the implementation details of the baseline methods (Section B.2.2).

B.2.1 Input Data

The following per-player features are used as input to our player location prediction model, with a separate feature set constructed for each player in order to predict their future location:

- **Player location data -** `Past X location, Past Y location, Past X velocity, Past Y velocity`
- **Ball location data -** `Past X ball location, Past Y ball location, New X ball location, New Y ball location`
- **Contextual data -** `Event time difference, Player role, Team on the ball (boolean), Player on the ball (boolean)`

The numeric features are normalised using the `RobustScaler` from `scikit-learn`.[1] The event time difference is computed using a ball travel time model from the literature (Spearman, 2018). Player roles are one-hot encoded. Overall, this results in an input tensor of shape $(b, 22, 28)$, where b is the batch size (256 for this model), 22 corresponds to the number of players and 28 to the features (including encodings). The target is the player's next (x, y) location.

B.2.2 Model Structure and Training

The model structure matches the model introduced in Chapter 3, with the exception that we use an input sequence length of one instead of five. The graph component of the model is fully connected across players. The model is trained over 100 epochs with a learning rate of 0.0005 and a batch size (b) of 256. Model hyperparameters were selected empirically through trial and adjustment; no formal hyperparameter tuning strategy was applied.

[1] `https://scikit-learn.org/stable/.`

Machine learning baselines, used as comparators to the transition function model, are constructed using the `scikit-learn` package. The graph neural network (GNN) baseline extracts only the GNN component of the full model and is trained under the same setup as the complete model. Testing is conducted with approximately a 95/5 train/test split, corresponding to two fully unseen games for testing.

B.3 Further Implementation Details

In this section, we present additional implementation details, including the software packages used to form subteams (Section B.3.1) and the attack location prediction model used in Section 4.5.4 (Section B.3.2).

B.3.1 Subteam Formation

As described in Section 4.3.3, potential subteams are identified using proximity-based K-means clustering. At each timestep, clustering is based solely on agent positional data. The `scikit-learn` package is used for both K-means clustering and silhouette coefficient computation to assess cluster quality. To improve stability, each clustering run is initialised 50 times, keeping the solution that minimises inertia. This number of initialisations was determined empirically, balancing between cluster reliability and computational efficiency.

B.3.2 Attack Location Prediction Model

The attack location prediction model (introduced in Section 4.5.4) uses defender and subteam contributions as inputs to a neural network that predicts the most likely zone of the next pass or shot. The pitch is discretised into a 6×4 grid (24 zones), and model outputs correspond to probabilities over these zones. The model is implemented in Keras.[2] Its specification is summarised below.

Inputs (features) Subteam contribution features and individual defender contribution features, aligned to the 6×4 zone grid.

Model Architecture (Keras) `Conv2D(32, (3,3), activation='relu')` $\rightarrow$ `Flatten()` $\rightarrow$ `Dense(64, activation='relu')` $\rightarrow$ `Dense(24, activation='sigmoid')`

Targets Zone-wise probabilities over 24 pitch zones for the next pass or shot event.

Training setup 25 epochs, batch size = 64, learning rate = 0.01.

[2] https://keras.io/

Appendix C

Appendix for Evaluating Off-Ball Player Contributions

In this appendix section, we provide details on the implementation of our *GAPP* model presented in Chapter 5 (Section C.1), as well as information about the dataset used to validate our *GAPP* approach and the new defensive metrics for evaluating off-ball player contributions (Section C.2).

C.1 Implementation Details

In this section, we explain the compute resources used in this work (Section C.1.1), and give further details on our input feature preparation (Section C.1.2).

C.1.1 Compute Resources

This work was implemented using Python 3.8.15, with our *GAPP* model and baseline graph neural network method developed in PyTorch.[1] Model training was conducted on a remote graphics processing unit (GPU) service using an NVIDIA T4 GPU with 51 GB of video random access memory (VRAM), which was facilitated through Google Colab[2]. Training the *GAPP* model required a maximum of three hours for each fold. Inference for the *GAPP* model, used to compute the defender influence (DI) and defender performance (DP) metrics, was performed on a high-performance computing (HPC) facility. We utilised 360 parallel nodes, each with 4 GB of random access memomry (RAM), to handle the large-scale dataset used in our analysis. The specific HPC facility was the IRIDIS High Performance Computing Facility at the University of Southampton.

[1] `https://pytorch.org/`
[2] `https://colab.research.google.com/`

The XGBoost baseline was implemented using the `XGBoost` library.[3] The XGBoost model has the following hyperparameters: max_depth = 6, eta = 0.1, subsample = 0.8, colsample_bytree = 0.8. The Random forest model for the Dauxais (Dauxais and Gautrais, 2018) baseline was implemented using the `scikit-learn` library[4] and uses 200 trees (matching the original paper) with a max_depth of 6.

C.1.2 Feature Preparation

The categorical features in our input feature set to the *GAPP* model (Section 5.3.1) are represented as binary indicators, as all of them are boolean variables. For numeric features, we apply normalisation using the `StandardScaler` from the `scikit-learn` library.

C.2 Dataset Details

In this section, we provide additional details about the dataset used in this work. As noted in Section 6.7, the dataset consists of 306 games from the 2023/24 English Premier League (EPL) season. However, not all games from the season are included, and we normalise for the number of events each team participated in. Table C.1 lists the teams included in the dataset along with the corresponding number of games.

TABLE C.1: Game count for each team in our EPL 2023/24 dataset.

Team	Game Count	Team	Game Count
Liverpool	33	Chelsea	31
West Ham	33	Brentford	31
Burnley	33	Man United	31
Aston Villa	32	Nottingham Forest	30
Fulham	32	Arsenal	30
Man City	32	Everton	30
Luton Town	32	Tottenham	29
Bournemouth	32	Newcastle	28
Crystal Palace	31	Sheffield United	27
Brighton	31	Wolves	24

Our dataset includes both on-ball events and tracking data for these games.

Tracking Dataset The tracking dataset used for training and validating our *GAPP* model, as well as for computing the DI and DP performance metrics, records the locations of all players and the ball at a frequency of 30 Hz. Each tracking frame also

[3] https://xgboost.readthedocs.io/en/stable/python/
[4] https://scikit-learn.org/stable/

includes a frame number, timestamp, and unique identifiers for each player and team. Our datasets use relative coordinates, ranging from the bottom-left corner of the pitch (0,0) to the top-right corner (105,68), from the perspective of the attacking team. The attacking direction is configured to always be from left to right.

Event Dataset Our dataset includes the timestamp and location of each on-ball event, along with the event type and the player responsible for the action. The possible event types can be categorised into the following:

- **Attacking events -** `Ball carry, Cross, Pass, Shot`
- **Defensive Events -** `Challenge, Clearance, Rebound`

The outcome of an event, such as the location of the next pass, is determined by the location of the subsequent on-ball event in the dataset.

The event and tracking data used in Chapter 5 were provided by Gradient Sports. For inquiries regarding access to these data, the company can be contacted directly at: `https://www.gradientsports.com/`.

Appendix D

Appendix for Optimising Short- and Long-Term Team Selection

In this appendix section, we describe the compute resources used for the team selection model in Chapter 6 (Section D.1). The model hyperparameter ranges used for hyperparameter tuning are provided in Section D.2. We provide further details on the datasets used in Chapter 6 in Section D.3. Finally, we present an overview of all injury feature contributions in Section D.4.

D.1 Compute Resources

This work was implemented in Python 3.8.15. For the match reward model (Section 6.5), player-level features were standardised using the `StandardScaler` from the `scikit-learn`[1] package. Match outcomes were predicted using the `PoissonRegressor` from `scikit-learn`. The pre-game policy employed the `LogisticRegression` model from `scikit-learn`, while the in-game policy used the same class with the `multi_class` parameter set to `multinomial` for multi-class classification.

For the injury probability model (Section 6.4), numerical features were also standardised using `StandardScaler`. The model itself was implemented with `XGBClassifier` from the xgboost[2] package. Machine learning baselines used for comparison were constructed using `scikit-learn`. All experimental graphs were generated using the `matplotlib`[3] library.

As described in Section 6.7.6, our experimental analysis involves simulating 500 seasons for each team and strategy in the 2018/19 English Premier League (EPL) season. Given

[1] `https://scikit-learn.org/stable/`
[2] `https://xgboost.readthedocs.io/en/stable/python/`
[3] `https://matplotlib.org/`

the computational complexity of our problem and the resulting runtime demands of Monte Carlo tree search (MCTS), these simulations were executed in parallel on a compute cluster using dual 2.0 GHz Intel Xeon Gold 6138 processors, with each simulation allocated ∼4GB of random access memory (RAM). The specific high performance computing facility was the IRIDIS High Performance Computing Facility at the University of Southampton. Each simulated season follows a standardised procedure:

- For every match, the pre-game MCTS algorithm is run for three minutes to select the lineup with the highest estimated value, considering only actions that have been selected at least 200 times to reduce the impact of noise.
- During each match, at every timestep, MCTS is used with a 30-second search to determine the number of substitutions, requiring that an action has been selected more than the total number of selections divided by the number of possible actions plus one.
- Once the number of substitutions is determined, a two-minute MCTS search is performed to select the specific player to substitute, considering only the top 20 most frequently selected actions.

D.2 Model Hyperparameters

Optimal hyperparameters for the injury probability model and the comparative baselines presented in Section 6.4 were found using grid-search cross-validation. Below, we list the hyperparameter ranges explored for each baseline, along with the best-performing values selected.

XGBoost **Grid search ranges:**

- `max_depth`: [3, 4, 5]
- `learning_rate` (eta): [0.01, 0.015, 0.02, 0.025, 0.03, 0.04, 0.05]
- `n_estimators`: [100, 200, 400, 600]
- `alpha`: [1, 5, 10, 20, 50]

Optimal values: max_depth = 3, learning_rate = 0.01, n_estimators = 400, alpha = 20

Logistic Regression **Grid search ranges:**

- `C`: [0.1, 1, 10]
- `solver`: ['liblinear', 'newton-cholesky']

Optimal values: C = 0.1, solver = 'newton-cholesky'

Random Forest **Grid search ranges:**

- `max_depth`: [3, 4, 5]
- `n_estimators`: [100, 200, 300, 400]
- `min_samples_split`: [2, 3, 4]

Optimal values: max_depth = 4, n_estimators = 300, min_samples_split = 4

Neural Network **Grid search ranges:**

- `hidden_layer_sizes`: [(10,), (25,), (50,), (10, 10), (25, 25), (50,50)]
- `learning_rate_init`: [0.01,0.02,0.03,0.04,0.05]
- `alpha`: [0.001, 0.01, 0.1, 1]

Optimal values: hidden_layer_sizes = (10,), learning_rate_init = 0.05, alpha = 1

For the pre-game real-world policy logistic regression model (Section 6.6.3.1), we used the default logistic regression parameters and `lbfgs` solver, as these provided robust performance in preliminary experiments. For the in-game real-world policy, the multi-class logistic regression models (Section 6.6.3.2) were implemented using the `multinomial` option for the `multi_class` parameter and the `lbfgs` solver. The Poisson regression models used for game and timestep outcome prediction, as well as the baseline Maher models, (Section 6.5) were implemented with `PoissonRegressor` using `alpha` = 1×10^{-9} and `tol` = 1×10^{-9}.

For our MCTS algorithm, the exploration constant c_{puct} was set to 2 for both the pre-game MCTS and the substitution number stage of the in-game MCTS, while a higher value of $c_{\text{puct}} = 200$ was used for the player substitution selection stage. These values were determined empirically by monitoring the convergence of the Q-values across multiple runs and selecting the values that yielded the most stable and high-quality suggestions. The time budgets allocated for each MCTS run (see Section D.1) were also chosen empirically, with the goal of achieving a practical balance between computational efficiency and the quality of the MCTS suggestions.

D.3 Dataset Details

In this paper, we use three key data sources:

On-ball Events Data We utilise detailed on-ball event data from StatsBomb for the 2017/18 and 2018/19 EPL seasons. This dataset includes match schedules, player selections, minutes played, player positional roles (e.g., Goalkeeper, Midfielder), team formations, final scorelines, and granular on-ball actions (such as shots, passes, and dribbles), along with their timing and location on the pitch. We use this data to extract features for both the injury probability model and the match prediction model. Additionally, the match schedule information is used to replicate the real-world sequence and timing of games within our Markov decision process (MDP) framework. Event data from StatsBomb can be publicly accessed at: `https://statsbomb.com/`.

Injury Data We collect injury data from Transfermarkt for all players participating in the 2017/18 and 2018/19 EPL seasons. This dataset covers each player's injury history, including injury type, date, and duration. We process this data to generate injury-related features for the injury probability model (see Section 6.4) and to create the target variable for model training (i.e., whether a player was injured for each game). This injury data is publicly accessible from Transfermarkt at: `https://www.transfermarkt.co.uk/`.

Wage Data To assess the financial implications of injuries in our experimental evaluation (Section 6.7.7), we use estimated EPL wage data for all players from the 2018/19 season, sourced from Capology. This wage data can be publicly accessed at: `https://www.capology.com/`.

D.4 Injury Feature Importance

Figure D.1 shows the contribution, denoted as the Shapley Additive Explanations (SHAP) value, of all features used in the player injury probability model towards player injury risks across our whole validation dataset.

Clubs can use this information to explain why players are at high injury risk. SHAP values can be extended to single-player predictions, allowing clubs to understand which features contribute to a high risk for a specific player in a particular game. Additionally, clubs may use this information to adapt training and recovery sessions to mitigate the factors that are contributing most to the player's high injury risk.

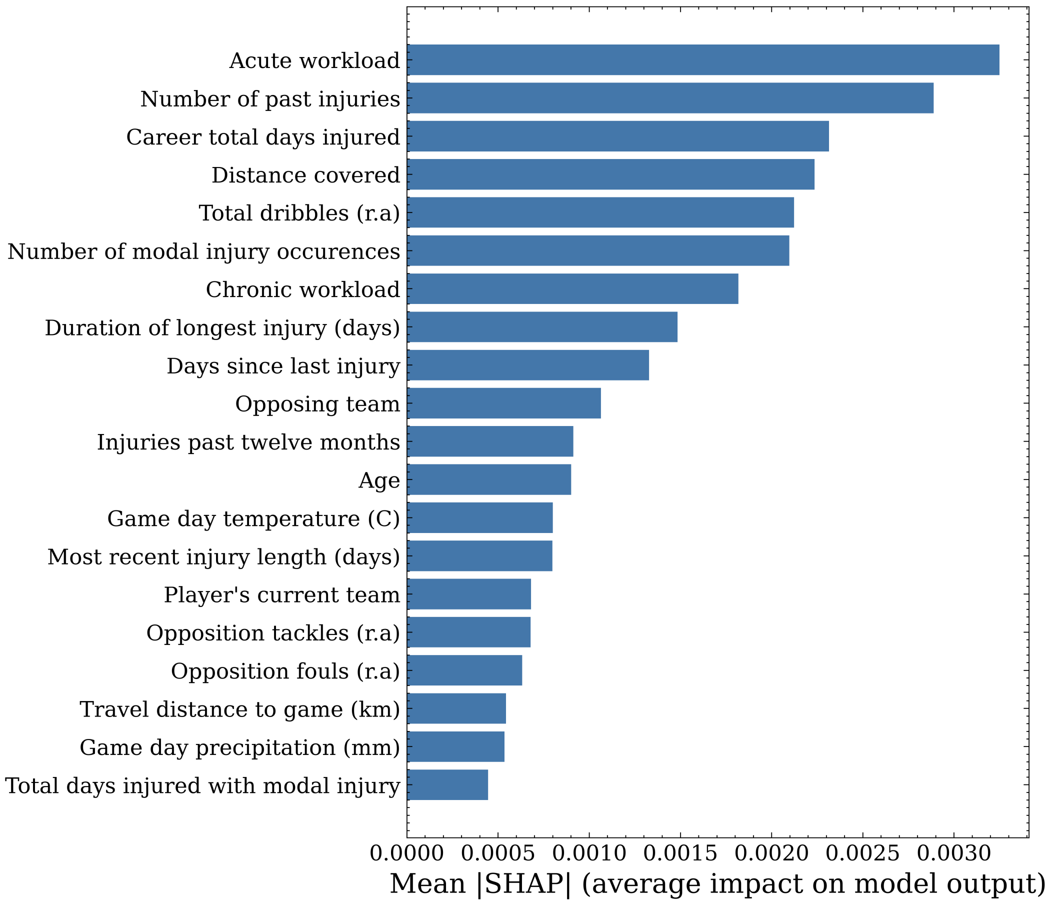

FIGURE D.1: SHAP values of all features for the injury risk model.

References

Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.

Guillermo Alcantara-Jimenez and Julio B Clempner. Repeated stackelberg security games: Learning with incomplete state information. *Reliability Engineering & System Safety*, 195:106695, 2020.

Sofia Amador, Steven Okamoto, and Roie Zivan. Dynamic multi-agent task allocation with spatial and temporal constraints. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multiagent Systems*, pages 1495–1496, 2014.

Aris Anagnostopoulos, Luca Becchetti, Carlos Castillo, Aristides Gionis, and Stefano Leonardi. Online team formation in social networks. In *Proceedings of the 21st international conference on World Wide Web*, pages 839–848, 2012.

Tone Arnold and Ulrich Schwalbe. Dynamic coalition formation and the core. *Journal of economic behavior & organization*, 49(3):363–380, 2002.

Ashay Aswale and Carlo Pinciroli. Heterogeneous coalition formation and scheduling with multi-skilled robots. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5402–5409. IEEE, 2023.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47:235–256, 2002.

Llorenç Badiella, Pedro Puig, Carlos Lago-Peñas, and Martí Casals. Influence of red and yellow cards on team performance in elite soccer. *Annals of Operations Research*, 325(1): 149–165, 2023.

Bing Bai, Jian Liang, Guanhua Zhang, Hao Li, Kun Bai, and Fei Wang. Why attentions may not be interpretable? In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 25–34, 2021.

Chris AB Baker, Sarvapali Ramchurn, WT Luke Teacy, and Nicholas R Jennings. Planning search and rescue missions for uav teams. In *Proceedings of the Twenty-second European Conference on Artificial Intelligence*, pages 1777–1778, 2016.

Inci M Baytas, Cao Xiao, Xi Zhang, Fei Wang, Anil K Jain, and Jiayu Zhou. Patient subtyping via time-aware lstm networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 65–74, 2017.

Ryan Beal, Timothy J Norman, and Sarvapali D Ramchurn. Artificial intelligence for team sports: a survey. *The Knowledge Engineering Review*, 34, 2019.

Ryan Beal, Georgios Chalkiadakis, Timothy J Norman, and Sarvapali D Ramchurn. Optimising game tactics for football. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 141–149, 2020a.

Ryan Beal, Narayan Changder, Timothy Norman, and Sarvapali Ramchurn. Learning the value of teamwork to form efficient teams. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7063–7070, 2020b.

Ryan Beal, Georgios Chalkiadakis, Timothy J Norman, and Sarvapali D Ramchurn. Optimising long-term outcomes using real-world fluent objectives: An application to football. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*, pages 196–204, 2021.

Ryan James Beal. *Artificial intelligence in team sports*. PhD thesis, University of Southampton, 2022.

Michael Beetz, Bernhard Kirchlechner, and Martin Lames. Computerized real-time analysis of football games. *IEEE pervasive computing*, 4(3):33–39, 2005.

Michael Beetz, Suat Gedikli, Jan Bandouch, Bernhard Kirchlechner, Nico V Hoyningen-Huene, and Alexander Perzylo. Visually tracking football games based on tv broadcasts. In *Proceedings of the 20th international joint conference on Artifical intelligence*, pages 2066–2071, 2007.

Håkan Bengtsson, Jan Ekstrand, and Martin Hägglund. Muscle injury rates in professional football increase with fixture congestion: an 11-year follow-up of the uefa champions league injury study. *British journal of sports medicine*, 47(12):743–747, 2013.

Aïcha BenTaieb and Ghassan Hamarneh. Predicting cancer with a recurrent visual attention model for histopathology images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 129–137, 2018.

Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

Alina Bialkowski, Patrick Lucey, Peter Carr, Yisong Yue, and Iain Matthews. Win at home and draw away: Automatic formation analysis highlighting the differences in home and away team behaviors. In *Proceedings of the 8th MIT Sloan Sports Analytics Conference*, pages 1–7, 2014a.

Alina Bialkowski, Patrick Lucey, Peter Carr, Yisong Yue, Sridha Sridharan, and Iain Matthews. Identifying team style in soccer using formations learned from spatiotemporal tracking data. In *IEEE international conference on data mining workshop*, pages 9–14. IEEE, 2014b.

Alina Bialkowski, Patrick Lucey, Peter Carr, Yisong Yue, Sridha Sridharan, and Iain Matthews. Large-scale analysis of soccer matches using spatiotemporal tracking data. In *IEEE international conference on data mining*, pages 725–730. IEEE, 2014c.

Filippo Bistaffa, Alessandro Farinelli, Georgios Chalkiadakis, and Sarvapali D Ramchurn. A cooperative game-theoretic approach to the social ridesharing problem. *Artificial Intelligence*, 246:86–117, 2017.

Samuel S Blackman. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine*, 19(1):5–18, 2004.

Laura Bowen, Aleksander Stephan Gross, Mo Gimpel, Stewart Bruce-Low, and Francois-Xavier Li. Spikes in acute: chronic workload ratio (acwr) associated with a 5–7 times greater injury rate in english premier league football players: a comprehensive 3-year study. *British journal of sports medicine*, 54(12):731–738, 2020.

Lotte Bransen and Jan Van Haaren. Player chemistry: Striving for a perfectly balanced soccer team. In *Proceedings of the 14th MIT Sloan Sports Analytics Conference*, 2020.

Marc Brittain and Peng Wei. Autonomous air traffic controller: A deep multi-agent reinforcement learning approach. In *ICML Workshop: RL for Real Life*, 2019.

Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.

Blazej Bulka, Matthew Gaston, and Marie Desjardins. Local strategy learning in networked multi-agent team formation. *Autonomous Agents and Multi-Agent Systems*, 15: 29–45, 2007.

Luca Capezzuto, Danesh Tarapore, and Sarvapali Ramchurn. Anytime and efficient coalition formation with spatial and temporal constraints. In *European Conference on Multi-Agent Systems*, pages 589–606, 2020.

Luca Capezzuto, Danesh Tarapore, and Sarvapali D Ramchurn. Large-scale, dynamic and distributed coalition formation with spatial and temporal constraints. In *European Conference on Multi-Agent Systems*, pages 108–125. Springer, 2021.

Albert V Carron, Steven R Bray, and Mark A Eys. Team cohesion and team success in sport. *Journal of sports sciences*, 20(2):119–126, 2002.

Julen Castellano, David Álvarez, Bruno Figueira, Diogo Coutinho, and Jaime Sampaio. Identifying the effects from the quality of opposition in a football team positioning strategy. *International Journal of Performance Analysis in Sport*, 13(3):822–832, 2013.

Tristan Cazenave and Nicolas Jouandeau. On the parallelization of uct. In *Computer games workshop*, 2007.

Georgios Chalkiadakis and Craig Boutilier. Sequentially optimal repeated coalition formation under uncertainty. *Autonomous Agents and Multi-Agent Systems*, 24:441–484, 2012.

Guillaume M Jb Chaslot, Mark HM Winands, H Jaap van den Herik, Jos WHM Uiterwijk, and Bruno Bouzy. Progressive strategies for monte-carlo tree search. *New Mathematics and Natural Computation*, 4(03):343–357, 2008.

Jie Chen, Yuqian Guo, Zhifeng Qiu, Bin Xin, Qing-Shan Jia, and Weihua Gui. Multiagent dynamic task assignment based on forest fire point model. *IEEE Transactions on Automation Science and Engineering*, 19(2):833–849, 2021.

Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD conference on knowledge discovery & data mining*, pages 785–794, 2016.

Jonathan Cohen and Abdel-Illah Mouaddib. Monte-carlo planning for team re-formation under uncertainty: Model and properties. In *IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 458–465. IEEE, 2018.

Cloe Cummins, Mitchell Welch, Brendan Inkster, Balin Cupples, Dan Weaving, Ben Jones, Doug King, and Aron Murphy. Modelling the relationships between volume, intensity and injury-risk in professional rugby league players. *Journal of science and medicine in sport*, 22(6):653–660, 2019.

Felipe Dambroz, Filipe Manuel Clemente, and Israel Teoldo. The effect of physical fatigue on the performance of soccer players: A systematic review. *PloS one*, 17(7): e0270099, 2022.

Yann Dauxais and Clément Gautrais. Predicting pass receiver in football using distance based features. In *MLSA 2018-5th Workshop on Machine Learning and Data Mining for Sports Analytic of ECML/PKDD*, pages 1–7, 2018.

Jesse Davis, Lotte Bransen, Laurens Devos, Arne Jaspers, Wannes Meert, Pieter Robberechts, Jan Van Haaren, and Maaike Van Roy. Methodology and evaluation in sports analytics: challenges, approaches, and lessons learned. *Machine Learning*, 113 (9):6977–7010, 2024.

Anthony Dearden, Yiannis Demiris, and Oliver Grau. Tracking football player movement from a single moving camera using particle filters. In *The 3rd European Conference on Visual Media Production*, pages 29–37, 2006.

Tom Decroos, Jan Van Haaren, Vladimir Dzyuba, and Jesse Davis. Starss: a spatio-temporal action rating system for soccer. In *Machine Learning and Data Mining for Sports Analytics ECML/PKDD 2017 workshop*, volume 1971, pages 11–20. Springer, 2017.

Tom Decroos, Jan Van Haaren, and Jesse Davis. Automatic discovery of tactics in spatio-temporal soccer match data. In *Proceedings of the 24th ACM SIGKDD conference on knowledge discovery & data mining*, pages 223–232, 2018.

Tom Decroos, Lotte Bransen, Jan Van Haaren, and Jesse Davis. Actions speak louder than goals: Valuing player actions in soccer. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1851–1861, 2019.

Tom Decroos, Jesse Davis, et al. Valuing on-the-ball actions in soccer: a critical comparison of xt and vaep. In *Proceedings of the AAAI-20 Workshop on Artifical Intelligence in Team Sports*, pages 1–8. AI in Team Sports Organising Committee, 2020.

Alexandre Dellal, Carlos Lago-Peñas, Ezequiel Rey, Karim Chamari, and Emmanuel Orhant. The effects of a congested fixture period on physical performance, technical activity and injury rate during matches in a professional soccer team. *British journal of sports medicine*, 49(6):390–394, 2015.

Pedro Demasi and J de O Adriano. Online coevolution for action games. *Int. J. Intell. Games & Simulation*, 2(2):80–88, 2003.

Emir Demirović, Nicolas Schwind, Tenda Okimoto, and Katsumi Inoue. Recoverable team formation: Building teams resilient to change. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, pages 1362–1370, 2018.

Ali Dorri, Salil S Kanhere, and Raja Jurdak. Multi-agent systems: A survey. *IEEE Access*, 6:28573–28593, 2018.

Jan Ekstrand and Benno M Nigg. Surface-related injuries in soccer. *Sports medicine*, 8: 56–62, 1989.

Greg Everett, Ryan Beal, Tim Matthews, Tim Norman, and Gopal Ramchurn. Contextual expected threat using spatial event data. In *StatsBomb Innovation in Football Conference, London, United Kingdom*, 2023a.

Gregory Everett, Ryan J Beal, Tim Matthews, Joseph Early, Timothy J Norman, and Sarvapali D Ramchurn. Inferring player location in sports matches: Multi-agent spatial imputation from limited observations. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 1643–1651, 2023b.

Gregory Everett, Ryan Beal, Tim Matthews, Timothy J Norman, and Sarvapali D Ramchurn. The strain of success: A predictive model for injury risk mitigation and team success in soccer. In *Proceedings of the 18th MIT Sloan Sports Analytics Conference*, 2024.

Gregory Everett, Ryan J Beal, Tim Matthews, Timothy J Norman, and Sarvapali D Ramchurn. Evaluating defensive influence in multi-agent systems using graph attention networks. In *IEEE 12th International Conference on Data Science and Advanced Analytics*, 2025a.

Gregory Everett, Ryan J Beal, Tim Matthews, Timothy J Norman, and Sarvapali D Ramchurn. Optimising spatial teamwork under uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23168–23176, 2025b.

Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, Andrew Lemieux, et al. Deploying paws: Field optimization of the protection assistant for wildlife security. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 16, pages 3966–3973, 2016.

Javier Fernandez and Luke Bornn. Wide open spaces: A statistical technique for measuring space creation in professional soccer. In *Proceedings of the 12th MIT Sloan Sports Analytics Conference*, 2018.

Javier Fernández, Luke Bornn, and Dan Cervone. Decomposing the immeasurable sport: A deep learning expected possession value framework for soccer. In *Proceedings of the 13th MIT Sloan Sports Analytics Conference*, 2019.

Javier Fernández, Luke Bornn, and Daniel Cervone. A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. *Machine Learning*, 110(6):1389–1427, 2021.

Leander Forcher, Stefan Altmann, Leon Forcher, Darko Jekauc, and Matthias Kempe. The use of player tracking data to analyze defensive play in professional soccer-a scoping review. *International Journal of Sports Science & Coaching*, 17(6):1567–1592, 2022.

Leander Forcher, Tobias Beckmann, Oliver Wohak, Christian Romeike, Ferdinand Graf, and Stefan Altmann. Prediction of defensive success in elite soccer using machine learning-tactical analysis of defensive play using tracking data and explainable ai. *Science and Medicine in Football*, 8(4):317–332, 2024.

Alexander Franks, Andrew Miller, Luke Bornn, and Kirk Goldsberry. Characterizing the spatial structure of defensive skill in professional basketball. *The Annals of Applied Statistics*, 9(1):94–121, 2015.

Andrea Galassi, Marco Lippi, and Paolo Torroni. Attention in natural language processing. *IEEE transactions on neural networks and learning systems*, 32(10):4291–4308, 2020.

Matthew E Gaston and Marie DesJardins. Agent-organized networks for dynamic team formation. In *Proceedings of the 4th International Conference on Autonomous Agents and Multiagent Systems*, pages 230–237, 2005.

Floris Goes, Edgar Schwarz, Marije Elferink-Gemser, Koen Lemmink, and Michel Brink. A risk-reward assessment of passing decisions: comparison between positional roles using tracking data from professional men's soccer. *Science and Medicine in Football*, pages 1–9, 2021a.

Floris R Goes, Matthias Kempe, Laurentius A Meerhoff, and Koen APM Lemmink. Not every pass can be an assist: a data-driven model to measure pass effectiveness in professional soccer matches. *Big data*, 7(1):57–70, 2019.

Floris R Goes, Michel S Brink, Marije T Elferink-Gemser, Matthias Kempe, and Koen APM Lemmink. 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):523–532, 2021b.

FR Goes, M Kempe, J van Norel, and KAPM Lemmink. Modelling team performance in soccer using tactical features derived from position tracking data. *IMA Journal of Management Mathematics*, 32(4):519–533, 2021c.

Haifeng Gong, Jack Sim, Maxim Likhachev, and Jianbo Shi. Multi-hypothesis motion planning for visual object tracking. In *International Conference on Computer Vision*, pages 619–626. IEEE, 2011.

Sam Gregory, Sam Robertson, Robert Aughey, Bartholomew Spencer, and Jeremy Alexander. Assigning goal-probability value to high intensity runs in football. *Plos one*, 19(9): e0308749, 2024.

Tyler Gunn and John Anderson. Effective task allocation for evolving multi-robot teams in dangerous environments. In *IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, volume 2, pages 231–238. IEEE, 2013.

Tyler Gunn and John Anderson. Dynamic heterogeneous team formation for robotic urban search and rescue. *Journal of Computer and System Sciences*, 81(3):553–567, 2015.

Monoem Haddad, Georgios Stylianides, Leo Djaoui, Alexandre Dellal, and Karim Chamari. Session-rpe method for training load monitoring: validity, ecological usefulness, and influencing factors. *Frontiers in neuroscience*, 11:612, 2017.

Martin Hägglund, Markus Waldén, Henrik Magnusson, Karolina Kristenson, Håkan Bengtsson, and Jan Ekstrand. Injuries affect team performance negatively in professional football: an 11-year follow-up of the uefa champions league injury study. *British journal of sports medicine*, 47(12):738–742, 2013.

Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

Jack Harmer, Linus Gisslén, Jorge del Val, Henrik Holst, Joakim Bergdahl, Tom Olsson, Kristoffer Sjöö, and Magnus Nordin. Imitation learning with concurrent actions in 3d games. In *IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8. IEEE, 2018.

Sandro Hauri, Nemanja Djuric, Vladan Radosavljevic, and Slobodan Vucetic. Multi-modal trajectory prediction of nba players. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1640–1649, 2021.

Jan Held, Anthony Cioppa, Silvio Giancola, Abdullah Hamdi, Christel Devue, Bernard Ghanem, and Marc Van Droogenbroeck. Towards an ai-powered video assistant referee system (vars) for association football. In *Scandinavian Conference on Image Analysis*, pages 295–309. Springer, 2025.

Nobuyoshi Hirotsu and Michael Wright. Using a markov process model of an association football match to determine the optimal timing of substitution and tactical decisions. *Journal of the Operational Research Society*, 53(1):88–96, 2002.

Nobuyoshi Hirotsu, Masamitsu Ito, Chikara Miyaji, Koji Hamano, and Azuma Taguchi. Modeling tactical changes of formation in association football as a non-zero-sum game. *Journal of Quantitative Analysis in Sports*, 5(3):1–15, 2009.

Kieran Howle, Adam Waterson, and Rob Duffield. Injury incidence and workloads during congested schedules in football. *International journal of sports medicine*, 41(02): 75–81, 2020.

Jinlong Hu, Lijie Cao, Tenghui Li, Shoubin Dong, and Ping Li. Gat-li: a graph attention network based learning and interpreting method for functional brain network classification. *BMC bioinformatics*, 22:1–20, 2021.

Billy T Hulin, Tim J Gabbett, Daniel W Lawson, Peter Caputi, and John A Sampson. The acute: chronic workload ratio predicts injury: high chronic workload may decrease injury risk in elite rugby league players. *British journal of sports medicine*, 50(4):231–236, 2016.

Ahmed Hussein, Eyad Elyan, Mohamed Medhat Gaber, and Chrisina Jayne. Deep imitation learning for 3d navigation tasks. *Neural computing and applications*, 29(7): 389–404, 2018.

Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2375–2384, 2019.

Susanne Jauhiainen, Jukka-Pekka Kauppi, Tron Krosshaug, Roald Bahr, Julia Bartsch, and Sami Äyrämö. Predicting acl injury using machine learning on data from an extensive screening test battery of 880 female elite athletes. *The American Journal of Sports Medicine*, 50(11):2917–2924, 2022.

Chris Jones and K Barber. Team formation strategies in a dynamic large-scale environment. *Massively Multi-Agent Technology*, pages 92–106, 2008.

Bilal Kartal, Ernesto Nunes, Julio Godoy, and Maria Gini. Monte carlo tree search for multi-robot task allocation. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 4222–4223, 2016.

Ryota Katayanagi and Toshiharu Sugawara. Efficient team formation based on learning and reorganization and influence of communication delay. In *IEEE 11th International Conference on Computer and Information Technology*, pages 563–570. IEEE, 2011.

Matthias Kempe. Machine learning for injury prediction. In *Artificial Intelligence and Machine Learning in Sports Science*, pages 275–283. Springer, 2025.

Sujeong Kim, Stephen J Guy, Wenxi Liu, Rynson WH Lau, Ming C Lin, and Dinesh Manocha. Predicting pedestrian trajectories using velocity-space reasoning. In *Algorithmic Foundations of Robotics X: Proceedings of the Tenth Workshop on the Algorithmic Foundations of Robotics*, pages 609–623. Springer, 2013.

Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Stephanie A Kovalchik. Player tracking data in sports. *Annual Review of Statistics and Its Application*, 10(1):677–697, 2023.

Robert Kraemer and Karsten Knobloch. A soccer-specific balance training program for hamstring muscle and patellar and achilles tendon injuries: an intervention study in premier league female soccer. *The American journal of sports medicine*, 37(7):1384–1393, 2009.

Kristen L Kucera, Stephen W Marshall, Donald T Kirkendall, PM Marchak, and William E Garrett. Injury history as a risk factor for incident injury in youth soccer. *British journal of sports medicine*, 39(7):462–462, 2005.

Karol Kurach, Anton Raichuk, Piotr Stańczyk, Michał Zając, Olivier Bachem, Lasse Espeholt, Carlos Riquelme, Damien Vincent, Marcin Michalski, Olivier Bousquet, and Sylvain Gelly. Google research football: A novel reinforcement learning environment. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:4501–4510, 04 2020. .

Hoang Le, Peter Carr, Yisong Yue, and Patrick Lucey. Data-driven ghosting using deep imitation learning. In *Proceedings of the 11th MIT Sloan Sports Analytics Conference*, 2017a.

Hoang M Le, Yisong Yue, Peter Carr, and Patrick Lucey. Coordinated multi-agent imitation learning. In *International Conference on Machine Learning*, pages 1995–2003. PMLR, 2017b.

Namhoon Lee and Kris M Kitani. Predicting wide receiver trajectories in american football. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016.

Daniel Link, Steffen Lang, and Philipp Seidenschwarz. Real time quantification of dangerousity in football using spatiotemporal tracking data. *PloS one*, 11(12):e0168768, 2016.

Da Liu, Liqian Dou, Ruilong Zhang, Xiuyun Zhang, and Qun Zong. Multi-agent reinforcement learning-based coordinated dynamic task allocation for heterogenous uavs. *IEEE Transactions on Vehicular Technology*, 72(4):4372–4383, 2022a.

Jia Liu, Xiaofeng Tong, Wenlong Li, Tao Wang, Yimin Zhang, and Hongqi Wang. Automatic player detection, labeling and tracking in broadcast soccer video. *Pattern recognition letters*, 30(2):103–113, 2009.

Xue Liu, Haokai Sun, and Rongkun Ye. Prediction of crowd flow in key areas based on a spatial-temporal convolutional model. In *IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1161–1166. IEEE, 2022b.

Yibing Liu, Haoliang Li, Yangyang Guo, Chenqi Kong, Jing Li, and Shiqi Wang. Rethinking attention-model explainability through faithfulness violation test. In *International Conference on Machine Learning*, pages 13807–13824. PMLR, 2022c.

Sergio Llana, Pau Madrero, Javier Fernández, and FC Barcelona. The right place at the right time: Advanced off-ball metrics for exploiting an opponent's spatial weaknesses in soccer. In *Proceedings of the 14th MIT Sloan Sports Analytics Conference*, 2020.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Yining Lu, Ayoosh Pareek, Ophelie Z Lavoie-Gagne, Enrico M Forlenza, Bhavik H Patel, Anna K Reinholz, Brian Forsythe, and Christopher L Camp. Machine learning for predicting lower extremity muscle strain in national basketball association athletes. *Orthopaedic Journal of Sports Medicine*, 10(7):23259671221111742, 2022.

Patrick Lucey, Alina Bialkowski, Peter Carr, Eric Foote, and Iain Matthews. Characterizing multi-agent team behavior from partial team tracings: Evidence from the english premier league. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, pages 1387–1393, 2012.

Patrick Lucey, Alina Bialkowski, Mathew Monfort, Peter Carr, and Iain Matthews. Quality vs quantity: Improved shot prediction in soccer using strategic features from spatiotemporal data. In *Proceedings of the 9th MIT Sloan Sports Analytics Conference*, 2015.

Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

Yudong Luo. *Inverse reinforcement learning for team sports: Valuing actions and players*. PhD thesis, Applied Sciences: School of Computing Science, 2020.

Michael J Maher. Modelling association football scores. *Statistica Neerlandica*, 36(3): 109–118, 1982.

Aritra Majumdar, Rashid Bakirov, Dan Hodges, Suzanne Scott, and Tim Rees. Machine learning for understanding and predicting injuries in football. *Sports Medicine-Open*, 8 (1):1–10, 2022.

Shane Malone, Adam Owen, Matt Newton, Bruno Mendes, Kieran D Collins, and Tim J Gabbett. The acute: chonic workload ratio in relation to injury risk in professional soccer. *Journal of science and medicine in sport*, 20(6):561–565, 2017.

Mehrtash Manafifard, Hamid Ebadi, and H Abrishami Moghaddam. A survey on player tracking in soccer videos. *Computer Vision and Image Understanding*, 159:19–46, 2017.

Francesco Marchetti, Federico Becattini, Lorenzo Seidenari, and Alberto Del Bimbo. Multiple trajectory prediction of moving agents with memory augmented networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):6688–6702, 2020.

Tim Matthews, Sarvapali Ramchurn, and Georgios Chalkiadakis. Competing with humans at fantasy football: Team formation in large partially-observable domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, pages 1394–1400, 2012.

Desmond McEwan and Mark R Beauchamp. Teamwork in sport: A theoretical and integrative review. *International Review of Sport and Exercise Psychology*, 7(1):229–250, 2014.

James McInerney, Alex Rogers, and Nicholas R Jennings. Improving location prediction services for new users with probabilistic latent semantic analysis. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 906–910, 2012.

James McInerney, Jiangchuan Zheng, Alex Rogers, and Nicholas R Jennings. Modelling heterogeneous location habits in human populations for location prediction under data sparsity. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 469–478, 2013.

Nazanin Mehrasa, Yatao Zhong, Frederick Tung, Luke Bornn, and Greg Mori. Deep learning of player trajectory representations for team activity analysis. In *Proceedings of the 11th MIT Sloan Sports Analytics Conference*, 2018.

Tiago Mendes-Neves, João Mendes-Moreira, and Rosaldo JF Rossetti. A data-driven simulator for assessing decision-making in soccer. In *EPIA Conference on Artificial Intelligence*, pages 687–698. Springer, 2021.

Charbel Merhej, Ryan J Beal, Tim Matthews, and Sarvapali Ramchurn. 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*, pages 3394–3403, 2021.

Alessio Monti, Alessia Bertugli, Simone Calderara, and Rita Cucchiara. Dag-net: Double attentive graph neural network for trajectory forecasting. In *25th international conference on pattern recognition (ICPR)*, pages 2551–2558. IEEE, 2021.

Dolev Mutzari, Jiarui Gan, and Sarit Kraus. Coalition formation in multi-defender security games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5603–5610, 2021.

Bret R Myers. A proposed decision rule for the timing of soccer substitutions. *Journal of Quantitative Analysis in Sports*, 8(1), 2012.

Claude Nadeau and Yoshua Bengio. Inference for the generalization error. *Advances in neural information processing systems*, 12, 1999.

Huda Dheyauldeen Najeeb and Rana Fareed Ghani. A survey on object detection and tracking in soccer videos. *MJPS*, 8(1), 2021.

Hiroshi Nakahara, Kazushi Tsutsui, Kazuya Takeda, and Keisuke Fujii. Action valuation of on-and off-ball soccer players based on multi-agent deep reinforcement learning. *IEEE Access*, 11:131237–131244, 2023.

Nobuyasu Nakano, Tetsuro Sakura, Kazuhiro Ueda, Leon Omura, Arata Kimura, Yoichi Iino, Senshi Fukashiro, and Shinsuke Yoshioka. Evaluation of 3d markerless motion capture accuracy using openpose with multiple video cameras. *Frontiers in sports and active living*, 2:50, 2020.

Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. A comprehensive overview of large language models. *ACM Transactions on Intelligent Systems and Technology*, 16(5):1–72, 2025.

Thanh H Nguyen, Yongzhao Wang, Arunesh Sinha, and Michael P Wellman. Deception in finitely repeated security games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2133–2140, 2019.

Zhaoyang Niu, Guoqiang Zhong, and Hui Yu. A review on the attention mechanism of deep learning. *Neurocomputing*, 452:48–62, 2021.

Itsuki Noda, Hitoshi Matsubara, Kazuo Hiraki, and Ian Frank. Soccer server: A tool for research on multiagent systems. *Applied Artificial Intelligence*, 12(2-3):233–250, 1998a.

Itsuki Noda, Sho'ji Suzuki, Hitoshi Matsubara, Minoru Asada, and Hiroaki Kitano. Robocup-97: The first robot world cup soccer games and conferences. *AI magazine*, 19 (3):49–49, 1998b.

Tenda Okimoto, Nicolas Schwind, Maxime Clement, Tony Ribeiro, Katsumi Inoue, and Pierre Marquis. How to form a task-oriented robust team. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 395–403, 2015.

Michael Boyer O'Leary and Jonathon N Cummings. The spatial, temporal, and configurational characteristics of geographic dispersion in teams. *MIS quarterly*, pages 433–452, 2007.

Mark Ollis, Wesley H Huang, and Michael Happold. A bayesian approach to imitation learning for robot navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 709–714. IEEE, 2007.

Shayegan Omidshafiei, Daniel Hennes, Marta Garnelo, Zhe Wang, Adria Recasens, Eugene Tarassov, Yi Yang, Romuald Elie, Jerome T Connor, Paul Muller, et al. Multiagent off-screen behavior prediction in football. *Scientific reports*, 12(1):1–13, 2022.

John W Orchard and John W Powell. Risk of knee and ankle sprains under various weather conditions in american football. *Medicine and science in sports and exercise*, 35 (7):1118–1123, 2003.

Oluwatoyosi BA Owoeye, Mitchell J VanderWey, and Ian Pike. Reducing injuries in soccer (football): an umbrella review of best evidence across the epidemiological framework for prevention. *Sports medicine-open*, 6(1):46, 2020.

James Parker, Ernesto Nunes, Julio Godoy, and Maria Gini. Exploiting spatial locality and heterogeneity of agents for search and rescue teamwork. *Journal of Field Robotics*, 33(7):877–900, 2016.

Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, pages 895–902, 2008.

Lokukaluge P Perera, Carlos Guedes Soares, et al. Ocean vessel trajectory estimation and prediction based on extended kalman filter. In *The Second International Conference on Adaptive and Self-Adaptive Systems and Applications*, pages 14–20, 2010.

Paul Power, Hector Ruiz, Xinyu Wei, and Patrick Lucey. 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, pages 1605–1613, 2017.

Mengshi Qi, Jie Qin, Yu Wu, and Yi Yang. Imitative non-autoregressive modeling for trajectory forecasting and imputation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12736–12745, 2020.

David Radke and Alexi Orchard. Presenting multiagent challenges in team sports analytics. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 1781–1785, 2023.

Pegah Rahimian and Laszlo Toka. Inferring the strategy of offensive and defensive play in soccer with inverse reinforcement learning. In *8th Workshop on Machine Learning and Data Mining for Sports Analytics*, 2021.

Pegah Rahimian and Laszlo Toka. Optical tracking in team sports: A survey on player and ball tracking methods in soccer and other team sports. *Journal of Quantitative Analysis in Sports*, 18:35–57, 2022.

Pegah Rahimian, Afshin Oroojlooy, and Laszlo Toka. Towards optimized actions in critical situations of soccer games with deep reinforcement learning. In *IEEE 8th international conference on data science and advanced analytics (DSAA)*, pages 1–12. IEEE, 2021.

Pegah Rahimian, Hyunsung Kim, Marc Schmid, and Laszlo Toka. Pass receiver and outcome prediction in soccer using temporal graph networks. In *International Workshop on Machine Learning and Data Mining for Sports Analytics*, pages 52–63. Springer, 2023.

Pegah Rahimian, Balazs Mark Mihalyi, and Laszlo Toka. In-game soccer outcome prediction with offline reinforcement learning. *Machine Learning*, 113(10):7393–7419, 2024a.

Pegah Rahimian, Jan Van Haaren, and Laszlo Toka. Towards maximizing expected possession outcome in soccer. *International Journal of Sports Science & Coaching*, 19(1): 230–244, 2024b.

Pegah Rahimian, Jernej Flisar, and David Sumpter. Automated explanation of machine learning models of footballing actions in words. *Journal of Sports Analytics*, 11: 22150218251353089, 2025.

N Rahnama, T Reilly, and A Lees. Injury risk associated with playing actions during competitive soccer. *British journal of sports medicine*, 36(5):354–359, 2002.

Sarvapali D Ramchurn, Maria Polukarov, Alessandro Farinelli, Cuong Truong, and Nicholas R Jennings. Coalition formation with spatial and temporal constraints. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, volume 3, pages 1181–1188, 2010.

Sarvapali D Ramchurn, Trung Dong Huynh, Feng Wu, Yukki Ikuno, Jack Flann, Luc Moreau, Joel E Fischer, Wenchao Jiang, Tom Rodden, Edwin Simpson, et al. A disaster response system based on human-agent collectives. *Journal of Artificial Intelligence Research*, 57:661–708, 2016.

Ermanno Rampinini, Franco M Impellizzeri, Carlo Castagna, Aaron J Coutts, and Ulrik Wisløff. Technical performance during soccer matches of the italian serie a league: Effect of fatigue and competitive level. *Journal of science and medicine in sport*, 12(1): 227–233, 2009.

Ezequiel Rey, Joaquín Lago-Ballesteros, and Alexis Padrón-Cabo. Timing and tactical analysis of player substitutions in the uefa champions league. *International Journal of Performance Analysis in Sport*, 15(3):840–850, 2015.

Pieter Robberechts, Maaike Van Roy, and Jesse Davis. un-xpass: Measuring soccer player's creativity. In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*, pages 4768–4777, 2023.

Nikki Rommers, Roland Rössler, Evert Verhagen, Florian Vandecasteele, Steven Verstockt, Roel Vaeyens, Matthieu Lenoir, Eva D'Hondt, and Erik Witvrouw. A machine learning approach to assess injury risk in elite youth football players. *Medicine and science in sports and exercise*, 52(8):1745–1751, 2020.

Christopher D Rosin. Multi-armed bandits with episode context. *Annals of Mathematics and Artificial Intelligence*, 61(3):203–230, 2011.

Alessio Rossi, Luca Pappalardo, Paolo Cintia, F Marcello Iaia, Javier Fernández, and Daniel Medina. Effective injury forecasting in soccer with gps training data and machine learning. *PloS one*, 13(7):e0201264, 2018.

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Guillaume Sartoretti, Justin Kerr, Yunfei Shi, Glenn Wagner, TK Satish Kumar, Sven Koenig, and Howie Choset. Primal: Pathfinding via reinforcement and imitation multi-agent learning. *IEEE Robotics and Automation Letters*, 4(3):2378–2385, 2019.

Paul Scerri, Alessandro Farinelli, Steven Okamoto, and Milind Tambe. Allocating tasks in extreme teams. In *Proceedings of the 4th International Conference on Autonomous Agents and Multiagent Systems*, pages 727–734, 2005.

Verena Schmid and Karl F Doerner. Ambulance location and relocation problems with time-dependent travel times. *European journal of operational research*, 207(3):1293–1303, 2010.

Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. Protect: An application of computational game theory for the security of the ports of the united states. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, pages 2173–2179, 2012.

Yong-Min Shin, Siqing Li, Xin Cao, and Won-Yong Shin. Faithful and accurate self-attention attribution for message passing neural networks via the computation tree viewpoint. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 20461–20469, 2025.

Shaurya Shriyam and Satyandra K Gupta. Task assignment and scheduling for mobile robot teams. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 51807, page V05AT07A075. American Society of Mechanical Engineers, 2018.

Gregor Siegert, Paweł Banyś, Cristina Sáez Martínez, and Frank Heymann. Ekf based trajectory tracking and integrity monitoring of ais data. In *2016 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 887–897. IEEE, 2016.

David Silver, James Bagnell, and Anthony Stentz. High performance outdoor navigation from overhead data using imitation learning. *Robotics: Science and Systems IV, Zurich, Switzerland*, 1, 2008.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

Jelena Simeunović, Baptiste Schubnel, Pierre-Jean Alet, Rafael E Carrillo, and Pascal Frossard. Interpretable temporal-spatial graph attention network for multi-site pv power forecasting. *Applied Energy*, 327:120127, 2022.

Karun Singh. Introducing expected threat(xt), 2019. URL `https://karun.in/blog/expected-threat.html`.

Arunesh Sinha, Fei Fang, Bo An, Christopher Kiekintveld, and Milind Tambe. Stackelberg security games: Looking beyond a decade of success. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pages 5494–5501. International Joint Conferences on Artificial Intelligence Organization, 2018.

William Spearman. Beyond expected goals. In *Proceedings of the 12th MIT Sloan Sports Analytics Conference*, pages 1–17, 2018.

William Spearman, Austin Basye, Greg Dick, Ryan Hotovy, and Paul Pop. Physics-based modeling of pass probabilities in soccer. In *Proceeding of the 11th MIT Sloan Sports Analytics Conference*, 2017.

Jochim Spitz, Johan Wagemans, Daniel Memmert, A Mark Williams, and Werner F Helsen. Video assistant referees (var): The impact of technology on decision making in association football referees. *Journal of Sports Sciences*, 39(2):147–153, 2021.

Pieter Spronck, Marc Ponsen, Ida Sprinkhuizen-Kuyper, and Eric Postma. Adaptive game ai with dynamic scripting. *Machine Learning*, 63(3):217–248, 2006.

NN Sriram, Buyu Liu, Francesco Pittaluga, and Manmohan Chandraker. Smart: Simultaneous multi-agent recurrent trajectory prediction. In *European Conference on Computer Vision*, pages 463–479. Springer, 2020.

Michael Stöckl, Thomas Seidl, Daniel Marley, and Paul Power. Making offensive play predictable-using a graph convolutional network to understand defensive performance in soccer. In *Proceedings of the 15th MIT Sloan Sports Analytics Conference*, 2021.

Chen Sun, Per Karlsson, Jiajun Wu, Joshua B Tenenbaum, and Kevin Murphy. Stochastic prediction of multi-agent interactions from partial observations. In *International Conference on Learning Representations*, 2018.

Karl Tuyls, Shayegan Omidshafiei, Paul Muller, Zhe Wang, Jerome Connor, Daniel Hennes, Ian Graham, William Spearman, Tim Waskett, Dafydd Steel, et al. Game plan: What ai can do for football, and what football can do for ai. *Journal of Artificial Intelligence Research*, 71:41–88, 2021.

Maaike Van Roy, Pieter Robberechts, Wen-Chi Yang, Luc De Raedt, and Jesse Davis. A markov framework for learning and reasoning about strategies in professional soccer. *Journal of Artificial Intelligence Research*, 77:517–562, 2023.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.

Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

Baoxiang Wang, Tongfang Sun, and Xianjun Sam Zheng. Beyond winning and losing: Modeling human motivations and behaviors with vector-valued inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 15, pages 195–201, 2019a.

Binru Wang, Yuan Zhang, Zhi-Hua Zhou, and Sheng Zhong. On repeated stackelberg security game with the cooperative human behavior model for wildlife protection. *Applied Intelligence*, 49:1002–1015, 2019b.

Jing Wang, Jianping Cai, Xiaohang Yue, and Nallan C Suresh. Pre-positioning and real-time disaster response operations: Optimization with mobile phone location data. *Transportation research part E: logistics and transportation review*, 150:102344, 2021.

Zhe Wang, Petar Veličković, Daniel Hennes, Nenad Tomašev, Laurel Prince, Michael Kaisers, Yoram Bachrach, Romuald Elie, Li Kevin Wenliang, Federico Piccinini, et al. Tacticai: an ai assistant for football tactics. *Nature Communications*, 15(1):1–13, 2024.

Feng Wu and Sarvapali D Ramchurn. Monte-carlo tree search for scalable coalition formation. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, pages 407–413, 2020.

Xu Xie, Chi Zhang, Yixin Zhu, Ying Nian Wu, and Song-Chun Zhu. Congestion-aware multi-agent trajectory prediction for collision avoidance. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 13693–13700. IEEE, 2021.

Xiaohan Xing, Fan Yang, Hang Li, Jun Zhang, Yu Zhao, Mingxuan Gao, Junzhou Huang, and Jianhua Yao. An interpretable multi-level enhanced graph attention network for disease diagnosis with gene expression data. In *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 556–561. IEEE, 2021.

Haifeng Xu, Benjamin Ford, Fei Fang, Bistra Dilkina, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, et al. Optimal patrol planning for green security games with black-box attackers. In *International Conference on Decision and Game Theory for Security*, pages 458–477. Springer, 2017.

Ming Xu, James Orwell, and Graeme Jones. Tracking football players with multiple cameras. In *International Conference on Image Processing*, volume 5, pages 2909–2912. IEEE, 2004.

Raymond A Yeh, Alexander G Schwing, Jonathan Huang, and Kevin Murphy. Diverse generation for multi-agent sports games. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4610–4619, 2019.

A Yiannakos and V Armatas. Evaluation of the goal scoring patterns in european championship in portugal 2004. *International Journal of Performance Analysis in Sport*, 6 (1):178–188, 2006.

Zhengyu Yin, Albert Xin Jiang, Milind Tambe, Christopher Kiekintveld, Kevin Leyton-Brown, Tuomas Sandholm, and John P Sullivan. Trusts: Scheduling randomized patrols for fare inspection in transit systems using game theory. *AI magazine*, 33(4): 59–59, 2012.

Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. Gnnexplainer: Generating explanations for graph neural networks. *Advances in neural information processing systems*, 32:9240–9251, 2019.

Lan You, Siyu Xiao, Qingxi Peng, Christophe Claramunt, Xuewei Han, Zhengyi Guan, and Jiahe Zhang. St-seq2seq: A spatio-temporal feature-optimized seq2seq model for short-term vessel trajectory prediction. *IEEE Access*, 8:218565–218574, 2020.

Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pages 3634–3640, 2018.

Guofu Zhang, Zhaopin Su, Xiaoxiao Song, Zixuan Gao, Miqing Li, and Xin Yao. Solving overlapping coalition structure generation in task-based settings. *Journal of Artificial Intelligence Research*, 82:2125–2166, 2025.

Jinhua Zhang, Hui Li, Peng Cheng, and Jie Yan. Interpretable wind power short-term power prediction model using deep graph attention network. *Energies*, 17(2):1–16, 2024.

Yongqi Zhang, Zhanke Zhou, Quanming Yao, Xiaowen Chu, and Bo Han. Adaprop: Learning adaptive propagation for graph neural network based knowledge graph reasoning. In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*, pages 3446–3457, 2023.

Tianxiang Zhao, Dongsheng Luo, Xiang Zhang, and Suhang Wang. Towards faithful and consistent explanations for graph neural networks. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 634–642, 2023.

Yang-Qing Zhao and Hui Zhang. Analysis of goals in the english premier league. *International Journal of Performance Analysis in Sport*, 19(5):820–831, 2019.

Weida Zhong, Qiuling Suo, Xiaowei Jia, Aidong Zhang, and Lu Su. Heterogeneous spatio-temporal graph convolution network for traffic forecasting with missing values. In *IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*, pages 707–717. IEEE, 2021.

Zhanke Zhou, Jiangchao Yao, Jiaxu Liu, Xiawei Guo, Quanming Yao, Li He, Liang Wang, Bo Zheng, and Bo Han. Combating bilateral edge noise for robust link prediction. *Advances in Neural Information Processing Systems*, 36:21368–21414, 2023.

Zhanke Zhou, Yongqi Zhang, Jiangchao Yao, Quanming Yao, and Bo Han. Less is more: One-shot subgraph reasoning on large-scale knowledge graphs. *The Twelfth International Conference on Learning Representations.*, 2024.