
Citation:

Sawczuk, T. (2023) Characterising and Analysing Performance Sequences in Rugby League Using Match Events Data [Online]. Leeds Beckett University. Available from:
<https://figshare.leedsbeckett.ac.uk/articles/thesis/Characterising_and_analysing_performance_sequences_in_rugby_league_using_match_events_data/23256440/1> [Accessed 13 June 2023]. - [Link](#)

Link to Leeds Beckett University Research Data and Thesis Repository record:

[10.25448/lbu.23256440.v1](https://figshare.leedsbeckett.ac.uk/articles/thesis/Characterising_and_analysing_performance_sequences_in_rugby_league_using_match_events_data/23256440/1)

Item Type:

Thesis

Licence:

CC BY 4.0

The aim of the Leeds Beckett University Research Data and Thesis Repository is to provide open access to the outputs and data from our research, as required by funder policies and permitted by publishers and copyright law.

The Leeds Beckett University Research Data and Thesis Repository holds a wide range of outputs and data, each of which has been checked for copyright, licenses and any relevant embargo periods have been applied by the Research Services team.

We operate on a standard take-down policy. If you are the author or publisher of an output and you would like it removed from the repository or believe there to be any issues with copyright please [contact us](#) and we will investigate on a case-by-case basis.

Characterising and Analysing Performance Sequences in Rugby League Using Match Events Data

Dr Thomas Sawczuk

A thesis submitted in partial fulfilment of the requirements of Leeds
Beckett University for the degree of

Doctor of Philosophy

School of Built Environment, Engineering and Computing

November 2022

Dr Thomas Sawczuk

Characterising and analysing performance sequences in rugby league using match events data

Keywords: rugby league, performance analysis, markov, kernel density estimation, wasserstein distance, bayesian

Abstract

The ability to accurately evaluate player and team performances in professional sport is particularly valuable. Doing so provides competitive advantages include extracting important information regarding the tactical strategies of future oppositions and producing player rating systems. A common method of evaluating player and team performances is via expected possession value (EPV) models. EPV models assign a value to every location and/or action on the pitch, which reflects the probability of points being scored within a given time period.

EPV models have been produced in several sports, including football, basketball and ice hockey. However, there is limited research surrounding these models in rugby league. Rugby league has a unique set of rules, including a six tackle attacking set and five possible scoring options at the end of a possession. These two factors, alongside the poor data availability in the sport ensure that the majority of previous methods cannot be adapted for use in rugby league. Therefore the aim of this thesis was to develop new methodologies evaluating player and team performances in rugby league.

In the first section of this thesis (studies 1 and 2), previous Markov models using zonal approaches were applied, adapted and extended in rugby league to provide insights into player and team performances. Six EPV models were produced with varying zone sizes using Markov Reward Processes. The Kullback-Leibler Divergence was used to evaluate the zone sizes which could

reproduce future team attacking performances. The model was then extended to incorporate actions and context nodes using Markov Decision Processes. Novel methods of evaluating player and team performances were also produced.

In the second section (studies 3 and 4), novel models producing smooth pitch surfaces were developed. The spatial trends of team attacking performances were evaluated using Kernel Density Estimation. Two novel Wasserstein distance metrics were used to provide valuable insights into team performances. A novel approach to the estimation of individual possession outcomes was also proposed using a Bayesian mixture model approach. The model used linear and bilinear interpolation techniques for its weights to produce a smooth pitch surface. Novel performance metrics evaluating player and team performances were also created.

The research provides new methodologies for use within rugby league, providing zonal and smooth EPV models through which player and team performances can be evaluated. Professional experts were impressed with the results they provided and validated their use within the sport.

Declaration

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

Thomas Sawczuk

Acknowledgements

Doing a second PhD in a different subject, immediately after completing your first, is hard. Starting it just before COVID-19 and the impact that had on everybody made it even harder. There were a number of occasions where I nearly didn't make it and the truth is that I wouldn't have made it without the support that I had.

With that in mind, as I complete my PhD journey, I have to acknowledge my supervisors, Ben, Anna and Jan. Ben, thank you for seven years of trust and guidance across two PhDs. My life would look very different right now if you hadn't taken a chance on me in 2015 and I'll never forget that. Anna and Jan, thank you for everything you've done for me throughout this PhD. Anna, for guiding me after the difficult start to the project, for making the right call to move away from computer vision, for helping shape my thought processes and for recognising when I just needed a break and letting me take it. Jan, for being patient every time it took me 45 minutes to explain that I meant the same thing as you in countless meetings, for helping me to develop my mathematical precision and for providing a brilliant Bayesian mind despite not being a statistician!

Unlike my first PhD, this was a more isolated experience so there's a much smaller pool of colleagues to thank. Partly due to COVID, partly due to me knowing what I needed to do to get this done. Thanks to the CARR group for the times we spent together over the last few years and a big thank you to Victor for the work we did together mainly at the start, trying to figure out the Narya computer vision workflow. I know you had your own stuff to work with then so thank you for the help you gave me at that time in particular.

Finally, thanks to Aimee, Wally, Muffin and the Sawczuk family. Aimee for the support you've given me and for telling me to take a break now. I won't

be doing a third PhD after this one. I promise. Wally and Muffin for being the best first dogs a man could ever ask for and making me take breaks from the PhD to walk you! The rest of the family, thanks for the Skype game nights, the surprise birthday party, the general support and Xbox nights. I wouldn't have completed this process without you.

Publications and Presentations

Books and Journals

- Thomas Sawczuk, Anna Palczewska, Ben Jones. Development of an expected possession value model to analyse team attacking performances in rugby league. *PLOS One*, 16(11): e0259536, 2021
- Thomas Sawczuk, Anna Palczewska, Ben Jones, Jan Palczewski. A Bayesian expected possession value model to evaluate player and team performances. In progress, plan to submit to sport journal.

Conference proceedings

- Thomas Sawczuk, Anna Palczewska, Ben Jones. "Markov Decision Processes with contextual nodes as a method of assessing attacking player performance in rugby league" in *20th UK Workshop on Computational Intelligence*, Aberystwyth, 2021.
- Thomas Sawczuk, Anna Palczewska, Ben Jones, Jan Palczewski. "Use of Kernel Density Estimation to understand the spatial trends of attacking possessions in rugby league" in *21st UK Workshop on Computational Intelligence*, Sheffield, 2022.

Presentation and posters

- "Markov Decision Processes with contextual nodes as a method of assessing attacking player performance in rugby league", UKCI 2021 Aberystwyth Presentation, September 2021
- "Use of Kernel Density Estimation to understand the spatial trends of attacking possessions in rugby league", UKCI 2022 Sheffield Presentation, September 2022

- ”Development of an expected possession value model to analyse rugby league team’s attacking performances”, Internal England Performance Unit Presentation, October 2020
- ”Evaluating team and player performances in rugby league using Markov Decision Processes”, Internal Carnegie Applied Rugby Research Centre Presentation, August 2021
- ”Gaining greater insights into rugby league using Stats Perform event level data”, External StatsPerform Presentation, April 2022
- ”Expected Possession Values in rugby league”, External Catapult Presentation, May 2022
- ”Gaining insights into rugby league match performances using advanced statistical methods”, Internal Carnegie School of Sport Postgraduate Conference Presentation, June 2022
- ”Understanding player and team performances in rugby league” External Leeds Rhinos Presentation, October 2022

Contents

1 Introduction1

1.1	Background.....	1
1.2	Motivation.....	3
1.3	Goals and Objectives.....	5
1.4	Thesis Scope	6
1.5	Thesis Structure.....	7

2 Literature Review11

2.1	Data in Team Sports	11
2.2	Team and Player Performance Analysis in Sport	13
2.2.1	Accumulated Count Data Analysis in Team Sports.....	14
2.2.2	Action by Action Data Analysis in Team Sports.....	19
2.2.2.1	Spatial Trends Analysis in Team Sports.....	19
2.2.2.2	Expected Possession Value Models in Team Sports.....	24
2.3	Rugby League	35
2.3.1	Overview	35
2.3.2	Team and Player Performance Analysis in Rugby League.....	41
2.3.2.1	Accumulated Count Data Analysis in Rugby League.....	41
2.3.2.2	Play by Play Data Analysis in Rugby League.....	43
2.4	Summary.....	45
2.4.1	Spatial Trends Models.....	46
2.4.2	Expected Possession Value Models.....	50
2.5	Methodological Overview.....	52

3 Markov Reward Processes for the Quantification of Team Attacking Performances55

3.1	Introduction.....	55
3.2	Methodology	56
3.2.1	Data	56
3.2.2	Data Preprocessing	58
3.2.3	Markov Reward Processes.....	62
3.2.3.1	Expected Possession Value.....	64
3.2.4	Configuration of Zone Sizes.....	66
3.2.4.1	Kullback-Leibler Divergence.....	68
3.2.5	Z-score Analysis.....	70
3.3	Results.....	71
3.3.1	EPV models.....	71
3.3.2	Reproducibility of Future Attacking Performances.....	79
3.3.3	Quantification of Teams' Attacking Performances.....	82
3.4	Summary.....	85

4 Markov Decision Processes for the Evaluation of Player and Team Performances87

4.1	Introduction.....	87
4.2	Methodology	88
4.2.1	Data	88
4.2.2	Data Preprocessing	89
4.2.3	Markov Decision Processes	94
4.2.4	Action Impact Ratings	98
4.2.5	Team State, Action Values	98
4.3	Results	99
4.3.1	Evaluation of Player Performances	99
4.3.1.1	Action Impact Ratings	100
4.3.1.2	Contextual Insights	102
4.3.2	Evaluation of Team Performances	102

4.4	Summary	105
5	Kernel Density Estimation and Wasserstein Distance Evaluation of the Spatial Trends of Team Attacking Performances	109
5.1	Introduction	109
5.2	Methodology	110
5.2.1	Data	110
5.2.2	Data Preprocessing	111
5.2.3	Kernel Density Estimation	117
5.2.4	Wasserstein Distance	120
5.2.4.1	Normalised Axis Wasserstein Distance	122
5.2.4.2	Directional Wasserstein Distance	124
5.3	Results	125
5.3.1	Quantification of the Spatial Trends of Team Attacking Performances	126
5.3.2	Evaluation of Between and Within Team Differences	127
5.3.2.1	Identification of Suitable p-norm	128
5.3.2.2	Normalised Axis Wasserstein Distance Differences at the Between Team Level	130
5.3.2.3	Directional Wasserstein Distance Differences at the Between Team Level	135
5.3.2.4	Evaluation of Within Team Differences	137
5.4	Summary	143
6	A Bayesian Approach to the Evaluation of Team and Player Performances	145
6.1	Introduction	145
6.2	Methodology	146
6.2.1	Data	146
6.2.2	Data Preprocessing	147
6.2.3	Bayesian Mixture Model	149
6.2.3.1	Centre Weights	152
6.2.3.2	EPV Calculation	154
6.2.3.3	Modelling Approach	155

6.2.4	Expected Points Scored	158
6.2.5	Actual vs Expected Ratings	159
6.3	Results	160
6.3.1	Estimation of Individual Possession Outcomes	161
6.3.2	Evaluation of Team Performances	164
6.3.2.1	Expected Points Scored	171
6.3.3	Evaluation of Player Performances	172
6.4	Summary	174
7	Conclusions	176
7.1	Rugby League	177
7.2	Research Contributions	178
7.2.1	Markov Reward Processes for the Quantification of Team Attack- ing Performances	178
7.2.2	Markov Decision Processes for the Evaluation of Player and Team Performances	179
7.2.3	Kernel Density Estimation and Wasserstein Distance Evaluation of the Spatial Trends of Team Attacking Performances	180
7.2.4	A Bayesian Approach to the Evaluation of Team and Player Per- formances	182
7.2.5	Summary of contributions	183
7.3	Limitations	184
7.4	Future Work	186
	References	189
	Appendix A Ethical Approval	198
	Appendix B List of actions and definitions coded by Opta/Stats Perform 2019- 2020	200
	Appendix C Opta/Stats Perform User Manual for 2021 Data	215
	Appendix D Leeds Rhinos Impact Statement	259

List of Figures

- 1.1 Overview of thesis outlining most important literature, timeline and how the work links to the literature. Note: Objective 1 was completed by Chapter2, but was not included in the schematic as it is not a specific study.8

- 2.1 Woods(2016) violin plot of performance indicators significantly associated with final league position. Model was completed using season accumulated counts, dots in this plot represent match accumulated counts and show the variability inherent within the use of accumulated match counts on a match by match basis. League position is denoted in alphabetical order (e.g. A = 1st, B = 2nd etc). 17

- 2.2 Whitaker et al.(2021) posterior plot of goal abilities (i.e. those actions attempting to score goals) for Sturridge and Reed. 18

- 2.3 Spearman et al.(2017) pitch control probability plot. On the left, blue regions are those controlled by the team in possession of the ball, red regions are controlled by the opposition team. White areas are contested (i.e. not controlled by either team). On the right, the shaded region represents the area of the pitch controlled by blue number 7. The trail behind every player, the ball represents their position over the past 3 seconds. 21

- 2.4 **Fernández and Bornn(2018)** pitch control plot. In the left frame, a player moves backwards from the 18 yard line to occupy a space of value with higher control. In the middle frame, the player attacks a space in front of him, dragging three defenders towards him. In the third frame, the green circled player is free because of the movement in the middle frame. The player on the 18 yard line makes a forward run into space so he can receive a crossed pass..... 22
- 2.5 **Mallepalle et al.(2020)** quarterback KDE plot (left) and pass completion probability surface (right)..... 23
- 2.6 **Cervone et al.(2016)** EPV plot for a sample possession. Immediate, rapid changes are caused by actions such as passes and shots. Smaller, more incremental changes in EPV are induced by changes in players' locations. The black line slightly smooths the EPV evaluations at each time point (grey dots)..... 26
- 2.7 Plot outlining the decomposed MDP methodology employed by **Fernández et al.(2019,2021)**. Figure taken from **Fernández et al.(2019)**. EPV plot highlights three specific situations (A, B and C). Blue areas represent areas of low EPV, red areas represent high EPV. Blue arrows represent pass chosen by the player, green arrows represent best possible pass according to the model. Yellow dots represent players on the attacking team (i.e. in possession), blue dots represent defending team..... 27
- 2.8 **Fernández et al.(2021)** evaluation of single player's EPV added when passing (left) and receiving (right) passes from different players. Plot shows probability of passes being selected by size of player circle and value of pass through the arrow size and colour. Metrics are normalised by the minutes the two players played together and multiplied by 90 minutes to provide a match value. 28
- 2.9 **Fernández et al.(2021)** plot identifying the EPV generated by teams in different formations against a specific team. The green circle represents the ball location. 29
- 2.10 **Decroos et al.(2019)** plot of VAEP for different actions. Plot shows an overview of the total contribution per 90 minutes for different types of actions for different players in the 2016/17 (a) and 2017/18 (b) seasons.... 29

2.11 Bransen and Haaren(2020) joint chemistry plot showing value generated between pairs of players on the same team. Green lines represent high attacking chemistry links, dark red lines reflect low attacking chemistry links.....	30
2.12 Routley and Schulte(2015) box plot of action impact values across different contexts. Left shows the action impact values for goals, where a high value is better for the team. Right shows the action impact values for penalties, where a high value is worse for the team. Central mark of the box plot is the median, the edges represent the 25th and 75th percentiles and the whiskers are at approximately 2.7 standard deviations.	31
2.13 Liu et al.(2020) action values for four different actions (shot, pass, cross and tackle) across the pitch.	33
2.14 Singh(n.d.) xT plot.....	34
2.15 Plot of 70m x 120m rugby league pitch. Assumes the team in possession of the ball is attacking in an upwards direction. Own refers to the possession team; Opp refers to the opposition team.....	36
2.16 Kempton et al.(2016) EPV plot.	44
3.1 Data preprocessing workflow	59
3.2 Counts for the actions beginning a play according to the Opta coding system in 2019. All actions began a play on at least 100 occasions. Complete passes are clearly the most frequent action.	60
3.3 Data analysis workflow.....	62
3.4 Depiction of relationship between EPV-308 zones and columns/rows used for statistical analysis. Top figure represents the 5m x 5m EPV-308 zones. Bottom left figure represents the columns; column match EPVs were produced by summing the match EPV for all zones from the EPV-308 in the column. Bottom right figure represents the rows; row match EPVs were produced by summing the match EPV for all zones from the EPV-308 in the row. Pitch lines represent own try area, half way line and opposition try area for team attacking in upwards direction.	67

3.5	Depiction of similarities and differences between EPV-308 (A), EPV-77 (B) and EPV-19 (C) in the 30m closest to the opposition try line. Each zone from the EPV-77 and EPV-19 is a weighted average of the EPV-308 zones they are composed of. Dotted line represents the opposition 20m line.	69
3.6	EPV-308 (A), EPV-77 (B), EPV-37 (C), EPV-19 (D), EPV-13 (E), EPV-9 (F). These are the six EPV-models produced in this study, where the number represents the number of zones present within the model. Lines represent the 20m line and 50m line. All zone values are shaded to the same scale.	72
3.7	KL Divergence values for EPV-77, EPV-37, EPV-19, EPV-13 and EPV-9. A lower value indicates greater similarity in reward distributions between previous matches and the subsequent match. EPV-308 not included as values could not be calculated for the first seven matches due to no non-infinity values being present. Line provides mean value, shaded area indicates standard deviation.	81
3.8	EPV-19 zones numbered so they can be distinguished from each other. Where numbers are repeated, both sides of the pitch make up the same zone (e.g. zone 2 is comprised of the widest ~ 5m on both sides of pitch, between 10m and 70m from the team's own try line.	83
3.9	Z-score analysis of teams' attacking performances for the whole 2019 Super League season. Numbers 1-19 reflect the zone numbers in Figure 3.8. A greater value indicates a greater proportion of EPV was obtained from the zone than the average Super League 2019 team. Distances are measured from the teams own try line.	84
4.1	Data preprocessing workflow	90
4.2	Standard rugby league pitch (A). Lines represent try line, 20m from the try line and half way line. Rugby league pitch split into the 19 states used within this analysis (B).	92
4.3	Boxplot showing variation in action values for each action across the 19 states with no context nodes included (A), period as context (B) and match score as context (C)). Circles depict outliers, calculated as those values more than 1.5*interquartile range beyond the upper quartile.	100

LIST OF FIGURES

4.4	Team state, action value plot for Hull. All state, action values are centred by the league average value. Green values indicate team is better at the action in a specific state than the league average team; red indicate the team is worse.....	105
4.5	Team state, action value plot for St Helens. All state, action values are centred by the league average value. Green values indicate team is better at the action in a specific state than the league average team; red indicate the team is worse.....	106
5.1	Data preprocessing workflow	112
5.2	Location of 99,966 actions used in this study. Dots representative of location only; quantity of actions at each location is not displayed in this plot.	115
5.3	Workflow of analyses conducted within this study	117
5.4	Standardised, unweighted Wasserstein metrics for $p = 1$ (left), $p = 2$ (centre) and $p = 3$ (right). Distances are calculated from central point.	122
5.5	Kernel Density Estimation plot for whole league data at bandwidth 4.10 .	127
5.6	Comparison between whole league attacking KDE and each team's overall KDE. Green areas are areas represent areas where the team has a higher density than the whole league, red areas have lower densities.	128
5.7	Visualisation of the diagonal distance when moving mass from the centre of the pitch relative to the horizontal/vertical components of the movement. Left figure shows that the diagonal distance covered to move mass from the opposition 20m to the team's own try line is significantly greater than the horizontal distance. Right figure shows that the distance covered is much more even between diagonal and vertical movements, reducing the influence of diagonal movements on the overall Wasserstein distance when mass is moved vertically relative to horizontally	131
5.8	Advanced Wasserstein metrics for $p = 1$ (left), $p = 2$ (centre) and $p = 3$ (right). Plots are shown for x-factor = 6, y-factor = 1 and highlight the rectangular shape of $p = 3$ compared to the diamond for $p = 1$ and oval for $p = 2$. Distances are calculated from central point.....	132

5.9	Quadrant plot for Wasserstein distance using factor 2 for x and y directions. y-axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x-axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.	132
5.10	Quadrant plot for Wasserstein distance using factor 2 for x and y directions. y-axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x-axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.	133
5.11	Quadrant plot for Wasserstein distance using 4.8x, 2.8y factors. y-axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x-axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.....	133
5.12	Quadrant plot for Wasserstein distance using 6.0x, 3.5y factors. y-axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x-axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.....	134
5.13	Quadrant plot showing directional Wasserstein distance at the between team levels. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. Blue lines represent median values.....	136
5.14	Comparison between team-opponent KDE and team's overall KDE for Castleford. Top left plot shows overall KDE for Castleford. For all other plots, green areas are areas represent areas where the team has a higher density against a specific opponent than the overall KDE, red areas have lower densities.	138

5.15	Quadrant plot showing directional Wasserstein distance for Castleford against different opponents. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. Blue lines represent median values.....	139
5.16	Comparison between team-opponent KDE and team's overall KDE for Catalans. Top left plot shows overall KDE for Catalans. For all other plots, green areas are areas represent areas where the team has a higher density against a specific opponent than the overall KDE, red areas have lower densities.	140
5.17	Quadrant plot showing directional Wasserstein distance for Catalans against different opponents. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. Blue lines represent median values.	141
5.18	Zoomed in quadrant plot showing directional Wasserstein distance for Castleford against different opponents. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. The opponent Leeds is removed. The co-ordinates for Leeds would be (9.41, -0.58) and (9.41, 7.14) for the left and right plot respectively. Blue lines represent median values, including Leeds.	142
6.1	Location of 30 centres used in this study. Three centres (not plotted) were included in the try area at $x \in \{0, 35, 70\}$, equivalent to the left, middle and right centres in the field of play. No y coordinate was considered for these centres.	153
6.2	Whole league model mean pitch surface plot. 4pt Try and 6pt Try refer to unconverted and converted tries respectively. With the exception of EPV, probabilities are plotted. Smooth pitch surface for possession outcome probabilities is calculated using Equation 6.3 for each x, y location on the pitch. EPV for each location is calculated using Equation 6.10. Brighter areas represent higher values.	162

6.3	Whole league model EPV standard deviation plot. 4pt Try and 6pt Try refer to unconverted and converted tries respectively. Whole league KDE is reproduced from Chapter 5 to provide an understanding of location densities. Brighter colours indicate higher values.....	164
6.4	Castleford smooth pitch surface plot from team attacking model. Green areas represent higher value for more favourable event (i.e. greater probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively	165
6.5	Catalans smooth pitch surface plot from team attacking model. Green areas represent higher value for more favourable event (i.e. greater probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively	166
6.6	Wigan smooth pitch surface plot from team attacking model. Green areas represent higher value for more favourable event (i.e. greater probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively	167
6.7	Castleford smooth pitch surface plot from team defending model. Green areas represent higher value for more favourable event (i.e. lower probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively	168
6.8	Catalans smooth pitch surface plot from team defending model. Green areas represent higher value for more favourable event (i.e. lower probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively	169
6.9	Wigan smooth pitch surface plot from team defending model. Green areas represent higher value for more favourable event (i.e. lower probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively	170

LIST OF FIGURES

A.1	Screenshot of ethical approval email for studies in this thesis	199
-----	---	-----

List of Tables

2.1	Possession outcomes in rugby league	37
2.2	Possession terms used in this thesis, ordered from lowest level (i.e. a single action) to highest (i.e. all actions within a specified period of time).	38
2.3	Rugby league player positions and their general responsibilities. Brackets represent number of players in that position on the pitch.....	40
2.4	Summary of studies considering variations of action by action or play by play data analyses across all sports.....	47
3.1	Details of the four nodes provided by Opta in each rugby league fixture xml file.....	56
3.2	Important match data variable definitions from the 2019 Opta dataset.	57
3.3	List of events, which could end an attacking possession.....	61
3.4	First two attacking possessions of the preprocessed dataset. FXID is the fixture ID; PosNum refers to the time point of the play within the possession sequence.....	62
3.5	Linear mixed model derived average column match EPV for EPV-308 columns. Bracketed values are metres included for column and 95% confidence intervals for column match EPV	73
3.6	Linear mixed model derived average row match EPV for EPV-308 rows. Bracketed values are metres included for row and 95% confidence intervals for row match EPV.....	73

3.7	EPV-308 symmetrically aggregated column match EPVs. Bracketed values are metres included for column and 95% confidence intervals for column EPV from linear mixed model.....	74
3.8	EPV-308 sequentially aggregated row match EPVs. Bracketed values are metres included for row and 95% confidence intervals for row EPV from linear mixed model.	74
3.9	Statistical comparison between columns at an SESOI of 0.5 for EPV-37. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 0.5. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally	75
3.10	Statistical comparison between rows at an SESOI of 0.5 for EPV-37. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 0.5. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area.	75
3.11	Statistical comparison between columns at an SESOI of 1.0 for EPV-19. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 1.0. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally	76

- 3.12 Statistical comparison between rows at an SESOI of 1.0 for EPV-19. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 1.0. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area. 76
- 3.13 Statistical comparison between columns at an SESOI of 1.5 for EPV-13. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 1.5. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally 77
- 3.14 Statistical comparison between rows at an SESOI of 1.5 for EPV-13. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 1.5. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area. 77
- 3.15 Statistical comparison between columns at an SESOI of 2.0 for EPV-9. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 2.0. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally 78

3.16	Statistical comparison between rows at an SESOI of 2.0 for EPV-9. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 2.0. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area.	78
3.17	Percentage of non-infinity zones for each model, providing the percentage of matches where the complete set of the subsequent match's zones were visited in the previous n matches. Values are mean (standard deviation) percentage across all clubs.	80
4.1	Definitions of grouped actions used in this study.....	91
4.2	Example episode sequence. FxID refers to fixture identification number, PLID refers to player identification number.....	93
4.3	Top 20 normalised player action impact ratings for the 2020 Super League season, including goal kicks. Tries, Try Assists, Metres and Goals are provided as references of statistics currently provided for player performances. To protect anonymity, reference statistics are provided based on the whole season as: T-5 (top 5); T-10 (top 10), T-20 (top 20) and 20+ (outside top 20). Players completed a minimum of 150 actions across the season. Player ID brackets provide position in the table when goal kicks are not included in the action impact rating.	101
4.5	Top 20 normalised player action impact ratings for the 2020 Super League season, not including goal kicks. Tries, Try Assists, Metres and Goals are provided as references of statistics currently provided for player performances. To protect anonymity, reference statistics are provided based on the whole season as: T-5 (top 5); T-10 (top 10), T-20 (top 20) and 20+ (outside top 20). Players completed a minimum of 150 actions across the season. Player names are anonymised, brackets provide position in the table when goal kicks are included in the action impact rating	103

4.7	Contextual action impact ratings for players nominated for Man of Steel award. Player names are anonymised.....	104
5.1	Passing example to show the action definitions utilised by Opta in 2021 ..	111
5.2	Preprocessing categories and actions included within them in this study ..	113
5.3	Sample possession within the dataset.....	116
5.4	Descriptive data for action counts in the data subsets. IQR is interquartile range; Min/Max refers to minimum/maximum number of observations observed.	116
5.5	Weighted Wasserstein distance metrics considered in this study	124
5.6	Wasserstein distances for W_1 , W_2 and W_3 at the between team level. Distances represent differences between the team's overall KDE and the whole league KDE. Final column ($W_1 - W_3$) is the difference between the W_1 and W_3 distances	129
6.1	Descriptive data for data subsets. IQR is interquartile range; Min/Max refers to minimum and maximum number of observations.....	149
6.2	Sample possession within the dataset. Data includes the teams involved in the possession, the player ID, the x, y coordinates of the action, the possession number (PosNum) and the possession outcome (PosCat; in this case no try for all rows).	149
6.3	Prior α values for whole league Bayesian model. Centre coordinates are provided, alongside α values for each possession outcome. C.Try refers to converted try (6 points) and U.Try refers to unconverted try (4 points). ..	156
6.4	α prior distribution values for team attacking and defending models. α values calculated via maximum likelihood estimation from posterior distribution of the whole league model. Centre coordinates are provided, alongside α values for each possession outcome. C.Try refers to converted try (6 points) and U.Try refers to unconverted try (4 points). Values rounded to 2 decimal places for brevity	157
6.5	Abbreviated example of weights used in this study. Only first 8 centres are shown, column titles represent centre coordinates. All other centres ($n = 25$) are assigned weights of 0 based on the x,y location of these example actions.	161

6.6	Expected points scored across the season compared to actual points scored at the team overall level. Exp Att refers to expected points scored as computed using the attacking team EPV; Exp Def refers to expected points scored as computed using the defending team EPV; Diff refers to the difference between points scored and the expected points scored (positive means more points were scored than expected). Teams are provided in league position order	172
6.7	Top 20 player ratings as assessed by the AE ratings (Equation 6.12). Tries, Try Assists, Metres and Goals are provided as references of statistics currently provided for player performances. To protect anonymity, reference statistics are provided based on the whole season as: T-5 (top 5); T-10 (top 10), T-20 (top 20) and 20+ (outside top 20).	173
1	Opta actions and definitions for the 2019 Super League season	201
2	Additional Opta actions and definitions for the 2020 Super League season. All actions described in 2019 table were coded in 2020, as well as these additional actions.	213

Introduction

1.1 Background

In the past two decades there has been an exponential increase in the amount of player data generated by sports teams. This increased data availability has been driven by the widespread use of wearable technology (e.g. global positioning system (GPS) units and heart rate monitors) and the introduction of event level match play data annotation and computer vision technologies. GPS units allow practitioners to understand the distances players run in a match and the accelerations and decelerations that they experience. Heart rate monitors allow the heart rate of players to be monitored alongside locomotor information to understand how hard a player is working physically. Event level data provides information surrounding the events or actions that occur within match play (e.g. passes, runs and kicks). In most sports, including rugby league, it is collected via human annotation. Human annotation involves an individual watching matches and coding each event that occurs in the match. It can happen either live or retrospectively. Recently, event level information has been extended and improved in some sports, most notably football, by the use of computer vision technologies, which allow more data to be provided for each event (e.g. the location of other players in the same video frame, rather than just the player performing the action ([StatsBomb,2022](#))).

With the influx of wearable technology data across all team sports, the number of published sports science research studies has exploded in recent years. Research providing physiological insights is therefore available across multiple sports ([Whitehead et al.,](#)

2018) or specific to individual sports such as football (Whitehead et al.,2018), hockey (Gabbett,2010) or rugby (Dalton-Barron et al.,2020). Unlike wearable technology data, which is widely available across all sports, the use of event level data to provide tactical insights is less common. Within the scope of this PhD, tactical insights refer to understanding team and player performances through the use of match event data (i.e. the actions players or teams perform on the pitch). Advanced tactical analyses are becoming increasingly common in those sports which are early adopters of computer vision technologies as a method of event level data annotation (e.g. football and basketball (Cervone et al.,2016;Decroos et al.,2019;Fernández et al.,2021;Liu et al.,2020)). Through the use of computer vision technologies, these sports can produce the large quantities of data necessary for advanced methods such as deep learning (Fernández et al.,2021) or deep reinforcement learning (Liu et al.,2020) to be conducted. They have been able to produce significant insights, which in some sports have resulted in drastic tactical changes (e.g. basketball, where there has been a large increase in 3-point shooting since research showed that it was a more efficient method of scoring than attempting 2-point shots (Cohen,2016)). In other sports, with low data availability, the analysis of match event level data is still uncommon.

Rugby league is an example of a sport with low match event level data availability. It is a field based collision sport, where two teams of 13 players attempt to score points over two 40 minute halves. Within each half, the teams attempt to progress as far as possible towards the opposition try line within their allotted set of six tackles. The sequence of actions leading to a tackle is called a 'play', and the sequence of six plays is termed a 'set'. At the end of each set, a team either scores points, receives an additional set of six plays due to a foul/error by the opposition, or loses possession due to a failed scoring attempt, a kick which is caught by the opposition, or a handover, where the opposition is given the ball to start their set of plays at the position of the sixth tackle. As such, each set may have as many as six tackles, but if one of the set-ending actions described above occurs, the set may be finish before the sixth play is completed. If the opposition commits fouls or makes errors, an attacking team may complete multiple sets within a possession before the opposition begins their next set of plays. There are four methods of scoring points: an unconverted try, where the ball is grounded beyond the opposition try line (4 points); a converted try, where the ball is grounded beyond the opposition try line and a subsequent conversion kick is scored (6 points); a penalty goal, where the ball is kicked

between the opposition posts from the floor (2 points); and a drop goal, where the ball is kicked between the posts from a player's hands (1 point). At the end of the match, the team which scores the most points is declared the winner. Since 2019, if the teams' points are level at the end of the match, an additional 10 minutes of extra time (5 minutes in each half) have been played on a 'golden point' basis, whereby the first team to score a point of any kind is declared the winner.

1.2 Motivation

Rugby league has a rich history of using wearable technology to monitor and improve athlete performance (Dalton-Barron et al.,2020;Glassbrook et al.,2019) and has recently embraced the introduction of instrumented mouth guards as a method of understanding concussion incidence (Tierney et al.,2021;Tooby et al.,2022). However, much less research has taken place regarding the tactical demands of the game (Kempton et al.,2016; Parmar et al.,2018;Woods et al.,2017). There are several reasons for this. First, there is reduced data availability compared to other sports, such as basketball or football. Action by action data annotation has only recently been introduced in rugby league, with computer vision technologies which can handle the regular occlusion of the ball/players not yet developed for the sport. Further to this, the GPS units used by the teams regularly have signal issues in enclosed places such as the stadiums where matches are played, which can result in lost data. They also aren't accurate enough (Song et al.,2022) to be used for the same type of tracking and match event data analysis previously employed in football (Fernández et al.,2021) or basketball (Cervone et al.,2016). Previous analysis using event level only data has been conducted in football (Decroos et al.,2019) but the authors used separate models to evaluate the probability of a goal being scored or conceded. Rugby league has five possible possession outcomes, so adapting this model to the sport would result in an excessive requirement of eight separate models (one for each outcome scored and one for each outcome conceded). A further modelling consideration is the episodic nature of rugby league, which is ideally suited to the Markovian approach used byKempton et al.(2016). Their approach combines all possible possession outcomes into a single value, but this limits the flexibility of the model as information regarding each individual possession outcome is lost. It is further limited by the requirement to identify areas on the pitch through which data is aggregated. In rugby league, a sport

with a standard 100m x 70m pitch, where points can be scored across the pitch width, it is extremely difficult to accurately identify these areas. This requirement is much simpler in basketball, where it has previously been employed and there are specific shooting areas which can be used (Cervone et al.,2016). Regardless of these concerns, a smooth pitch surface would be more desirable than a zone based approach. Unfortunately, the limited data availability in rugby league (typically 180 matches per season and one league of data available vs over 300 matches available per season per league across multiple European Leagues in football) has ensured that the machine learning (Decroos et al.,2019) and deep learning (Fernández et al.,2021;Liu et al.,2020) methodologies currently used to obtain smooth pitch surfaces in other sports cannot yet be adapted to rugby league. As a result of these differences, there is a need to develop a framework through which the tactical elements of rugby league can be analysed within the constraints of low data availability (i.e. using only the event level data available) and accounting for the multiple point scoring possession outcomes present within the sport.

Across all sports, the two key tactical elements through which insights can be provided are *team performances* and *player performances*. Evaluating opposition team performances can allow strategies for future matches to be produced, maximising the chances of beating the opposition. Understanding player performances objectively can help to identify players who provide better or worse value for the salary they are paid. Both elements are important in all sports, but in rugby league they are perhaps more important due to the salary cap employed which means that teams can only spend a specified amount of money on all players within their squad. This salary cap has the dual effect of reducing the differences in player quality between teams compared to other sports (e.g. football, where the richer clubs can offer better players significantly higher salaries) and increasing the importance of tactical strategies with regard to the match outcome as two evenly matched teams are more likely to be separated by their ability to work as a unit than by individual moments of brilliance.

With high quality player tracking and event level data available in football, basketball and ice hockey, the use of modern deep learning (Decroos et al.,2019;Fernández et al., 2021), deep reinforcement learning (Liu and Schulte,2018;Routley and Schulte,2015) and probabilistic (Cervone et al.,2016) methods to evaluate player and team performances has become more prominent. These models are able to provide excellent analyses of team and player performances, allowing coaches to identify strategies to beat opposition teams

and evaluate undervalued players, who may be good transfer targets. Unfortunately, the research in these sports is much more limited with respect to its applicability to rugby league due to the sport's lower quantity and quality of available data and its more complicated point scoring system. Indeed, action by action data only became available in rugby league in 2020, so the only studies considering any kind of spatial data at the onset of this thesis have been conducted using play by play data (Holbrook et al., 2019; Kempton et al., 2016). Kempton et al. (2016) used dynamic programming to estimate the value of different locations on the pitch dependent on the next scoring action. Their work combined all possession outcomes into a single value and utilised a zonal approach, the limitations of which are discussed above. Holbrook et al. (2019) used a deep learning approach to concurrently predict six outcomes, including: expected metres gained by the play; expected try scored within the play/set; and predicted match outcome. They were able to produce a smooth pitch surface, but required more than 250,000 play by play observations across three seasons of National Rugby League to do so using their deep learning approach. This size of dataset is not available in the Super League, which is the focus of this thesis. Furthermore, the application of both studies to team or player performance analysis is limited by their use of play-by-play data, rather than action by action data, which wasn't available at the time they were published. There is therefore a gap in the rugby league literature whereby a methodology, which can evaluate player and team performances on an action by action level, providing a smooth pitch surface and estimating the individual probabilities of possession outcomes is required.

This research will address the gap in the literature by providing new methodologies to evaluate player and team performances using action by action data in rugby league - a sport with low data availability. The project is part of a collaboration between the Rugby Football League, Carnegie School of Sport and School of Built Environment, Engineering and Computing at Leeds Beckett University. The overarching aim was to provide new insights into rugby league through the use of advanced data analytics methods.

1.3 Goals and Objectives

This PhD project aims to develop new methodologies to evaluate player and team performances in rugby league using match event level data. To achieve this goal, the following objectives will be met.

- O1 Investigate existing methodologies of evaluating player and team performance in team sports, focusing on their application to rugby league
- O2 Apply and adapt existing methodologies and metrics evaluating player and team performances in team sports to rugby league
- O3 Validate adapted versions of existing methodologies and metrics for usability and reliability in rugby league
- O4 Develop novel methodologies and metrics to evaluate player and team performances in rugby league
- O5 Validate proposed methodologies and metrics for usability and reliability with respect to evaluating player and team performances in rugby league

The data for the project was provided by the Rugby Football League and is proprietary data owned by Opta/Stats Perform. Throughout the duration of the PhD, StatsPerform purchased Opta and progressively improved their data collection processes allowing for more advanced studies to take place as shown within the goals of the project. Unfortunately, as data collection processes were improving, action definitions were also changed, ensuring that for each study, only one season of data could be used. Ethical approval was obtained prior to any analyses being conducted (Appendix A). The ethical approval allowed teams to be named, but not players. As such, all player names are pseudonyms within this thesis.

1.4 Thesis Scope

In order to achieve the objectives for this thesis, the current literature will be applied, adapted and extended within rugby league before a novel approach to the estimation of EPV is considered. In doing so, two types of pitch surfaces will be produced: zonal and smooth.

Zonal pitch surfaces will be produced by applying, adapting and extending work previously published in rugby league (Kempton et al., 2016). Markov models will be used to establish the value of zones on the pitch. The Kullback-Leibler Divergence will be used to empirically evaluate the most appropriate zone sizes for use in rugby league based

on their ability to reproduce future attacking performances. Empirically evaluating these zone sizes will extend previous literature (Kempton et al.,2016;Singh,n.d.), which has arbitrarily chosen its zone sizes. Initially, Markov Reward Processes will be used as their effectiveness has previously been shown in rugby league (Kempton et al.,2016) before concepts are adapted from football (Singh,n.d.) and ice hockey (Routley and Schulte,2015) to produce a Markov Decision Process model which incorporates context nodes. Novel z-score analysis of team performances will be shown, with action impact ratings adapted from ice hockey (Routley and Schulte,2015) used to evaluate player performances.

Smooth pitch surfaces will be produced to provide an understanding of the spatial trends of attacking performances and to produce a novel EPV model. First, previous work in American football (Mallepalle et al.,2020) will be adapted, applied and extended in rugby league to provide an understanding of the spatial trends of attacking performances. Kernel Density Estimation will be used to evaluate the probability that a team will perform an attacking action in any location on the pitch. These models will be produced at the whole league, between team and within team levels. The Wasserstein distance will be used to quantify differences between and within teams. Novel metrics will be produced by manipulating the cost function (normalised axis Wasserstein distance) and transport matrix (directional Wasserstein distance) properties of the Wasserstein distance. Subsequently, a novel Bayesian approach to the estimation of EPV will be shown. A Bayesian Mixture Model with custom weights for each centre calculated using linear and bilinear techniques will be used to estimate the probability of individual possession outcomes. These possession outcomes will be evaluated individually and an EPV measure will be derived using their real points scoring values. Models will be produced at the whole league, team attacking and team defending levels. Novel metrics evaluating team (expected points scored) and player (actual vs expected ratings) performances will be proposed.

1.5 Thesis Structure

Figure 1.1 is a schematic outlining a timeline for the project, the objectives and how each study fits within the literature. It shows the links between the literature and specific studies, and the applicability of different studies to rugby league using a traffic light system (red: not applicable; amber: elements applicable; green: fully applicable). Details sur-

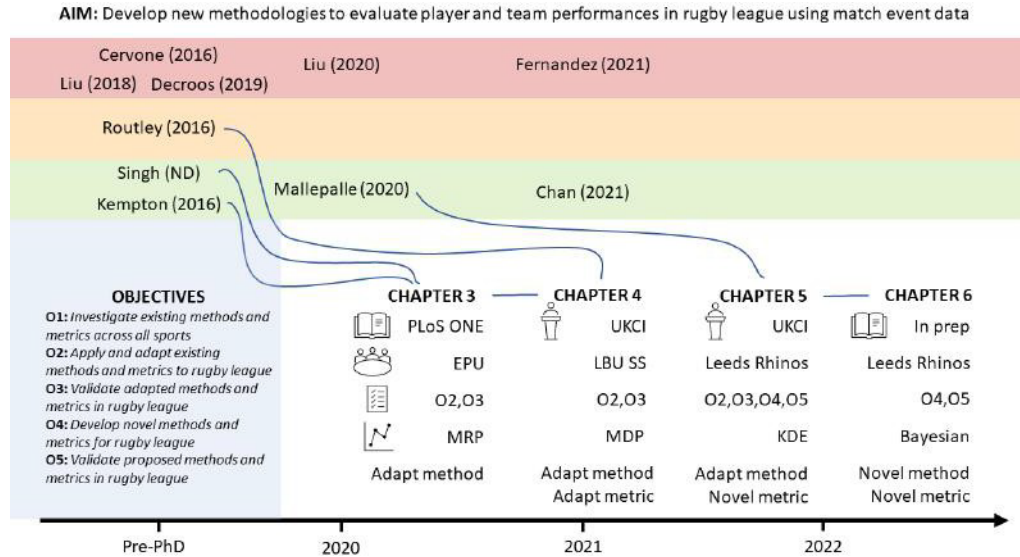


Figure 1.1: Overview of thesis outlining most important literature, timeline and how the work links to the literature. Note: Objective 1 was completed by Chapter 2, but was not included in the schematic as it is not a specific study.

rounding publications, presentations, the methods used and objectives achieved are also provided. To achieve the objectives described in Section 1.3, this PhD thesis is arranged into seven chapters and four appendices. Chapter 2 (the literature review) meets objective 1 by investigating existing methodologies in team sports; Chapters 3 and 4 meet objective 2 by applying, adapting and extending existing methodologies and metrics to rugby league; and Chapters 5 and 6 meet objectives 4 and 5 by developing novel methodologies and metrics to evaluate player and team performances in rugby league. The results from Chapters 3-4 and 5-6 were presented to coaches to evaluate their performances against objectives 3 and 5 respectively.

- Chapter 2 provides an overview of the elements of rugby league important to this thesis and investigates the sports analytics literature with respect to player and team performances in team sports. The sports analytics literature is evaluated with respect to its potential application to rugby league. The limitations of the rugby league literature are discussed.
- Chapter 3 describes the adaptation of an expected possession value (EPV) model

to analyse team attacking performances in rugby league using play by play data and a Markov Reward Process approach. The locations of 59,233 plays from 180 matches in the 2019 Super League season were used. The study applied, adapted and extended previous work in rugby league (Kempton et al.,2016) and football (Singh,n.d.). The Monte Carlo every visit algorithm was used to simulate EPVs for the EPV-308. Five further models were constructed from the EPV-308 and the KL Divergence was used to establish the models' ability to reproduce future team attacking performances. A novel method to highlight the value generated by teams in different zones on the pitch was proposed. It was published in PLOS One (Thomas Sawczuk, Anna Palczewska, Ben Jones. Development of an expected possession value model to analyse team attacking performances in rugby league. *PLOS One*, 16(11): e0259536, 2021) and presented to the England Rugby League Performance Unit for feedback.

- Chapter4 extends the EPV model from Chapter3 to analyse player and team performances using action by action data in rugby league for the first time. The locations of 77,045 actions from 63 matches in the 2020 Super League season were used. A Markov Decision Process approach, via the Monte Carlo every visit algorithm, was used to estimate action values conditional on their location. Context nodes were adapted from previous research in ice hockey (Routley and Schulte,2015) to provide further insights into player performances and allow individual team analyses to be conducted. An action impact rating was adapted from ice hockey (Routley and Schulte,2015) to evaluate the performances of players based on the actions they took. The work was presented at the UKCI 2021 conference and published in Advances in Computational Intelligence Systems (Thomas Sawczuk, Anna Palczewska, Ben Jones. Markov Decision Processes with contextual nodes as a method of assessing attacking player performance in rugby league. In: Jansen, T., Jensen, R., Mac Parthalain, N., Lin, CM. (eds) Advances in Computational Intelligence Systems. *UKCI 2021*. Advances in Intelligent Systems and Computing, vol 1409. Springer, Cham.). It was presented to the Leeds Beckett University Sport Science Department for feedback.
- Chapter5 evaluates the use of kernel density estimation (KDE) and the Wasserstein distance to understand the spatial trends of attacking performances in rugby league

between and within teams. 99,966 actions from 138 matches in the 2021 Super League season were used. The study applied, adapted and extended previous work in American football (Mallepalle et al., 2020) by using KDEs to understand the location densities of team actions across the whole league season and against specific opponents. Two novel metrics (the normalised axis Wasserstein distance and directional Wasserstein distance) were computed from the Wasserstein distance to compare the 2-dimensional location densities between and within teams. The results showed that differences between and within teams could be evaluated and summarised in a simple graph for benchmarking purposes. The work was presented at UKCI 2022 and will be published in a future edition of *Advances in Computational Intelligence Systems*. It was presented to coaches at Leeds Rhinos Rugby League club for feedback.

- Chapter 6 proposes a novel Bayesian approach to producing a smooth EPV model, which allows for better evaluation of player and team performances. 99,966 actions from 138 matches in the 2021 Super League season were used. A Bayesian Mixture Model approach, adopting centres and the use of linear and bi-linear interpolation techniques rather than zones, was used to generate an EPV model, which individually estimated the probability of each possession outcome in the same model for the first time. Models were produced at the whole league, individual team attacking and individual team defending levels. This approach was able to differentiate between the strengths and weaknesses of different teams across the pitch. It was combined with the KDE's from Chapter 5 to generate a novel expected points scored metric. Player performances were evaluated by a novel actual vs expected contribution to the final possession outcome metric. The work was presented to the Leeds Rhinos coaching staff and will be prepared for publication in a sports journal.
- Chapter 7 summarises the outcomes of this thesis. It highlights the contributions of each chapter, shows how they met the objectives of the project and considers future directions for the body of work.

Literature Review

This chapter presents a detailed discussion of the literature relevant to the objectives of this thesis. First, the types of data collected by professional sports teams are introduced. Next, the use of this data to evaluate team and player performances in the wider context of sports excluding rugby league is discussed. Rugby league and its rules are then introduced, and the literature evaluating player and team performances within the sport is evaluated. The summary section outlines the difficulties with adapting some elements of previous research in other sports to rugby league and identifies those methods which could be adapted to evaluate player and team performances in rugby league. Finally, a high level computational overview of the methods utilised in this thesis is provided.

2.1 Data in Team Sports

Recently, there has been a large increase in the quantity of data produced by sports teams regarding their players' performances in training and match play. The data obtained by these teams can broadly be split into two areas: physiological data and tactical data. Physiological data provides information surrounding the body of the athlete and can be used to understand how hard the athlete is working physically (McArdle et al.,2015). Tactical data provides information surrounding the technical and tactical events (e.g. passes, runs, kicks) that take place and can be used to understand the match play performances of players or teams (Hughes et al.,2019).

Physiological data is typically collected using wearable technologies, including global

positioning system (GPS) units, heart rate monitors and instrumented mouth guards. GPS units use satellite-based radio navigation to understand the location of players during training or match play. From this information, practitioners can understand how hard a player has worked in terms of total running demands, accelerations, decelerations or amount of time spent at/above/below a given running speed threshold. Heart rate monitors are typically used alongside GPS units to understand objectively how hard a given workload is for the athlete. This is achieved by comparing the heart rate to the workload the athlete has performed. Instrumented mouth guards are mouth guards fitted with accelerometers, which are worn by players to allow practitioners to quantify the magnitude of head impacts. With the exception of instrumented mouth guards, which are a recent innovation, wearable technologies have been available to sports practitioners for approximately 20 years. Consequently, there is a widespread body of research understanding the information these technologies can collect and how they can be used to help improve training practices. For example, research of this nature has been conducted across multiple sports ([Whitehead et al.,2018](#)) or specific to individual sports such as football ([Whitehead et al.,2018](#)), hockey ([Gabbett,2010](#)) or rugby ([Dalton-Barron et al.,2020](#)).

Tactical data is collected via event level data annotation and is usually limited solely to match play situations. The two key methods of annotating event level data are human annotation and computer vision technologies ([Pappalardo et al.,2019](#); [Vats et al.,2020](#)). Human annotation involves a human watching a video of a match and annotating information surrounding each event that occurs. This annotation can take place either live (i.e. as the match is happening) or retrospectively (i.e. after the match has been completed). The data collected typically contains an estimated location of the event and details regarding the event type, the player who completed it and the event outcome. Computer vision technologies have recently been employed in some early adopting sports to enhance the detail of event level data annotation. For example, in basketball, it is now possible to understand the location of every player on the court using multiple cameras located around the stadium perimeter ([Barker,2016](#)). Similarly, in football, it is possible to know the location of every player in the same television frame as an action being performed if the television feed is available ([StatsBomb,2022](#)). The same computer vision technology allows players to be tracked and provides a method through which their locations can be collected for tactical analyses. Such information is extremely valuable as GPS units cannot estimate a player's location accurately enough to be used for any spatial data based tactical

2.2 Team and Player Performance Analysis in Sport

analyses. Unlike wearable technologies, the availability of high quality event level data is much more varied between sports. In those sports where computer vision technologies have been employed, there is excellent data availability in the quality and quantity required to perform advanced modelling techniques such as deep learning (Decroos et al., 2019; Fernández et al., 2021) or deep reinforcement learning (Liu and Schulte, 2018; Routley and Schulte, 2015). In these sports where there is high data availability, some lower quality data has been made publicly available (Pappalardo et al., 2019), which has likely increased researchers' interest and boosted research in the sport. However, in sports where event level data is limited to human annotation, the data availability is considerably reduced and the analyses conducted are much more limited (Croft et al., 2017; Irvine and Kennedy, 2017; Parmar et al., 2018). Furthermore, with limited data available, data providers do not release their proprietary data freely so research within the sport is not as common because the data is less freely available.

2.2 Team and Player Performance Analysis in Sport

The two most important elements of tactical analyses within sport are understanding *team performances* and understanding *player performances*. Team performances can be measured in many ways (e.g. Did the team win it's match? Did the team score more tries than the opposition? Did the team complete more passes or tackles than the opposition?) but usually analyses of team performances are used to provide an underlying understanding of how well the team met its objectives across a set period of time (e.g. a match or season), without considering match results. Evaluating team performances can be used in many ways within sport including: identifying key performance indicators for teams; understanding key sequences of play for a given team; and evaluating likely areas on the pitch through which a team may create value. All three of these elements can provide significant tactical understanding as coaches attempt to improve their own teams and prepare strategies to beat opposition teams. Player performances refer to how much the actions completed by a player contributed to improving a team's performance. Player performances can be measured using standard count data (e.g. number of passes made, number of try assists) or using some form of value metric (e.g. actual goals vs. expected goals scored). Evaluating player performances can help a coach to understand whether a player is producing an output that is good or bad value for the cost of their salary. It can also

2.2 Team and Player Performance Analysis in Sport

help to identify transfer targets (i.e. players who may perform better than those currently at the club).

Although one of the earliest examples of analysing team tactical performances was in football (Hughes,1990), the underlying idea of utilising statistical analysis to evaluate player performances was popularised by the Oakland Raiders baseball team in Major League Baseball in the early 2000s (Lewis,2003). Their 'Sabermetrics' approach allowed them to far exceed the performance their budget should have provided by using analytics to understand underrated players. They did this by redefining the elements of player performances perceived to be valuable. For example, rather than considering hit percentages (Equation2.1) as the best measure of batting performance, they instead considered on-base percentage (Equation2.2). Their reasoning for this was that running around bases is how a team wins matches and this can be achieved in more ways than just hitting the ball. Similarly, hitting the ball does not necessarily lead to reaching any given base. Equations 2.1 and 2.2 can be defined as

$$\text{Hit\%} = \frac{\text{Hits}}{\text{At-bats}}, \quad (2.1)$$

where an at-bat is a player's turn to bat, and the total number of both hits and at-bats are used; and

$$\text{On-Base\%} = \frac{\text{Hits} + \text{Bases on balls} + \text{Hit by pitch}}{\text{At-bats} + \text{Bases on balls} + \text{Hit by pitch} + \text{Sacrifice flies}}, \quad (2.2)$$

where bases on balls refers to a batter being awarded a walk to first base due to poor pitching, hit by pitch is another situation where a batter is awarded a walk to first base and sacrifice flies refer to a batter hitting the ball in such a way that his chances of scoring are diminished markedly, but a team-mate (usually on third base) has an enormously improved chance of scoring. The total number of all measures is used.

2.2.1 Accumulated Count Data Analysis in Team Sports

Following the success of the Oakland Raiders, other sports have conducted similar analyses using accumulated count measures in an attempt to evaluate player and team performances. Accumulated count measures refer to counts of match actions, which are accumulated over a specific time period (usually a match) and are used to provide key performance indicators for coaches. These performance indicators can be used as bench-

2.2 Team and Player Performance Analysis in Sport

marks through which coaches evaluate changes to their tactical strategies (Herold et al., 2021).

Puente et al.(2015) used season accumulated basketball statistics to identify performance indicators for teams who made the playoffs in the Spanish Basketball League. They used 10 seasons of Spanish Basketball League data between 2003 and 2013, including a total of 3,060 basketball games. The data were preprocessed so that game-related statistics were accumulated at the season level (i.e. total counts of each statistic for each season were used). The final league ranking of the team and the number of wins within the season were also identified. Pearson correlation (Rodgers and Nicewander,1988) was used to assess the association between game-related statistics and the number of wins during the regular season. Multiple regression analysis (Jobson,1991) was used to provide a model explaining the number of in season wins using the variables identified as significant via the univariate Pearson correlation. One Way ANOVA (Fisher,1925) was used to evaluate whether there was a significant difference in game related statistics between playoff and non-playoff teams. The authors showed that the combination of all three forms of shooting accuracy (2 point, 3 point and free throw) accounted for 26% of the variance in the multiple regression model, with rebounds accounting for a further 23%. In total, the model accounted for 76% of the total variance in the number of wins achieved during a regular basketball season. Similar results were shown by the one way ANOVA, which indicated that 2 and 3 point, but not free throw, shooting accuracy were significant differentiators between those teams who did and did not reach the end of season championship playoffs. The results quantify the importance of having accurate shooters and proficient rebounders within a basketball team in the Spanish Basketball League and so could be used as a bench marking method by coaches when considering transfer targets.

Harrop and Nevill(2014) used match accumulated football statistics to identify team performance indicators with respect to winning football matches in League One (the third tier of English football). They used 46 matches from the same team in the 2012/13 season. No preprocessing was required as the raw match counts of several attacking and defensive variables were used, alongside the fixture outcome (win, loss or draw). Binary logistic regression (Menard,2009) (win vs loss and win vs draw in separate models) was used to identify those performance indicators important to winning matches. The results indicated that completing fewer passes in total ($P=0.006$, odds ratio(OR)=0.972) and having a higher percentage of successful passes ($P=0.042$, OR=1.243), were significantly related

2.2 Team and Player Performance Analysis in Sport

to match winning performances. Attempting more shots ($P=0.027$, $OR=1.296$) and fewer dribbles ($P=0.018$, $OR=0.695$) were also significant factors in the model, which correctly classified the result of 71.7% of in sample matches. No out of sample analysis was conducted. The results provide important bench marking information for the club in question, which their coach could use to understand their underlying performances but the applicability to all teams within the league is limited due to the single club sample used.

Woods(2016) used season accumulated Australian Rules football statistics to identify team performance indicators relevant to the team's final league position. He used 195 matches from the 2015 Australian Football League season. During preprocessing, the match counts of 12 performance indicators were summed to provide a season count for each performance indicator. Cumulative link mixed models (Hedeker and Gibbons, 1994) were used to evaluate the relationship between performance indicators and end of season league position. The results indicated that three performance indicators were associated with final league position: clearances ($P=0.002$, $\beta=-0.032$); hit-outs ($P<0.001$; $\beta=-0.034$); and inside 50m ($P=0.012$, $\beta=-0.020$). The negative associations for all three variables indicated that greater counts were indicative of a better final league position. The results provide a bench mark for performances that coaches may wish to use; however the violin plots provided by the author in the study (Figure2.1) show the difficulty with using any kind of accumulated count measure as a performance indicator due to the wide variability inherent within these counts on a match by match basis.

Robertson et al.(2016) used match accumulated Australian Rules football statistics to evaluate the influence of individual player performances on match outcomes. They used 198 games from the 2014 Australian Football League season. The data were preprocessed by transforming the match accumulated player statistics for 13 performance indicators into percentages of the team's overall match count for the respective performance indicator. These percentages were then used as ordered weights for each player so fourteen features were included for each performance indicator based on percentile ranks of the ordered weights. That is, the percentile ranks described the proportion of the team's overall performance indicator count accounted for by the relevant player. Generalised estimating equations (Hardin and Hilbe,2003) were used to describe the relationship between match outcome and the feature set of performance indicators. The authors used 10-fold cross validation with a random selection of 33% of the data; they did not specify whether this was an in or out of sample validation. The results showed that five performance indica-

2.2 Team and Player Performance Analysis in Sport




Figure 2.1: Woods(2016) violin plot of performance indicators significantly associated with final league position. Model was completed using season accumulated counts, dots in this plot represent match accumulated counts and show the variability inherent within the use of accumulated match counts on a match by match basis. League position is denoted in alphabetical order (e.g. A = 1st, B = 2nd etc).

tors were significantly related to match outcome at different percentile ranks: marks (25th percentile); disposal (25th and 50th percentile); goals (75th, 90th and 95th percentile); behinds (90th percentile); and inside 50s (95th percentile). The model correctly classified the match outcome with $63.9 \pm 4.2\%$ accuracy. The results provide a novel approach to understanding the distribution of player contributions to match outcome, suggesting that the more evenly spread the goal scorers are and the more disposals (total count of kicks and handballs) the lower percentile players perform, the greater the team's chance of winning. Information of this nature could be used to provide an understanding of squad depth and positions/areas through which recruitment teams may wish to target new players.

Whitaker et al.(2021) used match accumulated football statistics to infer the abilities of football players in the Premier League, working under the assumption that players involved in a higher quantity of actions are better at those actions (i.e. a player who passes more in a match is a better passer). Their data set included 1.2 million actions from the 2013/14 - 2014/15 seasons of the English Premier League. Match accumulated counts of actions, including passes and shots, were calculated for each player during preprocessing.

2.2 Team and Player Performance Analysis in Sport




Figure 2.2: Whitaker et al.(2021) posterior plot of goal abilities (i.e. those actions attempting to score goals) for Sturridge and Reed.

Also produced were a dichotomous home/away variable, identifying the location of the match, and a variable denoting the fraction of time played by the player. A Bayesian model following a Poisson distribution was utilised to consider latent player abilities, while accounting for the player and opponent teams' influence on typical match action counts. A variational inference approach was used to estimate the model parameters. The authors showed how increased playing time could help shift parameter estimates away from the mean level, and provided player rankings for two categories of actions: Goal (goal actions) and GoalStop (actions attempting to stop a goal). Figure 2.2 is taken from Whitaker et al.(2021) and shows two different player's goal abilities. It shows how the model is much more certain about the positive ability of Sturridge with respect to scoring goals because he is involved in more goal scoring actions. Conversely, there is much greater uncertainty surrounding Reed's scoring ability as he is involved in much fewer goal actions. Although the underlying assumption of the model that players involved in more actions are better at those actions is open to debate, the results of the study provide a novel methodology through which players can be rated and the certainty with which they are rated. This information could be used by coaches and scouts to evaluate their players with much greater detail than previously considered within the literature, where mean point estimates are typically considered. In sports with low data availability, such as rugby league, this Bayesian approach could also allow the uncertainty caused by low data availability to be estimated providing an appropriate method through which player and team performances could be evaluated.

2.2.2 Action by Action Data Analysis in Team Sports

The use of accumulated count data to evaluate player and team performances in sports is useful to the extent that it provides a range of key performance indicators through which match or season performances can be compared. Furthermore, it provides a bench mark through which coaches can evaluate tactical changes or player recruitment. However, the information provided to coaches is limited by the nature of accumulated count measures, which don't consider the context surrounding each action individually. They are therefore unable to provide a true indication of the value of an action being taken, or the probability of its success. For example, in football, a pass to a team mate in the opposition half is more likely to result in a goal than a pass to a team mate in the possession team's own half but a pass to an unmarked team mate in the opposition half is more likely to result in a goal than a pass to a marked team mate. Similarly, a pass to a team mate in the centre of the penalty area is more likely to result in a goal than a pass to a team mate wide on the pitch, but is also more likely to be unsuccessful. This information is not considered in accumulated count data analyses which, until recently, was a barrier to the widespread usage of advanced statistical analyses to rate player and team performances in team sports.

The recent introduction of *action by action* data, which provides spatial information for every action taking place on the pitch has changed the focus of research towards the development of more advanced statistical analyses to evaluate player and team performances. However, such data is only readily available in a limited number of sports. This section considers two types of these analyses: *spatial trends* analyses, which attempt to understand the pitch as a probability density function of where the player or team could pass or control the ball; and *expected possession value* (EPV) models, which provide a value for every location on the pitch based on the probability that a positive or negative scoring action will occur within a predefined period of time.

2.2.2.1 Spatial Trends Analysis in Team Sports

Despite its importance and ability to directly influence the decisions that players and coaches make, literature surrounding spatial trends analyses (Definition 2.2.1) is surprisingly sparse. In football, two studies have considered the concept of pitch control (Fernández and Bornn, 2018; Spearman et al., 2017). Pitch control can be defined as the degree or probability of control that a given player (or team) has on any specific

2.2 Team and Player Performance Analysis in Sport

position/location in the available playing area (Fernández and Bornn,2018). In American football, individual player pass location distributions have also been considered (Mallepalle et al.,2020) as a form of spatial trends analysis. To the authors knowledge, no other studies have been conducted in sport specifically considering the pitch as a probability density function to evaluate the spatial trends of player or team performances.

Definition 2.2.1. Spatial trends analyses estimate the probability that a team will control the ball in a given location or where a player or team is likely to pass the ball.

Spearman et al.(2017) adopted a physics-based approach to evaluating pitch control using match events and player tracking data from 38 matches played by a single club in the Premier League in the 2015/16 season. During preprocessing, the authors identified 10,875 passes to be used within the training and testing of their model. For each pass, the authors compute the ball's trajectory (i.e where the ball travelled), the time-to-intercept for each player for every location on the pitch (i.e. how long it would take a player to reach that location) and the time-to-control (i.e. the probability that a player can control the ball given how long he will be in its vicinity). A Bayesian model was used to estimate the probability of a pass being successful for every location on the pitch. The model had an overall accuracy of 80.5% when predicting pass success or failure and 67.9% when predicting the specific receiving player. A pitch control plot fromSpearman et al. (2017) is shown in Figure2.3. The plot shows areas of control for each team on the left, identifying where the player in control of the ball may have wished to pass the ball. On the right, it shows the area of pitch control for a single player (blue number 7) as a dark shaded area.

Fernández and Bornn(2018) evaluated pitch control using only player tracking data in football. Their approach used 2.4 million examples from 20 Spanish First and Second Division matches. During preprocessing, they extracted the x,y player location coordinates, speed and direction of movement, as well as their distance from the ball at a frame by frame level. Player influence areas were defined using a bivariate normal distribution, which accounted for the speed and direction of movement of the player. The influence area was scaled by the distance of the player to the ball so that players further away from the ball had less influence than those close to it. These areas were summed for each location at a team level to identify the overall team pitch control for the possession or opposition team. As only tracking data was used, no accuracy metrics were provided, but

2.2 Team and Player Performance Analysis in Sport

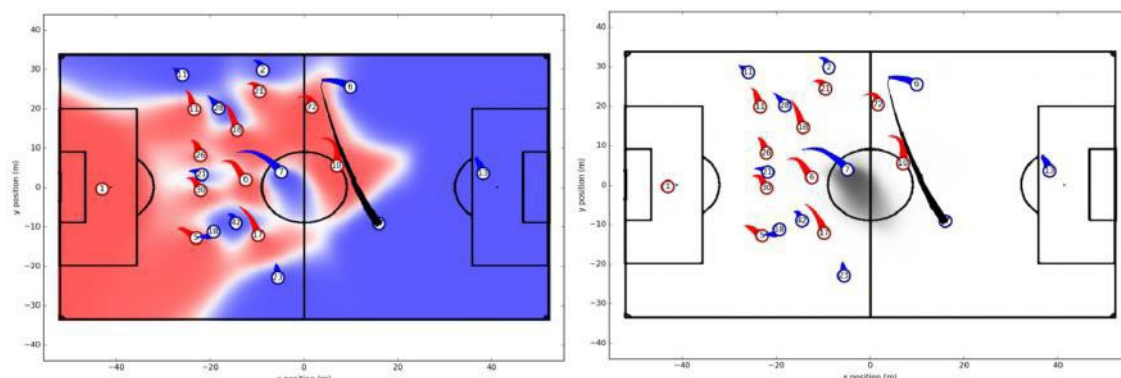


Figure 2.3: Spearman et al.(2017) pitch control probability plot. On the left, blue regions are those controlled by the team in possession of the ball, red regions are controlled by the opposition team. White areas are contested (i.e. not controlled by either team). On the right, the shaded region represents the area of the pitch controlled by blue number 7. The trail behind every player, the ball represents their position over the past 3 seconds.

the authors did show how the model could be used to evaluate the space gained or lost through player movements. Figure 2.4 provides a plot showing this usage.

Both studies considering pitch control (Fernández and Bornn, 2018; Spearman et al., 2017) provide extremely valuable insights into the concept of space occupation on a football pitch. They provide methods through which each pass or run can be valued with respect to its ability to generate space and compared to other options the player may have had. They also allow for the possibility that player profiles can be generated based on how risky the passes they make are. However, the ability to implement the models in rugby league is limited for two reasons. First, data availability - there is no tracking data available in rugby league. Second, and perhaps more importantly, the rules of rugby league limit the usefulness of this method. In rugby league, the ball can only move forwards through either a run or a kick. If the ball is thrown (the most common passing method), it must move backwards. Furthermore, defending players are almost always in front of attacking players, so when an attacking player throws the ball backwards, there is usually very little chance of it being intercepted by the defending player. As such, the pitch control model would mostly only be useful in rugby league when considering kick passes, which occur approximately 20 times less frequently than passes.

Away from the concept of pitch control, Mallepalle et al. (2020) also considered the

2.2 Team and Player Performance Analysis in Sport

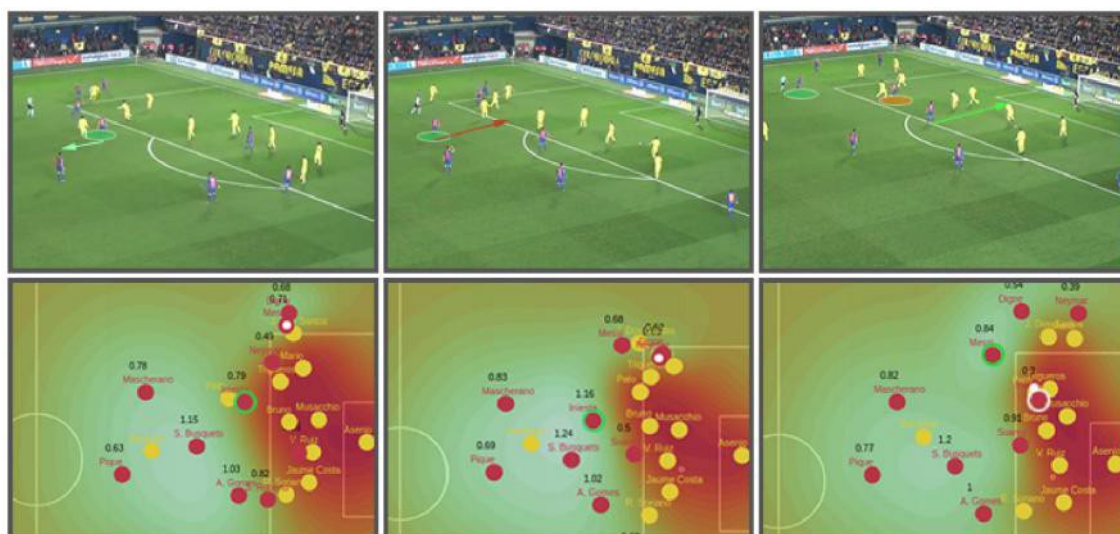


Figure 2.4: Fernández and Bornn(2018) pitch control plot. In the left frame, a player moves backwards from the 18 yard line to occupy a space of value with higher control. In the middle frame, the player attacks a space in front of him, dragging three defenders towards him. In the third frame, the green circled player is free because of the movement in the middle frame. The player on the 18 yard line makes a forward run into space so he can receive a crossed pass.

pitch as a probability density function by calculating the pass location probability distribution of quarterbacks in American football. The authors used 27,171 quarterback passes from the 2017 and 2018 National Football League seasons. The data used was already centred at the line of scrimmage and so the only preprocessing required was to limit the pitch to 10m behind the line of scrimmage and 55m in front of it. The authors used kernel density estimation (KDE; Rosenblatt(1956)) to provide a smooth passing probability surface and generalised additive models (Hastie and Tibshirani,1986) to estimate the probability of the pass being successful. As a purely observational study, no accuracy figures were reported. The authors showed how the KDE can be used to identify passing trends at a player or team level and how the generalised additive models can be used to evaluate players' or teams' completion rates compared to the league average. Figure 2.5 shows these comparisons for two quarterbacks. The left column shows the player's KDE, whereas the player's predicted completion probability surface is shown on the right. The idea of utilising KDEs to establish the spatial trends of player performances is unlikely




Figure 2.5: Mallepalle et al.(2020) quarterback KDE plot (left) and pass completion probability surface (right).

to work particularly well in rugby league given the wide variety of areas from which a player could receive the ball, which differs from the quarterback who always receives the ball in the centre of the pitch. However, there is significant potential to use the method to try and understand the spatial trends of team attacking performances, which may help reveal tendencies to move the ball in certain areas of the pitch. Surprisingly, no study has attempted this form of analysis at a team level in any sport so there is scope for its consideration within rugby league.

2.2.2.2 Expected Possession Value Models in Team Sports

Identifying the spatial trends of player and team performances provides an interesting, useful and visually appealing representation of the tactical strategies of players or teams as it evaluates the locations on a pitch they are likely to try and control or move the ball. However, the key limitation of these spatial trends analyses is that the pitch is treated as a uniform entity where players and teams are equally likely to wish to move the ball in any direction. In invasion team sports, such as basketball, football and rugby, where the opposition scoring area is at one end of the pitch, this is not the case as most movements made by a team are aimed at progressing towards the scoring area. Those areas closer to the opposition scoring area are therefore intrinsically of higher value as a team is more likely to score points/goals the closer they get to the scoring area.

The concept of EPV was introduced by [Cervone et al.\(2016\)](#) in basketball to quantify the value of controlling the ball in any location on the pitch (Definition 2.2.2). Unlike spatial trends analyses, EPV has received significant attention within the sports analytics literature and has since been adapted to football ([Decroos et al.,2019](#);[Fernández et al., 2021](#);[Singh,n.d.](#)) and ice hockey ([Liu and Schulte,2018](#);[Routley and Schulte,2015](#)), amongst other sports ([Bukiet et al.,1997](#);[Chan et al.,2021](#)). However, the name of the measure, the method used to calculate it and its subsequent usefulness to decision making in sports is highly variable. Broadly, four methods have been used to analyse EPV models: *deep learning* approaches; *Markovian* approaches (e.g. reinforcement learning); *machine learning* approaches (e.g. CatBoost algorithms); and *Bayesian* analysis.

Definition 2.2.2. EPV provides an instantaneous snapshot of the possession's value ([Cervone et al.,2016](#)). This value is typically measured in terms of the probability of scoring goals/points conditional on the location.

[Cervone et al.\(2016\)](#) used tracking and event level data to produce an EPV model in basketball. The model analysed over 1 billion space-time observations from the 2013/14 NBA season. The data provided to the authors for analysis were pre-processed to identify the player in possession of the ball, their speed and acceleration, whether they were defended by the opposition at that time point and the court zone they were in. The authors produced a multi-resolution stochastic model following semi-Markovian principles, which allowed movements of the ball (i.e. match events) and movements of players (i.e. tracking data) to be analysed together with respect to their impact on EPV. In order to

2.2 Team and Player Performance Analysis in Sport

make the model computationally tractable, player data was grouped together using player similarity matrices, allowing the authors to share spatial information regarding action preferences between players within their model. The model was able to provide an EPV for every action or movement on the court. Figure 2.6 visualises this process, showing the changes in EPV that occur during a passage of play. The authors showed how the model could be used to rate player performances using an EPV-added estimate, which compared the EPV when the player received the ball to the EPV when the player had completed an action or lost the ball. They used this to provide a ranking for the best and worst players in the 2013/14 season. The model provided a comprehensive introduction to the concept of EPV, with elements which would work in rugby league (e.g. adopting Markovian principles to value multiple point scoring actions at the end of the possession with a single number). However, data availability issues in rugby league make it difficult to implement the model directly in the sport (no tracking data is available in rugby league). Furthermore, the method through which the EPV remains constant while an action takes place is not suited to rugby league. Whereas in basketball this works because actions (e.g. passes or shots) are typically less than 1 second in duration, in rugby league kick action can frequently take 3-5 seconds as a team attempts to gain territory. Consequently, holding the EPV constant for this duration of time, while the ball and players progress as far as 60 metres down the pitch is inappropriate for the requirements of rugby league.

Fernández et al.(2021) produced a comprehensive deep learning EPV model in football. They used 633 English Premier League matches, which took place between the 2013/14 and 2014/15 seasons. A significant amount of preprocessing was used to convert the x, y player location co-ordinates and match event level data into appropriate features for their model. Rather than solely providing the model with the player's coordinates on the pitch, information such as the distance and angle to the opposition goal was included, alongside the player's location relative to the ball. Furthermore, contextual information was included, for example the player's pitch control (i.e. the probability they would control the ball if it was moved to a given location), the pressure lines of the opponent (i.e. information about their formation) and the number of players who would be outplayed if the ball was moved to a given location. The authors employed a 'decomposed Markov Decision Process (MDP)' approach, which allowed them to individually model each element of the sport (pass, kick and run) using the features most relevant to the action specified - such an approach would not be possible if modelling via MDP or reinforcement learning

2.2 Team and Player Performance Analysis in Sport




Figure 2.6: Cervone et al.(2016) EPV plot for a sample possession. Immediate, rapid changes are caused by actions such as passes and shots. Smaller, more incremental changes in EPV are induced by changes in players' locations. The black line slightly smooths the EPV evaluations at each time point (grey dots).

as the value of all actions is estimated together within that framework. Machine learning methodologies appropriate to the decomposed question were used: logistic regression was used to identify pass and turnover probabilities; action likelihoods used convolutional neural networks built on the pitch control model (Fernández and Bornn,2018) described in Section2.2.2.1; and pass and ball drive expectations were calculated using deep neural networks. Figure2.7provides a sample of the model's usage.

The authors provide multiple applications for the model. They use an EPV added estimate similar to that employed byCervone et al.(2016) but extend the model to consider the EPV added when passing to a different players (Figure2.8). They also discuss how the model can be used to evaluate the EPV generated by teams in different formations against a specific opponent, providing insights into how to prepare strategies against an opposition team (Figure2.9). The work ofFern ández et al.(2021) provides an excellent gold standard approach to the use of deep learning within sport as it breaks down different elements of football to gain incredible insights into the sport at both a player and team level. In its current guise it isn't suitable for use in rugby league, as it utilises a binary outcome measure (goal/no goal) whereas in rugby league there are 5 possible possession outcome, but adapting the model to this scoring difference would not prove difficult. A

2.2 Team and Player Performance Analysis in Sport

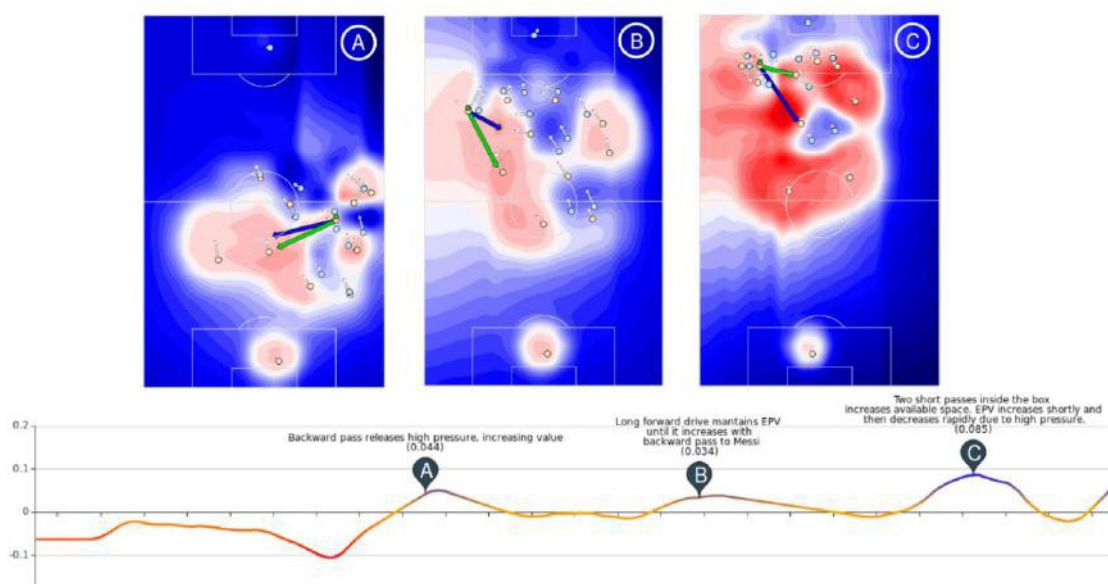


Figure 2.7: Plot outlining the decomposed MDP methodology employed by Fern ández et al.(2019,2021). Figure taken from Fern ández et al.(2019). EPV plot highlights three specific situations (A, B and C). Blue areas represent areas of low EPV, red areas represent high EPV. Blue arrows represent pass chosen by the player, green arrows represent best possible pass according to the model. Yellow dots represent players on the attacking team (i.e. in possession), blue dots represent defending team.

larger issue is the quantity and quality of data required to produce the models, which isn't available in rugby league where there is no tracking data available of any type and action by action location data has only been available for two seasons (totalling 201 matches). As such, although the work of Fern ández et al.(2021) is a gold standard to aspire to, it is not possible to reproduce in rugby league with the current state of data availability.

Decroos et al.(2019) valuing actions by estimating probabilities (VAEP) framework provides an alternative method through which football actions can be evaluated, using match event data only. The authors used 11,565 matches from five top European football leagues between 2012-2018 (totalling 14.4 million actions). The data were preprocessed to provide model features related to the previous 3 actions. These features were grouped into descriptive features (e.g. type of action, location of action, body part used and action result), complex features (e.g. the distance and angle to the goal and the distance covered during the action in x and y directions) and game context features (e.g. number of goals

2.2 Team and Player Performance Analysis in Sport




Figure 2.8: Fernández et al.(2021) evaluation of single player's EPV added when passing (left) and receiving (right) passes from different players. Plot shows probability of passes being selected by size of player circle and value of pass through the arrow size and colour. Metrics are normalised by the minutes the two players played together and multiplied by 90 minutes to provide a match value.

scored by the player's team and opponent in the game after a specific action). A CatBoost algorithm was used to evaluate the probability of scoring and conceding within the next 10 actions in separate models. The results were not visualised on a pitch due to the context features involved, but the authors showed how the model could be used to identify potentially undervalued player. Furthermore, they showed how the model could be used to provide a deeper understanding of players' ability to provide value through the different actions they attempt (Figure 2.10).

Recently, the model has been extended to show how the VAEP framework can be used to provide player ratings based on the pressure experienced by a player within a game state by manipulating the context features (Bransen et al., 2019) and to provide an understanding of the chemistry between players on the same team, where chemistry is

2.2 Team and Player Performance Analysis in Sport

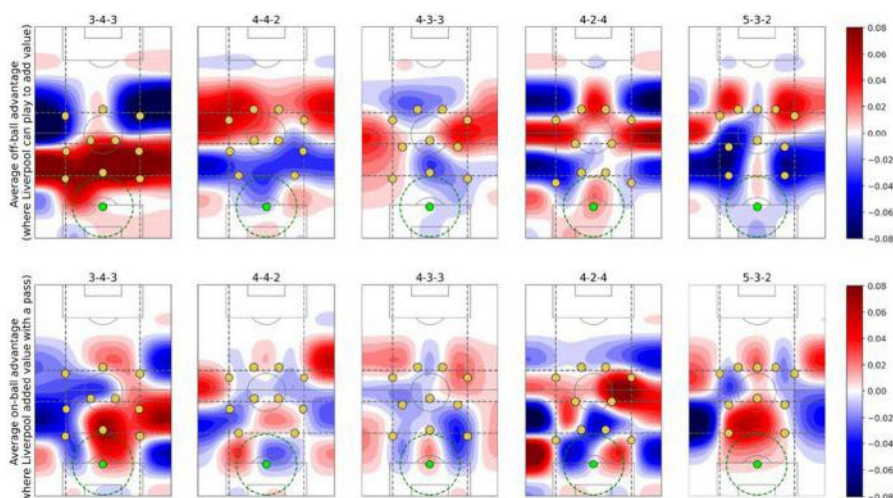


Figure 2.9: Fernández et al.(2021) plot identifying the EPV generated by teams in different formations against a specific team. The green circle represents the ball location.




Figure 2.10: Decroos et al.(2019) plot of VAEP for different actions. Plot shows an overview of the total contribution per 90 minutes for different types of actions for different players in the 2016/17 (a) and 2017/18 (b) seasons.

defined as the joint-value generated between two players (Bransen and Haaren,2020). Figure2.11depicts an attacking chemistry plot provided byBransen and Haaren(2020). The VAEP is an interesting approach, which accounts for a wide variety of action details,

2.2 Team and Player Performance Analysis in Sport




Figure 2.11: Bransen and Haaren(2020) joint chemistry plot showing value generated between pairs of players on the same team. Green lines represent high attacking chemistry links, dark red lines reflect low attacking chemistry links.

including the context of the match. These context features could provide interesting insight if used within rugby league. However, in the model's current form, it would require an impractical eight models to be calculated based on the points which can be scored in rugby league: two each for converted try; unconverted try; penalty goal and drop goal (one for the possession team and one for the opposition team). Furthermore, the amount data used by the authors to calculate this model was more than 10 times greater than that available in rugby league, limiting the ability to replicate the method in the sport.

Routley and Schulte(2015) utilised a Markov Game model (Shapley,1953) approach to determine the value of each action within an ice hockey match by estimating its contribution to the probability of either team scoring the next goal or conceding the next penalty. They used 2.8 million actions from the 2006/07 NHL season. In order to convert the event level data into an MDP framework, a significant amount of preprocessing occurred. Episodes were defined as the period of match play between goals, penalties or stoppages and were described by three contextual features (goal differential, manpower (i.e. number of players on each team) differential and period of match). Eight actions

2.2 Team and Player Performance Analysis in Sport




Figure 2.12: Routley and Schulte(2015) box plot of action impact values across different contexts. Left shows the action impact values for goals, where a high value is better for the team. Right shows the action impact values for penalties, where a high value is worse for the team. Central mark of the box plot is the median, the edges represent the 25th and 75th percentiles and the whiskers are at approximately 2.7 standard deviations.

were isolated and provided with two markers (the team performing the action and the zone in which it occurred). A state was then described by the contextual features and a playing history (i.e. each action that took place within the game state). Consequently, each state essentially played the role of an episode within their model, as all the actions taken within the episode were described in the state. A reward of 1 was provided when a goal was scored or a penalty was received. The authors used dynamic programming via the value iteration algorithm to calculate the values of different actions within the model. They then produced box plots of action impact values (Figure 2.12) to show their variation across different contexts. Players were rated by comparing the action impact value of the action taken by a player to the state value. The Markov Game approach used by Routley and Schulte(2015) is not suitable for rugby league because a tackle by the opposition team does not mean they will receive possession of the ball. However, as with Decroos et al.(2019), the use of context nodes to identify match situation and value player actions is a particularly interesting concept, which could be used within rugby league.

Liu and Schulte(2018) advanced the work of Routley and Schulte(2015) in ice hockey by using a deep reinforcement learning algorithm (Sutton and Barto,2018) to learn the model's action values. They used 3.3 million actions from the 2015/16 NHL season. A similar, but more comprehensive, set of features to Routley and Schulte(2015) was

2.2 Team and Player Performance Analysis in Sport

produced during preprocessing. The reinforcement learning approach allowed thirteen actions to be considered (rather than eight), alongside ten features (including continuous x and y co-ordinates, rather than zones, an infinite score differential (rather than -8 to +8 values) and other elements such as the angle between the puck and goal and the velocity of the puck). Unlike their previous work, the authors focused solely on goals scored/-conceded for their rewards in this study, with 1, 0 or -1 being the reward for the home, neither or away team scoring. Episodes were defined as starting at the beginning of the game or immediately after a goal and finished with a goal or the end of the game. To estimate the action values, the authors used a long-short term memory neural network. The Sarsa on-policy temporal difference prediction method was used to train the neural network weights. A goal impact metric was calculated for players and correlated well with success measures currently used in ice hockey (e.g. goals and assists). The authors also showed how the model could be correlated with current salaries to identify players who produced the best (or worst) value for the money they are paid. The study provides an informative model, but cannot be applied to rugby league due to the data availability and Markov Game application issues described above.

Liu et al.(2020) adopted a more advanced model free deep reinforcement learning approach with their work in football. Their model learnt from 4.6 million actions across ten European leagues from the 2017/18 season. The data was preprocessed to provide a similar set of features toLiu and Schulte(2018), however continuous features were used rather than zones for location information, to allow the model to generalise findings to areas of the pitch not visited in the training process. The states were again composed of this set of features and an action history. The reward used was also the same as previously (+1 for home goal, -1 for away goal, 0 for no goal). Rather than solving the model as a single agent, the authors employed a two-tower design (one for the home team and one for the away team) to generate the action values. This allowed them to fit the data of the home and away teams separately. Each tower used a long short-term memory neural network, whose weights were trained via the Sarsa temporal difference prediction method. Figure2.13provides a plot of different action values across the pitch, some of which may not have been visited in the training process. The authors created a goal impact rating, similar to their work in ice hockey (Liu and Schulte,2018), to rank player performances and showed that it correlated well with future performances at a global level but required fine tuning when attempting to generalise to a single league (e.g. performances at the

2.2 Team and Player Performance Analysis in Sport

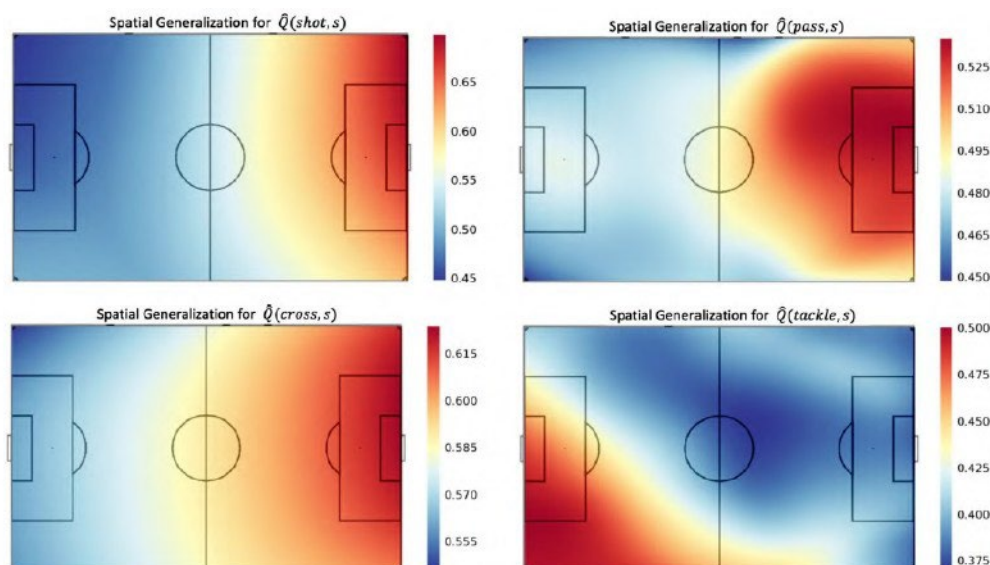


Figure 2.13: Liu et al.(2020) action values for four different actions (shot, pass, cross and tackle) across the pitch.

Championship level were poorly correlated with the global model). The model provides an extensive set of useful results and the two-tower design employed in the training process ensures that it could be possible to apply the methods to a rugby league environment. Unfortunately data availability issues within rugby league ensure that this is not currently possible.

Singh(n.d.) utilised Markovian principles (Howard,1971) to provide EPV values for a set of 192 zones across a football pitch, which he called "expected threat" (xT). He used match event data from the 2017/18 English Premier League season. The data were preprocessed so that only the zone for the action and its descriptor from the set of "move" (i.e. run with ball or pass) or "shoot" were considered. The only reward utilised was a goal being scored, valued as 1. Dynamic programming (Bellman,1954) was used to iteratively evaluate the xT for each zone, starting with a value of 0 for all zones. The author provided a heat map style interactive plot of the completed model (Figure2.14) and showed how the values could be used to identify threatening players by taking the difference in the team's expected threat before and after the player performed an action. Similarly, he showed how it could be used to evaluate the locations through which teams generated value and whether this occurred via passing or shooting in an interactive plot for the purposes of

2.2 Team and Player Performance Analysis in Sport




Figure 2.14: Singh(n.d.) xT plot

future tactical preparations. Although the xT model provides interesting insights into player and team performances, it is not clear whether the zones used were empirically or arbitrarily chosen. Similarly, the dynamic programming approach requires a complete specification of the transition matrix, which may not be computationally feasible if more than two actions are evaluated. However, the study provides a baseline approach, which could be adapted to the needs of rugby league.

Merhej et al.(2021) extended the work of Singh(n.d.) by using deep learning to evaluate defensive actions using only event level football data. They used 760 matches from 2 seasons of English Premier League football data (2017-2019). First, they preprocessed their data and produced an xT model specific to their data, following the previously published methodology (Singh,n.d.). Next, the authors produced a simple feature dataset of two previous actions, providing their x and y co-ordinates and their xT value. A multi-layer perceptron model was used to predict the xT, which an action would have achieved if a tackle or interception had not been made by the opposition. The authors showed how these predicted xTs could be used to value individual players' defensive performances by summing them and normalising to a 0-100 score. The study provides a practical method through which defensive actions can be evaluated and only requires event level data. However, it relies on a version of an EPV or xT model, which can be applied to player performances. Such a model is not currently available within rugby league, so it will need to be developed before the concepts of this study can be applied to the sport.

Chan et al.(2021) adopted a Markov Reward Process approach to evaluate team performances in American football. They used play-by-play data from the 2013/14 to

2016/17 National Football League seasons, totalling 164,299 plays. During preprocessing, the authors extracted the yard line and down to form the state space and the next play's down and yard line for the transition matrix estimation. Four rewards were calculated from the perspective of the attacking team: field goal (3 points); safety (-2 points); touchdown (6.97 points); and opposition scoring turnover touchdown (-6.97 points). Dynamic programming was used to estimate the transition matrix and solve the Markov Reward Process. A points gained metric was identified as the difference in value between the beginning and end of the play. It was used to evaluate team performances and different play types. The results showed that better teams performed better in the middle of the field and were able to use short pass plays more effectively than worse teams. Although only using data at a play by play rather than at the action by action level, the methodology employed by [Chan et al.\(2021\)](#) can be directly related to a rugby league environment: rather than four downs, rugby league would use six plays; and rather than the rewards used in this study, rugby league would use the five possession outcomes (no try, converted try, unconverted try, penalty goal and drop goal). Unlike American football though, a play can begin anywhere across the pitch in rugby league so rather than only considering the distance from the opposition try line, rugby league would have to create zones, which also consider how central or wide on the pitch the play begins.

2.3 Rugby League

2.3.1 Overview

Rugby League is a field-based collision sport, where 13 players on each team attempt to score more points than the opposition in a match played across two 40 minute halves. Matches are played on a pitch with dimensions 68m x 120m. The end 10m at both ends of the y-axis are designated as try areas, where the highest value points are scored. Consequently, few actions occur in the opposition try area when a team is attacking because as soon as the attacking team is in control of the ball in this area, they will attempt to ground it for a try. [Figure 2.15](#) provides a representation of the rugby league pitch. The try lines represent the beginning of the try areas, which extend to the end of the pitch. On each try line there is a 6m wide goal. Teams are able to score points if they kick the ball over this goal for either a drop goal, penalty kick or conversion kick. At the beginning of a half of




Figure 2.15: Plot of 70m x 120m rugby league pitch. Assumes the team in possession of the ball is attacking in an upwards direction. Own refers to the possession team; Opp refers to the opposition team.

match play or after a scoring action, play is restarted at the 50m line. The 20m line is an important visual signpost for players as it signifies a point where coaches typically want them to employ specific tactical strategies (e.g. set attacking or defensive plays to try and score or prevent points from being scored).

There are four methods of scoring points in rugby league: a *try*, where the ball is grounded beyond the opposition try line; a *conversion kick*, which is attempted after every try and is successful when the ball is kicked from the floor between the goal posts; a *penalty kick*, which is also successful when the ball is kicked from the floor between the goal posts, but is offered as a result of an opposition foul rather than a try being scored; and a *drop goal*, which is successful when the ball is kicked between the goal posts from the attacking players hands. A converted try (try scored and conversion kick successful) is worth 6 points, an unconverted try (try scored, but conversion kick unsuccessful) is worth 4 points, a penalty goal is worth 2 points and a drop goal is worth 1 point. Table 2.1 outlines the potential possession outcomes in rugby league. If the teams have scored equal points after 80 minutes of match play, two additional periods of 5 minutes is played

Table 2.1: Possession outcomes in rugby league

Action	Points Scored
Converted try	6
Unconverted try	4
Penalty Goal	2
Drop Goal	1
No score	0

to enable a 'golden point' to be scored. A 'golden point' is any point scoring event, which in this additional time wins the match for the team who scores it.

Rugby league is a sport which follows an episodic pattern. The first action of any match is a kick from the halfway line. The successful receipt of this kick begins a phase of attacking possession for the ball catcher's team and a phase of defensive possession for the kicker's team. In an attacking possession, a team attempts to progress the ball as far as possible towards the opposition try line to increase their chances of scoring points. Conversely, in the defensive phase of possession, the defending team attempts to stop the progression of the opposition team towards their try line. Typically, an *attacking possession* consists of a *set* of six plays as per the rules of rugby league. A *play* begins when the ball is taken from the opposition (e.g. by catching one of their kicks) or when the ball is rolled between the attacking player's legs after a tackle has been completed by the opposition (a play-the-ball). If a team is tackled for the sixth time within their set of six plays, the ball is handed to the defending team to begin their attacking set at the point where the tackle took place (a handover). Consequently, a team will normally kick the ball as far down the pitch as possible, or into the try area to try and score points, before the sixth tackle happens. In this situation, the defending team's attacking possession begins at the point where they catch the ball. Throughout this thesis, the term attacking possession is used. An attacking possession can consist of multiple consecutive sets for the same team. Such a situation typically occurs when the defending team fouls the attacking team, resulting in a new set of six plays beginning for the attacking team. An attacking possession could therefore feasibly contain as many as 20, 30 or 40 plays (the equivalent of 3-7 full sets) if the opposition continuously fouls the attacking team before their attacking set ends. Combining the results of all the attacking possessions a team

Table 2.2: Possession terms used in this thesis, ordered from lowest level (i.e. a single action) to highest (i.e. all actions within a specified period of time).

Term	Description
Action	Any action performed on the pitch (e.g. pass, kick, run, foul).
Play	A sequence of actions, which ends when a player on the team is tackled. The number of actions in a play is unlimited but typically does not exceed 20 as a player is usually tackled before then.
Set	A group of up to six plays for the same team. There cannot be more than six plays in the same set, but a team can complete consecutive sets in the same possession if the opposition fouls them.
Attacking Possession	A sequence of actions, plays or sets through which a team is in possession of the ball. Typically this will be the same as a set, but in the case of an opposition team foul, an attacking possession can contain multiple sets.
Attacking Performance	All the attacking possessions a team produces during a specified period of time (e.g. a single match, multiple matches against the same opponent, or across the season).

produces during a match can be considered a team's *attacking performance*. Table 2.2 details the possession terms used in this study and their relationship with each other.

From an attacking team perspective with a finite opportunity to score points per attacking possession, it is logical that teams will attempt to perform as many actions as possible with their best players, or perform as many actions as possible targeting the weaknesses of their opposition. Whilst such tactics provide attacking teams with the best opportunity to score points in matches, they also provide an opportunity for opposition teams to understand the tactical strategies a team may employ before the match based on their previous performances. Such information could be gleaned either from the proportion of actions occurring within a specific area, or *zone*, or the total/proportional value of actions completed within that zone. In practice, performance analysts do not currently obtain this information via data-driven processes. Instead, they watch videos of the opposition team's previous performances prior to a match and evaluate their attacking performances subjec-

tively. This information is then relayed to the head coach who will produce a strategy to try and beat the opposition. This process is both subjective (and open to individual bias, which can have positive or negative results) and time consuming. Therefore, an objective, data-driven method of understanding team attacking performances in rugby league could be beneficial from the perspective of preparing tactical strategies to try and beat future opponents.

Table 2.3 provides an overview of the roles and responsibilities of different players on the pitch (Dalton-Barron et al., 2022). In brief, there are three main elements to rugby league: *gaining metres*, *scoring tries* and *linking play together*. Props and second rows typically attempt to gain metres in their own half of the pitch to set up promising attacking positions for the team; centres, wingers and full backs attempt to score points by performing the last actions of the team's attacking plays and usually perform the first action after an opposition team has kicked the ball to counter attack; and the half pair and hooker link play together by playing the first actions at the beginning of most plays. With different roles and responsibilities, the players are evaluated at an individual level by coaches in different ways. For example, the main role of the hooker is to link play together, as such it wouldn't be useful to rate these players based on the number of tackles they make or the number of tries they score. The ability to evaluate players is pivotal though as rugby league operates under a salary cap, which limits the amount of money teams can spend on player salaries. Therefore, in order to remain competitive, teams must always seek out either the best quality players in their position, or those who are most underrated in terms of their salary cost.

To help coaches evaluate rugby league players, some basic *performance indicators* are provided by data providers (Opta, n.d.). These include number of tries scored, pass statistics (number of try assists and offloads), run statistics (number of carries, metres, tackle busts and clean breaks made, average metres gained per carry), kick statistics (number of attacking kicks and forty-twos made, drop goals and penalty/conversion goals scored, missed goals), defensive statistics (number of tackles made or missed) and error/foul statistics (number of errors made, penalties conceded, red and yellow cards). Although useful, these elements are usually provided as *accumulated counts* across fixtures, opponents or seasons, which provide limited context surrounding the value of the actions completed with regards to how they benefit the team. For example, gaining 10 metres is much more valuable if you are 9m away from the opposition try line as it provides you

Table 2.3: Rugby league player positions and their general responsibilities. Brackets represent number of players in that position on the pitch.

Position	Responsibility
Full back	The last line of defence, these players stand behind the other defenders and must chase down and tackle any attacking player who breaks through the first line of defence. In attack, they provide support runs in an attempt to create 'overloads', which provide better opportunities for tries to be scored.
Wingers	Typically the widest players on the team. In defence, they must catch and return high balls kicked into their try area and defend their opposite winger in open play. They are expected to finish attacking moves by exploiting space generated for them and scoring tries in attack.
Centres	This position is just inside the wingers and holds a very similar role. In defence, they defend their opposite centres and in attack they pose a threat in wide areas, but typically would provide an assist for a winger to score a try, rather than scoring themselves. They are less likely to catch a high ball in defence than wingers or full backs.
Half pair	The scrum-half and the stand-off, first and second-receivers of the ball at the start of a play. They are the playmakers of the team and are expected to direct their team-mates as to the tactics that are going to be used when the team is attacking.
Hooker	Typically the player who gets the ball from a play-the-ball and passes it to the half-pair. They are similarly creative to the half-pair and have the opportunity to run with the ball when they receive it if they think it's a better attacking option than passing.
Props	Typically the biggest and heaviest players on the team, they tend to be used as "battering rams" who run directly at the opposition to gain metres. This tactic is used particularly when the team is close to its own try line. Defensively, they attempt to stop their opposing prop gaining metres when they try the same tactic.
Second Row	Similar to the props, but placed wider on the pitch, these players work in conjunction with their centre and wing to try and make metres. Defensively, they perform a larger number of tackles across a wider range of locations than other players.
Loose Forward	A more free position, which varies between teams. Can play more like a prop (i.e. battering ram) or a hooker (i.e. playmaker) depending on the requirements of the coach.

with the opportunity to ground the ball for a try, than it is if you are stood on your own 20m line, where running 10 metres does not provide the same opportunity to score a try). Recently, action by action *spatial data* (i.e. x, y coordinates for each action) has become available, which could help to provide an understanding of this context. There is therefore scope for research, which can objectively rate player performances based on the actions they take and the locations in which they take them.

2.3.2 Team and Player Performance Analysis in Rugby League

Unlike the sports mentioned in Section 2.2, rugby league is an example of a sport with poor data availability. Indeed, action by action spatial data only became available in the Super League in 2020. As a result of this poor data availability, no study has analysed action by action data in rugby league. Prior to 2020, only the first action of each play was coded with spatial information, which meant that only play by play analyses similar to that of Chan et al.(2021) could be conducted (Holbrook et al.,2019;Kempton et al., 2016). Consequently, the majority of research considering team and player performances in rugby league is limited to the use of accumulated count data as performance indicators (Parmar et al.,2018;Wedding et al.,2020;Woods et al.,2017). This section highlights the elements of both types of analysis (i.e. using accumulated counts as performance indicators and play by play data), which aim to enhance the coaches' ability to evaluate player and team performances in rugby league.

2.3.2.1 Accumulated Count Data Analysis in Rugby League

Woods et al.(2017) analysed the relationship between 14 author selected performance indicator variables and the win/lose dichotomous match outcome and team league position at the end of the season. The authors used 376 observations from the 2016 National Rugby League season, from which 2 drawn matches were removed. The performance indicators were accumulated at the fixture level (i.e. counts of the performance indicators per match) for the match outcome analysis and the season level (i.e. counts of the performance indicators per season) for the league position analysis. A conditional inference classification tree (Hothorn et al.,2006) was used to establish the extent to which the 14 selected team performance indicators could be used to explain individual match outcomes. Ordinal regression, via cumulative link mixed models, was used to estimate

the relationship between performance indicators and the team's final league position. Five team performance indicators were used to predict match outcome by the conditional inference classification tree: try assists, all run metres, line breaks, dummy half runs and offloads. The model was able to accurately classify 91% of wins and 66% of losses in the data set. The authors showed that missed tackles were significantly related to a worse final league position by the ordinal regression, whereas kick metres and dummy half runs were associated with an improved final league position. Coaches could use the results as performance indicators to measure performance or influence tactical strategies.

[Parmar et al.\(2018\)](#) extended the previous work in rugby league ([Woods et al.,2017](#)) by utilising a more data driven approach within the selection of performance indicators. Their model attempted to predict match outcome using performance indicator counts accumulated at the fixture level. Their data was taken from the 2012-2014 Super League seasons and used a total of 567 matches. 45 action variables were used as performance indicator variables by subtracting the away team's performance from the home team's performance. The authors used principal component analysis ([Jolliffe,2002](#)) to reduce 45 performance indicators to 10 linear combinations. Logistic regression was used with these 10 performance indicators against a dichotomous match outcome (win/lose) dependent variable. Exhaustive chi-square automatic interaction decision trees ([Kass,1980](#)) were used to evaluate the importance of the principal components with respect to the outcome variable. The model accurately predicted match outcome with 88.3% accuracy in the training data set and 90.5% accuracy in the testing data set. The three most important performance indicators were identified as "making quick ground" (i.e. progressing quickly towards the opposition try line), "amount of possession" (i.e. how many actions the team performs with the ball) and "form" (i.e. whether the team won or lost a previous set of matches). By utilising a principal component analysis approach, the authors were able to group together different performance indicators and provide two key modifiable elements ("making quick ground" and "amount of possession"), which coaches could use to measure performance. The inclusion of "form" provided little useful information from a practical perspective.

[Wedding et al.\(2020\)](#) followed a similar data driven approach to understand seasonal changes in performance indicator variables for different playing positions. They used 34,047 observations from the 2015-2019 seasons. 48 performance indicators were identified at a player level and normalised to playing time prior to analysis. The authors used

principal component analysis to reduce the 48 performance indicators to 14 linear combinations. The 14 performance indicators accounted for 58.5% of the variance in the data. Two-step cluster analysis (Chiu et al.,2001) was used to group players into positional groups based on their performance indicators. Six positional groups were identified by the authors' data driven approach, compared to four groups highlighted by previous research (Sirotic et al.,2011). Discriminant analysis showed that players were correctly classified into the six positional groups with 62.9% accuracy using the 14 performance indicators. The authors were able to identify the combinations of performance indicators most relevant to each positional group, providing a set of performance indicators that coaches could use to measure performances at the accumulated count data level.

2.3.2.2 Play by Play Data Analysis in Rugby League

As mentioned in Section2.2.2, although accumulated count data can provide performance indicators which coaches can use to evaluate performances or change tactics, the studies provide few direct insights into the sport on an action by action basis. This is because accumulated counts disregard the context surrounding each individual action. However, whereas in other sports action by action and tracking data have become readily available and high quality analyses have been conducted (Fernández et al.,2021;Liu et al.,2020), at the onset of this PhD this data was not yet available, so these analyses could not yet be conducted in rugby league. Indeed, only play by play data (i.e. the location of the first action of each play) had been considered in the literature (Holbrook et al.,2019;Kempton et al.,2016).

Kempton et al.(2016) produced the first variation of an EPV model in rugby league based on the location of the play-the-ball (the first action of each play). The authors used 155,352 observations from 768 matches in the 2010-2013 seasons of the National Rugby League. The x, y coordinates of each play-the-ball were noted and assigned to one of seventy 10m x 10m zones on the pitch. The reward was calculated as the number of points scored by the next try scoring action plus an approximate value for the 2 point conversion kick based on the probability of it being scored. Dynamic programming was used to estimate the expected long-run value of the difference in points scored by the team in possession and its opposition. The model was run five times, once for each play within an attacking set. Figure2.16provides an EPV plot of the pitch. The results showed that




Figure 2.16: Kempton et al.(2016) EPV plot.

the closer a team was to the opposition try line irrespective of play number, the more points they were likely to score. Similarly, the closer the team was to the opposition try line at the start of the attacking set, the more likely they were to score points. The study provides a good introduction to the use of Markovian principles within the sport of rugby league but provided very few practical applications. Furthermore, the zone sizes appeared to be arbitrarily chosen. As such, there is scope for an empirical study which considers different zone sizes and attempts to provide more practical applications, similar to the insights provided in other sports (Chan et al.,2021;Singh,n.d.).

Holbrook et al.(2019) also used play-by-play data in rugby league to concurrently predict five possession/match outcomes. The authors used more than 250,000 play-the-ball actions from over 750 games in the 2015-2018 seasons of the National Rugby League. Rather than preprocessing the data in a specific manner, the authors used embedding layers to obtain high level feature information for specific elements of the model. The inputs to the model therefore included the raw values for score difference, time remaining and x,y co-ordinates of the play-the-ball. Also included were individual embedded features for the x,y position co-ordinates, tackle number with a back-to-back set dummy variable, one hot encoded team ID with season ID, and one hot encoded opponent ID with season ID. The authors used a mixture density neural network (Bishop,1994) to predict five outcomes on a play-by-play basis. These outcomes were: expected metres gained in the

next play; expected try during the next play; expected try in this set; win probability for the match; scoreline prediction for the match. The authors showed how the results could be used to evaluate team performances both in real-time and post-match. They did this by considering the expected metres gained and probability of scoring a try within a set in real-time. Offensive and defensive value over average metrics were also produced, which highlighted a team's ability to attack or defend better than expected measures across the season. These metrics were shown to be well related to the final league position of each team. The results of the study were mainly developed for entertainment purposes (i.e. use on television), but coaches could use them as performance indicators similar to the accumulated count data described above.

2.4 Summary

The aim of this PhD is to develop new methodologies to evaluate player and team performances in rugby league using match event level data. Within the literature, two broad areas of research were identified in this area: accumulated count data analyses and action by action (or play by play) data analyses. The key limitation of accumulated count data analyses, highlighted at the beginning of Sections 2.2.2 and 2.3.2.2, is that they do not provide contextual data surrounding the actions completed. For example, a 10m run in rugby league is much more valuable 9m from the opposition try line than 20m from the opposition try line as running from 9m provides the possibility for the ball to be grounded for a try, but running from 20m does not. Similarly, a pass to a player with no opponents in front of him is more valuable than a pass to a player who has three defenders directly in front of him, as the player receiving the ball has greater opportunity to progress towards the opposition try line with no opponents in front of him than he does with three opposition players to get through or around. Accumulated count data analyses are unable to account for all of this information, which reduces their validity with respect to evaluating player and team performances. Consequently, accumulated count data analyses will not be considered within this thesis.

Prior to the onset of this PhD, two studies considered play by play data analyses in rugby league (Holbrook et al., 2019; Kempton et al., 2016). These studies used x, y coordinate data from the first action of each play (the play-the-ball) to provide some spatial context within their models. Kempton et al. (2016) produced the first EPV model

in rugby league, which provided insights into the value of different locations on the pitch. They used dynamic programming with fixed zone sizes, which allowed greater insights than previously. However, they framed their analysis at a global level (i.e. what happens in the league, rather than what specific teams/players do) on a play by play level, limiting the usefulness of the results in practice. Furthermore, the zone sizes were arbitrarily chosen, rather than through data-driven analysis or expert insight (Cervone et al.,2016), and it is unclear how this affected the results obtained. The work provides a valuable basis from which future research within the sport could be conducted though, as they show how the episodic nature of rugby league can be harnessed within an MDP framework to gain interesting results. (Holbrook et al.,2019) adopted a deep learning approach to understand team performances in rugby league. They showed the expected metres gained in a set on a play by play basis, alongside the expected try probability. Furthermore, they extended the results to provide metrics evaluating teams' attacking and defensive performances. The work provides useful information to practitioners in the National Rugby League at a play by play level, but it does not consider individual action data or the actions taken by each team to achieve the success they did. As such, the insights it provides with respect to the development of tactical strategies is limited.

Table 2.4 provides a summary of the action by action data analyses currently available across all sports. The literature can be split into two main strands: spatial trends analyses and EPV models. Although EPV models have gained considerably more attention within the literature, there is scope for the two models to be used together to understand player and team performances, as shown by Fernández et al.(2021) who used their pitch control model within their comprehensive deep learning approach to calculating EPV in football. The following subsections analyse both strands of literature (spatial trends analyses and EPV models), evaluating their suitability for adaptation and validation within rugby league.

2.4.1 Spatial Trends Models

Recall from Definition 2.2.1 that spatial trends analyses attempt to estimate how likely is it that a team will control the ball in a given location, or where is a player or team likely to pass the ball. Three studies have considered these analyses in sport with respect to player and team performances (Fernández and Bornn,2018;Mallepalle et al.,2020;

Table 2.4: Summary of studies considering variations of action by action or play by play data analyses across all sports

Author	Data	Method	Outcome
Spearman et al. (2017)	38 English Premier League matches from the 2015/16 season, event and tracking data from 10,875 passes	Bayesian estimation of physics-based approach	Produced a pitch control model estimating the probability of a successful pass for every location on pitch
Fernández and Bornn (2018)	20 Spanish First and Second Division matches, equalling 2.4 million frames of tracking data only	Manipulated the bivariate normal distribution using speed and direction of player movements	Produced a pitch control model based on the probabilistic pitch influence of each player
Mallepalle et al. (2020)	27,171 quarterback passes from 2017-2018 NFL seasons	Kernel density estimation to produce smooth passing probability surface	Produced a model which estimated the probability of quarterbacks passing to different locations across the pitch
Cervone et al. (2016)	2013/14 NBA season action and tracking data ("over 1 billion space-time observations")	Bayesian approach following Markovian principles	Produced EPV model capable of rating player performances

Fernández et al. (2021)	633 English Premier League matches between 2013-2015, action and tracking data	Decomposed MDP deep learning approach to learn individual models for each action	Produced EPV model capable of rating player and team performances
Decroos et al. (2019)	11,565 European football matches between 2012-2018 (14.4 million actions)	CatBoost algorithm separately predicting goal scored or conceded within next 10 actions	Produced VAEP model capable of rating player performances
Routley and Schulte (2015)	2.8 million actions from the 2006/07 National Hockey League season	Markov Game model to evaluate action values	Produced an action impact rating to evaluate player performances in ice hockey
Liu and Schulte (2018)	3.3 million actions from the 2015/16 National Hockey League season	Deep reinforcement learning algorithm to evaluate action values	Produced a goal impact rating to evaluate player performances in ice hockey
Liu et al. (2020)	4.6 million actions across ten European football leagues in 2017/18	Two tower deep reinforcement learning algorithm to estimate action values for each team individually	Produced a goal impact rating to evaluate player performances in football
Singh (n.d.)	Matches from 2017/18 English Premier League	Dynamic programming to model transition matrix	Produced xT model, able to value player and team performances

Merhej et al. (2021)	760 English Premier League matches between 2017-2019	Multi-layer perceptron model to predict xT, which would have been achieved without a defensive action taking place	Produced a defensive action xT model which values defensive actions
Chan et al. (2021)	Play by play data from 2013/14-2016/17 National Football League seasons (164,299 plays)	Dynamic programming to model transition matrix of an MRP	Produced points gained metric able to value different types of plays and differentiate between teams' performances
Kempton et al. (2016)	155,352 play-the-ball observations from the National Rugby League between 2010-2013	Dynamic programming used to estimate the expected long-run value of the difference in points scored by the team in possession and its opposition	Provided an understanding of valuable play-the-ball locations for a team in possession
Holbrook et al. (2019)	More than 250,000 play-the-ball observations from the National Rugby League between 2015-2018	Mixture density neural network to predict five outcomes, including expected metres and try for the set and match score prediction	Produced a defence-adjusted value over average metric to evaluate attacking and defensive team performances

Spearman et al. (2017)). Although the methods used by **Fernández and Bornn (2018)** and **Spearman et al. (2017)** are extremely valuable within football, their value in rugby league is more limited due to the laws of the game, which dictate that all thrown passes must

be directed backwards. There would be great value in adapting the methods for use with kicking actions, but these actions occur 10 times less frequently than passes so the need for a model of this type in rugby league is questionable. Regardless, tracking data is not available in the sport so it is not yet possible to adapt and validate the methods previously used in football (Fernández and Bornn,2018;Spearman et al.,2017) in rugby league.

Unlike the football studies (Fernández and Bornn,2018;Spearman et al.,2017), the work ofMallepalle et al.(2020) in American football could be adapted and validated in rugby league with some small changes.Mallepalle et al.(2020) used KDEs to calculate quarterback pass probability distributions from 27,121 quarterback passes. In rugby league, adopting this analysis at a player level makes limited sense as thrown passes must travel backwards and kicks don't happen frequently enough for tactical value to be gleaned from the analysis. However, adopting a similar KDE approach at a team level could be extremely valuable as it may provide an insight into areas on the pitch where a team is most likely to control the ball and thus an indication of the spatial trends of their attacking performances. Furthermore, it may be possible to compare these distributions within teams against different opponents or between teams to identify teams who perform in a similar or dissimilar manner.Mallepalle et al.(2020) did not directly quantify the differences between players in their study so there is scope for the development of a novel metric, which is able to quantify the differences in spatial trends of attacking performances between and within teams. Coaches may be able to use this understanding to develop tactical strategies for future matches.

2.4.2 Expected Possession Value Models

Recall from Definition2.2.2that EPV models value player and team performances based on the probability of points/goals being scored conditional on the location of the action. The literature surrounding EPV models, whilst still developing, is much more widespread than the research considering the spatial trends of attacking performances. A wide variety of analytical methods have been adopted, including deep learning (Fernández et al.,2021;Merhej et al.,2021), machine learning (Decroos et al.,2019), Markovian approaches (Chan et al.,2021;Liu et al.,2020;Liu and Schulte,2018;Routley and Schulte,2015;Singh,n.d.) and Bayesian analysis (Cervone et al.,2016). However, a central tenet of most EPV models is the adoption of Markovian principles to some extent regardless of

the specific analytical methods employed. Such usage in other sports with more advanced analyses than rugby league is positive given [Kempton et al.\(2016\)](#) have already shown that a Markovian approach can be used in rugby league to generate interesting findings. Despite this, there are significant limitations relating to the application, adaptation and validation of EPV models from other sports in rugby league.

The models of [Fernández et al.\(2021\)](#) and [Cervone et al.\(2016\)](#) provide gold standard approaches from a deep learning and statistical analysis perspective. However, both employ tracking data, which is not currently available in rugby league and so they cannot currently be applied to the sport. Similarly, although the work of [Decroos et al.\(2019\)](#), [Liu and Schulte\(2018\)](#) and [Liu et al.\(2020\)](#) is valuable in the space of analysing event level only data, the amount of data required to run these models (11,565 matches ([Decroos et al.,2019](#)), 3.3 million actions ([Liu and Schulte,2018](#)), 4.6 million actions ([Liu et al., 2020](#))) is not available in rugby league. From a pure data availability perspective, the two most replicable models are [Singh\(n.d.\)](#) who used a Markov Decision Process approach for his xT model in football and [Chan et al.\(2021\)](#), who adopted a Markov Reward Process approach for their model in American football, similar to that used by [Kempton et al.\(2016\)](#) in rugby league. Whereas [Chan et al.\(2021\)](#) did not need to select two dimensional zones due to the rules of American football, [Singh\(n.d.\)](#) did and used arbitrary zone sizes in his xT model, similar to [Kempton et al.\(2016\)](#). A much more appropriate method of zone size selection would be using expert opinion ([Cervone et al.,2016](#)) or utilising a data driven approach, which has never previously been attempted. There is therefore scope for a sequence of studies within rugby league, which apply, adapt and validate the previous literature ([Kempton et al.,2016](#); [Singh,n.d.](#)) by first empirically determining the most suitable zone sizes for use within rugby league using a data driven approach, building upon the Markov Reward Process approach previously employed ([Chan et al., 2021](#); [Kempton et al.,2016](#)), and second, extend the Markov Reward Process approach to include actions, adopting a Markov Decision Process approach similar to [Singh\(n.d.\)](#).

Alongside the issue of data availability, some of the EPV models produced in other sports cannot be applied to rugby league due to its unique rules and playing style. For example, the episodic nature of rugby league, whereby a tackle by the opposition team does not result in a change in possession precludes the adaptation of the Markov Game approach suggested in ice hockey ([Liu and Schulte,2018](#); [Routley and Schulte,2015](#)). The two tower deep reinforcement learning approach suggested by [Liu et al.\(2020\)](#) as an

extension to these methods in football could be used in rugby league if a suitable quantity of data was available though. A further consideration is the point scoring structure of rugby league, which is much more complicated than football or ice hockey. In rugby league, there are five possible possession outcomes (Table 2.1), which limits the applicability of some of the models previously identified within the literature. For example, the individual probability estimation of scoring and conceding outcomes suggested by Decroos et al. (2019) would result in an impractical eight models being calculated. Conversely, aggregating the point scoring methods together as has been done in all approaches to date, where multiple scoring options are available (Cervone et al., 2016; Chan et al., 2021; Kempton et al., 2016) could result in the loss of valuable data with respect to opposition strategies. For example, if a team is more likely to score or concede a penalty goal or a converted try from a specific area of the pitch, it would be useful for a coach to be able to use that information to devise attacking or defending strategies for upcoming matches. Given the data available within rugby league, it could be difficult to identify an appropriate machine learning method to estimate the probability of individual possession outcomes, so a Bayesian approach, which provides parameter estimates conditional on prior knowledge and the data available may be more appropriate. A Bayesian approach has previously been used in the EPV literature (Cervone et al., 2016), albeit not to estimate individual possession outcome probabilities, and has been used to show uncertainty in player ratings based on the data available for them within accumulated count data analysis (Whitaker et al., 2021). This novel approach to the estimation of individual possession outcome probabilities could therefore provide the best analytical solution to the challenges imposed by low data availability sports. The flexibility of such a model could also provide significantly more scope to produce novel metrics evaluating player and team performances.

2.5 Methodological Overview

In Section 2.4, the difficulties associated with adapting some methods previously used in the literature to the unique characteristics of rugby league were highlighted. In this section, a high level overview of the methods which will be used in this thesis is provided. Four key methods will be considered: Markov models from the existing literature (Kempton et al., 2016; Routley and Schulte, 2015; Singh, n.d.) will be applied, adapted and

extended to rugby league in Chapters3and4; Kernel Density Estimation will be used to evaluate the spatial trends of attacking performances, adapting previous research in American football (Mallepalle et al.,2020) in Chapter5; Bayesian analysis will be used to produce a completely novel EPV model capable of estimating individual possession outcome probabilities in Chapter6; and distributional comparisons will be used to provide single value differences between pitch surfaces for the first time, via the Kullback-Leibler (KL) Divergence (Chapter3and Wasserstein Distance (Chapter5).

A Markov model is a stochastic process, which at a minimum, consists of a finite set of states. These states represent the models' understanding of the world and assume the Markov property ensuring that transition to future states is dependent only on the current state. A Markov Reward Process (MRP), as utilised in Chapter3, is a tuple of states, transition probability matrix, reward function and discount factor. It is used to calculate a value function for each state, by back propagating the rewards through the chain of states subject to the discount factor. A Markov Decision Process (MDP), as used in Chapter4, extends an MRP by adding an action to the framework. Within an MDP, the agent chooses an action in each state, which results in a transition to the next state. Consequently, MDPs calculate a value function for each state, action pair.

Kernel Density Estimation (KDE) is a method of probability density estimation, which applies kernel smoothing. Kernel smoothing allows a probability density to be estimated as a weighted function of its neighbouring values. The two most important elements of KDE are the kernel chosen (e.g. uniform, triangular, normal or cosine) and the bandwidth. The bandwidth parameter influences the smoothness of the KDE model and controls “overfitting” by establishing the amount of data smoothed at each point.

Bayesian analysis is a statistical approach to data analysis centred around Bayes Theorem. There are three main elements: the prior distribution; the likelihood function; and the posterior distribution. The prior distribution defines the current knowledge surrounding a model parameter. The likelihood function defines the method through which the data influences the posterior distribution. The posterior distribution balances the prior knowledge with the observed data to identify the most plausible parameter values given the evidence available. This evidence based approach to estimating parameter values allows complex model parameters to be estimated using significantly less data than the machine learning and deep learning approaches previously considered in sport (Decroos et al.,2019;Fernández et al.,2021;Liu et al.,2020).

Comparisons between distributions allow the differences in pitch surfaces between teams or matches to be evaluated using a single measure. In Chapter 3, the KL Divergence will be used. The KL Divergence is a measure used to understand the similarity between two discrete distributions. It is a type of statistical distance rather than a metric and compares an approximating distribution to a true distribution. KL Divergence values range from 0, indicating the two distributions are identical, to infinity, indicating no relationship between the two distributions. The Wasserstein distance, used in Chapter 5, is a distance function defined between probability distributions on a given metric space. It is closely related to optimal transport planning. The optimal transport plan describes the movement of mass between distributions with minimal cost, subject to a cost function. The Wasserstein distance is the total cost of all mass movements described by this transport plan subject to the cost function.

Markov Reward Processes for the Quantification of Team Attacking Performances

3.1 Introduction

This first study of this thesis proposes a model evaluating EPV in rugby league. The study applies and adapts the work of [Kempton et al.\(2016\)](#) in rugby league and [Singh \(n.d.\)](#) in football. Their research is adapted in this study by the implementation of a data driven approach to empirically evaluate different zone sizes, rather than their arbitrary selection ([Kempton et al.,2016](#);[Singh,n.d.](#)). This data driven approach compares different zone sizes' ability to replicate future team attacking performances in rugby league. A novel method through which team attacking performances can be quantified in rugby league is also proposed. The study was published in PLOS One (Thomas Sawczuk, Anna Palczewska, Ben Jones. Development of an expected possession value model to analyse team attacking performances in rugby league. *PLOS One*, 16(11): e0259536, 2021) and presented to the England Rugby League Performance Unit. The feedback provided by the Performance Unit provided validation for the model.

Table 3.1: Details of the four nodes provided by Opta in each rugby league fixture xml file

Node	Details
Match data	Details all the actions which took place in the match
Team data	Identifies the players who played in the match
Fixture data	Stores information regarding the fixture itself (e.g. scores, kick-off time, referee)
Play-the-ball data	Provides information for timings of every play-the-ball in the match

3.2 Methodology

In this section, the methodology for this study is presented. It describes the data used and the preprocessing steps required to prepare the data for analysis. The Markov Reward Process framework used in this study is defined and a method of calculating expected possession values is proposed. The configuration of different zone sizes and their comparison through the Kullback-Leibler Divergence is described. A novel method of comparing team performances through z-score analysis is outlined.

3.2.1 Data

In this study, event level match-play data were obtained from Opta (Stats Perform, London, UK) for the 2019 Super League season. The event level data provided by Opta was produced via human annotation of the actions taking place on the field of play during matches. In the 2019 season, 200 unique actions were coded by Opta. These actions included player actions (e.g. passes, kicks, runs, tackles) and miscellaneous actions (e.g. video referee try reviews). A full list of actions and their definitions is provided in Appendix B. The data for each match was provided by Opta in separate xml files. Each xml file had four nodes: match data; team data; fixture data; and play-the-ball data. Table 3.1 describes each of these nodes. Raw information was provided by Opta using unique identifiers rather than 'real' values.

The match data node contained 28 variables describing each action taking place on the pitch. Due to the coding processes used by Opta, 10 of the variables contained

Table 3.2: Important match data variable definitions from the 2019 Opta dataset.

Column	Definition
ID	Unique identifier for each action row
FXID	Unique identifier for the fixture
PLID	Unique identifier for the player
team_id	Unique identifier for the player's team
MatchTime	Time the action took place in seconds
x_coord	x coordinate of the action
y_coord	y coordinate of the action
action	Unique identifier for the group of action taking place
ActionType	Unique identifier for the action taking place
Actionresult	Unique identifier for the result of the action
Metres	Metres gained by the action
PlayNum	The play number within the attacking set
SetNum	The attacking set number within the match

missing data or no data at all; these variables were therefore removed. Variables removed for this reason included: psID; qualifier3, qualifier4; qualifier5; sequence id; player_advantage; score advantage; flag; advantage; and assoc player. Five further variables were removed because they didn't provide information relevant to the objectives of this study: ps timestamp and ps endstamp, which provided video-time information for when the match event occurred; period, which identified the period of the match; and x_coord.end and y_coord end, which identified the end locations of actions so the Metres variable could be calculated. Consequently, only 13 variables within the match data node were deemed to be useful for this analysis. Table 3.2 provides definitions of these variables.

Team data provided 8 details to identify the players on the pitch, including their player ID, team ID, player name, team name, shirt number, position on the team sheet and minutes played. Fixture data outlined 13 details about the match including the time and date it started, the match day week, IDs for the home and away teams, their full team names and scores at full time and information regarding the referee. Play-the-ball data provided technical information surrounding the contact time, completion of the tackle and the foot contact of each play-the-ball.

In the 2019 Super League season, 180 matches were completed across 12 Super

League clubs. A total of 372,173 match events were recorded across 13,574 sets (median 75.5 sets per match, interquartile range 72-79) and 59,233 plays (median 5 plays per set, interquartile range 3-6). Across the season, 1,369 tries were scored (1,013 successful conversion kicks, 356 unsuccessful conversion kicks), 271 penalty goals were attempted (239 successful, 32 unsuccessful) and 89 drop goals were attempted (42 successful, 47 unsuccessful).

3.2.2 Data Preprocessing

Figure 3.1 provides an overview of the data preprocessing completed within this study. The preprocessing steps converted the 180 raw xml files described above to a single file of 59,233 rows and 7 columns ready for analysis. Full details of each preprocessing step are provided below.

The match action data, provided in 180 xml files, were collected and integrated into a single dataset of 372,173 rows and 28 columns. To achieve this, the match data node variables were extracted from each xml file into separate tables and then concatenated together. The team data node variables were also extracted into tables for each fixture. These were used to map unique team identifiers to the real team name. Similarly, Opta's action definitions and unique identifiers (Appendix B) were used to obtain the real name of the coded actions.

Due to the coding protocol used by Opta in 2019, which aimed to calculate accumulated counts rather than promote action by action analysis, not all actions were provided with accurate coordinates. The first action of each play was always accurately coded, but all other actions' location in the play were coded according to the first action. As such, accurate location data could only be extracted on a play by play, rather than action by action basis. This is similar to the studies conducted by Kempton et al.(2016) and Holbrook et al.(2019). Furthermore, the actions were not always coded in order - sometimes a defensive action was coded before an attacking action had taken place within a play - so all defensive (e.g. tackles) and auxiliary (e.g. yellow card, video referee review) were removed from the dataset before any play by play information was extracted.

The dataset was filtered so that only the first action of each play (as coded by Opta) was included. In a rugby league match, the first action of a play is typically a catch of the ball from an opposition kick or pass interception, a receipt of the ball from a play-the-

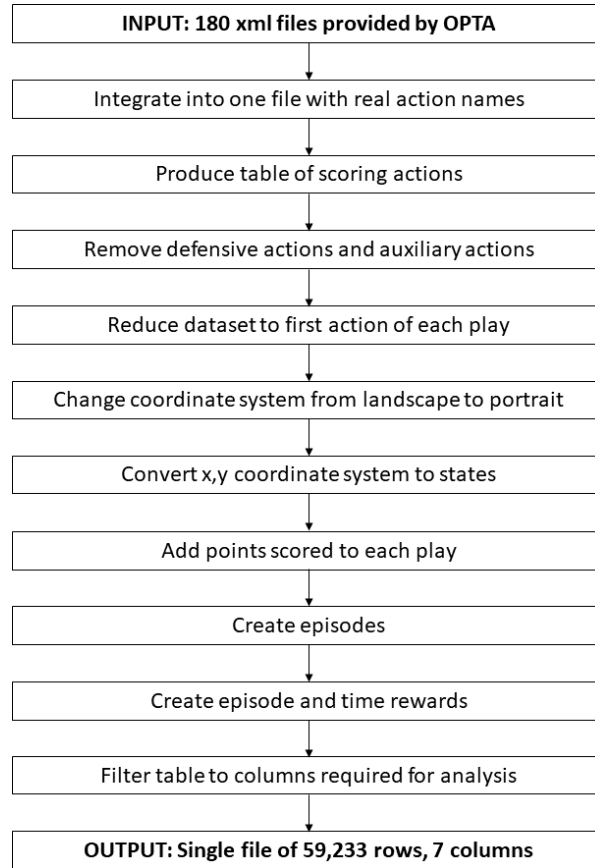


Figure 3.1: Data preprocessing workflow

ball in a handover or a kick (either for the touchline or towards the goal) after a penalty has been awarded. Opta did not code play-the-ball receptions in 2019, so the first action usually referred to the pass or run taken by a player after the play-the-ball reception. Figure 3.2 plots the frequency with which 14 unique actions began a play in this study. All 14 actions occurred on at least 100 occasions, but a completed pass was clearly the most frequent action beginning a play. A further 29 actions began a play on less than 100 occasions: Advantage; Other Carry; Normal; Tap Down; Forward; Play Carry; Off Target; Pick And Go; Chip; Bad Pass; Scraps; Charge Down; Bomb; Fast; Initial Break; Lost Ball Forced; Kicking Offence; Penalty Goal; Cross Pitch; Attacking Catch; In Goal Catch; In Goal Touchdown; Drop Goal; Carried Dead Ball; Key; Forward Pass; In Goal




Figure 3.2: Counts for the actions beginning a play according to the Opta coding system in 2019. All actions began a play on at least 100 occasions. Complete passes are clearly the most frequent action.

Tap Out; Support Carry; Carried In Touch. All actions could be broadly grouped into pass, catch, run, kick or foul categories.

In 2019, Opta provided data in continuous x, y coordinates for a 120m x 68m horizontal pitch. Previously, [Kempton et al.\(2016\)](#) binned these coordinates into 10m x 10m zones. To begin the analysis at a finer granularity in this study, the x, y coordinates for each play were initially binned into $\sim 5\text{m} \times 5\text{m}$ zones, resulting in 308 zones (14 along the x -axis, 24 along the y -axis). For ease of calculation and due to Opta only providing data to the nearest metre, the outer two columns were 4m wide rather than 5m.

All plays were assigned the value of the true points scored within the play: 6 points for converted try; 4 points for unconverted try; 2 points for penalty goal; 1 point for drop goal. If no points were scored within the play, a value of 0 was given.

Table 3.3: List of events, which could end an attacking possession

Event	Description
Handover	A completed sixth tackle by the opposition team, resulting in a change in possession
Kick at goal	Conversion, Penalty Goal or Drop Goal attempt
Foul	Any foul resulting in the opposition team receiving the ball (e.g. conceding a penalty)
Misplaced pass	A pass or tap down, which is intercepted by the opposition
Misplaced kick	Any kick not caught by the team in possession, including bombs/grubber kicks and positional kicks
Handling error	Any situation where the ball is lost from the player's possession (e.g. lost in contact, dropped catch)

The data were split into attacking possessions. An attacking possession was coded as a sequence of plays, which began when a team gained possession of the ball and ended when the team lost possession of the ball (i.e. due to an error, handover, field kick, penalty, drop goal or try). Any situation where a pass/kick missed its target player, but was not successfully collected by the opposition, resulted in the continuation of the attacking possession for the attacking team. Table 3.3 provides a list of the events, which could end an attacking possession and defines them. Using these definitions, the 59,233 plays were grouped into 10,156 attacking possessions (median length 4 plays per attacking possession, range 1-26 plays). The time point of the play within the attacking possession sequence was also identified, starting at 1 to represent the beginning of the possession.

After completing the preprocessing steps, the dataset was filtered to include only the columns required for analysis. As such, the 180 individual xml files were preprocessed into a single dataset of 59,233 rows and 7 columns. Table 3.4 provides an example of two attacking possessions from the first match of the 2019 Super League season. It shows Wigan lose the control of the ball in their second play before St Helens take control of the ball, progress towards the Wigan try line and score a converted try (6 points) in the fifth play of their attacking possession.

A workflow for the analyses completed in this study is shown in Figure 3.3. All elements of analysis are described in the following subsections.

Table 3.4: First two attacking possessions of the preprocessed dataset. FXID is the fixture ID; PosNum refers to the time point of the play within the possession sequence.

FXID	Team	Possession	PosNum	x	y	Points
124011	Wigan	1	1	3	6	0
124011	Wigan	1	2	4	6	0
124011	St Helens	2	1	10	18	0
124011	St Helens	2	2	2	20	0
124011	St Helens	2	3	5	19	0
124011	St Helens	2	4	7	20	0
124011	St Helens	2	5	13	20	6

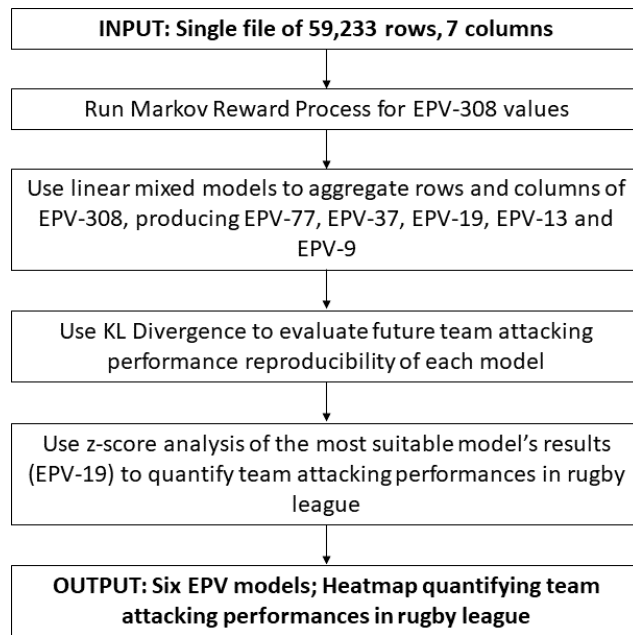


Figure 3.3: Data analysis workflow

3.2.3 Markov Reward Processes

The aim of this study was to apply, adapt and extend the work of [Kempton et al.\(2016\)](#) in rugby league and [Singh\(n.d.\)](#) in football by producing an EPV model with empirically evaluated zone sizes. As with previous literature ([Kempton et al.,2016](#);[Singh,n.d.](#)), a

Markovian approach was employed. A Markov Chain is a stochastic model, describing a sequence of possible events, in which the probability of each event depends only on the state attained in the previous event (Howard,1971). A Markov Reward Process (MRP) is stochastic process, which extends a Markov Chain by adding a reward to each state. It can be defined as follows:

Definition 3.2.1. A *Markov Reward Process* is a tuple (S, P, R, γ) where:

- S is a finite set of states
- P is a transition probability matrix
- R is a reward function
- γ is a discount factor ($\gamma \in [0, 1]$)

A state s provides information about the environment, through which a value function can be learnt (Sutton and Barto,2018). In an MRP, a Markov chain is represented by a finite sequence of states, termed an episode. Within rugby league, attacking possessions can be considered as episodes, whereby the location of an action or play can be considered the state. As a rugby league match takes place within a finite area and in a finite period of time, it is well suited to Markovian modelling. For a state s and successor state s' , the *state transition probability* is defined by:

$$P_{ss'} = P[S_{t+1} = s' | S_t = s]$$

The transition probability matrix P defines transition probabilities from all states s to all successor states s' . It can be estimated explicitly, through dynamic programming (Kemp-ton et al.,2016), or learned implicitly, using real sequences of possessions. The reward R represents the value of performing a specific transition between states. In rugby league, it can be defined as the true points outcome achieved by the play in the team's attacking possession.

In an MRP, the return G_t at time t is defined as the total discounted reward from time t :

$$G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}, \quad (3.1)$$

where k represents the number of time steps in the future and the discount factor $\gamma \in [0, 1]$ provides the present value of future rewards. The value of receiving reward R after $k + 1$ time-steps is $\gamma^k R$.

Subsequently, the value function $V(s)$ for state s can be defined as the expected return when starting from state s (Definition 3.2.2). It provides the long-term value of s , which in rugby league identifies the expected points obtained from state s at the end of the possession.

Definition 3.2.2. State value function $V(s)$ for state s is the expected return when starting from state s :

$$V(s) = \mathbf{E}[G_t | S_t = s], \quad (3.2)$$

There are three main methods of solving MRPs: dynamic programming, Monte Carlo learning or Temporal Difference learning. Dynamic programming (Bellman, 1954) is a model based learning method. It is a collection of algorithms, which can calculate the MRP when the transition matrix is known or can be estimated. Monte Carlo and Temporal Difference learning (Sutton and Barto, 2018) are model free learning methods. They do not require the transition matrix to be estimated as they learn from experience. Of the three methods, Monte Carlo learning is the only algorithm which utilises full episodes when updating its values. Consequently, it is the only algorithm which does not provide biased estimates (Sutton and Barto, 2018). This is advantageous in rugby league as complete samples of attacking possessions from the 2019 Super League season could be used to update state values using the algorithm.

3.2.3.1 Expected Possession Value

Recall the definition of EPV as an instantaneous snapshot of a possession's value (Definition 2.2.2). The value function $V(s)$ for state s provides this instantaneous value for state s based on the expected points obtained at the end of the possession. Consequently, within the MRP framework, the EPV for state s can be considered analogous to the value function for state s :

$$\text{EPV}(s) = V(s)$$

$\text{EPV}(s)$ was simulated using the Monte Carlo every visit algorithm for the reasons described in Section 3.2.3. The Monte Carlo every visit algorithm calculates the empirical

Algorithm 1 Monte Carlo algorithm for EPV calculation

Input : Episode dataset, including states and rewards, γ , empty returns list (Returns), empty counter list (Counter), empty state list (S)

Output: EPV, populated with results of MRP

```

for episode do
  for time-step, t do
     $G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$ 
    Returns(s) = Returns(s) +  $G_t$ 
    Counter(s) = Counter(s) + 1
  end
end
EPV(s) =  $\frac{\text{Returns}(s)}{\text{Counter}(s)}$ 

```

mean of each zone by summing the discounted rewards achieved by the zone and dividing by the total number of visits. The algorithm allows every visit to a zone to be valued, which is important within rugby league because there is no guarantee that plays will move between states as the opposition defence aims to stop progression up the pitch. Algorithm 1 provides pseudocode for the Monte Carlo every visit algorithm used to simulate the EPV for each state. The algorithm loops through each episode and time-step t , calculating the return G_t . The return is added to the returns list for state s and the counter list is incremented by 1. At the end of this loop, the total returns are divided by the counter count for each state to provide a list of EPVs.

In this study, the 5m x 5m zone of the location of the first action of each play was used as the state for the MRP. Episodes were defined as a sequence of these zone locations for each attacking possession performed by a team (Section 3.2.2). Unique episodes were created for every attacking possession. Rewards were calculated as the points outcome of the play: converted try scored (+6); unconverted try scored (+4); penalty goal scored (+2); drop goal scored (+1); loss of possession or missed goal attempt (0). In plays where none of these scoring events occurred, a reward of 0 was assigned. Each time-step within the attacking possession sequence was assigned a reward, so it was possible for a zone to receive multiple rewards in an episode if more than one play began in a given location within the same attacking possession. However, a point scoring reward was only obtained once per attacking possession (i.e. it was treated as a transition to a terminal state).

3.2.4 Configuration of Zone Sizes

EPV models provide the expected points obtained from state s at the end of the possession. In this study, the state refers to an area on the rugby league pitch in which the play begins. Previous studies in rugby league [Kempton et al.\(2016\)](#) and football ([Singh,n.d.](#)) have used arbitrarily selected zone sizes, which may or may not provide the most suitable configuration for use in the sport. This study aimed to adapt and extend previous work ([Kempton et al.,2016](#);[Singh,n.d.](#)) by empirically evaluating the most suitable zone size for use within the sport. In order to achieve this aim, six zone configurations of different granularities were evaluated by the Kullback-Leibler Divergence based on their ability to reproduce future team attacking performances.

Two fixed zone size configurations were used. The EPV-308 contained 5m x 5m zones and used the EPV values obtained from the Monte Carlo every visit algorithm as defined in Section 3.2.3.1. For comparison with the literature, the EPV-77 was produced using the same zones at [Kempton et al.\(2016\)](#) in the field of play, plus a further 7 zones in the team's own try area. The EPVs for the EPV-77 zones were calculated as the weighted average of the four 5m x 5m EPV-308 zones they encompassed.

Four further zone configurations were produced through statistical analysis of the EPV-308 column and row match EPVs. The match EPV, $EPV(s, m)$, obtained by zone s in match m was calculated as

$$EPV(s, m) = \frac{1}{|z \in m|} \sum_{z \in m} EPV(s), \quad (3.3)$$

where $z \in m$ refers to the set of all unique visits to zone s in match m .

The zones' match EPVs were summed at the column and row level to produce 14 column match EPVs for each match and 22 row match EPVs. This process was completed individually for each match. Figure 3.4 depicts the process, showing how the 308 zonal match EPVs were summed to create the row and column match EPVs.

Visual inspection of the initial column and row match EPV's showed that they could be smoothed based on their spatial similarity. Therefore, the fourteen columns were averaged at the 5m level symmetrically to form seven ~ 10m columns (i.e. the two widest columns were averaged, then the two second widest, through to the two most central columns) and the twenty two 5m rows were aggregated sequentially to produce eleven 10m rows similar




Figure 3.4: Depiction of relationship between EPV-308 zones and columns/rows used for statistical analysis. Top figure represents the 5m x 5m EPV-308 zones. Bottom left figure represents the columns; column match EPVs were produced by summing the match EPV for all zones from the EPV-308 in the column. Bottom right figure represents the rows; row match EPVs were produced by summing the match EPV for all zones from the EPV-308 in the row. Pitch lines represent own try area, half way line and opposition try area for team attacking in upwards direction.

to [Kempton et al.\(2016\)](#) (i.e. the bottom two rows were averaged, followed by the next two rows, through to the top two rows).

Linear mixed models were then used to evaluate whether the columns or rows could be aggregated further. In separate models, the column match EPV and row match EPV were added as dependent variables, with team and fixture ID added as random effects. To identify differences in their match EPVs, column and row indexes were added to their respective models as categorical fixed effects. Minimal effects testing ([Murphy and Myors, 1999](#)) was used in four separate models to determine whether two columns or rows could be combined against a smallest effect size of interest (SESOI) of 0.5, 1.0, 1.5 and 2.0

units of match EPV respectively. These values were chosen to provide a range of possible EPV models within the constraints of realistic point scoring differences that could happen (i.e. it is unlikely that a team would produce more than 2 points greater value from any given zone per match). Pairwise least square mean comparisons between consecutive columns or rows were used to evaluate whether they should be aggregated or differentiated. If the difference between two columns or rows was statistically significant (i.e. $P < 0.05$), they remained separate as the difference in match EPV generated by the two factors was greater than the SESOI. Otherwise, the column/row match EPVs were averaged and compared to the next column or row's match EPV. This iterative process was conducted independently for the columns and rows.

At the end of this iterative process, the remaining columns and rows were combined to form a grid. The zones were aggregated at a row level between -10m and 10m, but not at the column level. This decision was made because the zones within the row were visited infrequently relative to other zones on the pitch and so had highly variable values. At SESOI 0.5, 7 rows and columns were present, resulting in 37 zones (EPV-37). At SESOI 1.0, 4 rows and 6 columns were present, resulting in 19 zones (EPV-19). At SESOI 1.5, 4 rows and 4 columns were present, resulting in 13 zones (EPV-13). At SESOI 2.0, 3 rows and 4 columns were present, resulting in 9 zones (EPV-9). The aggregated zone values were calculated as a weighted average of the values of the EPV-308 zones they were composed of. Figure 3.5 highlights this process by depicting the similarities and differences between the EPV-308, EPV-77 and EPV-19 in the 30m closest to the opposition try line. It shows the symmetrical aggregation of columns for EPV-19, compared to the sequential aggregation for EPV-77.

3.2.4.1 Kullback-Leibler Divergence

To empirically evaluate the most suitable zone sizes for use within rugby league, a criterion measure was required. In this study, this was defined from a team performance analysis perspective. In order for a model to be useful in practice, it must bear some semblance to future performances. Therefore, the reproducibility of future team attacking performances by a set of previous matches was used to empirically evaluate the suitability of EPV model zone sizes in rugby league. The Kullback-Leibler (KL) Divergence (Kullback and Leibler, 1951) was used to evaluate this reproducibility. The KL Divergence

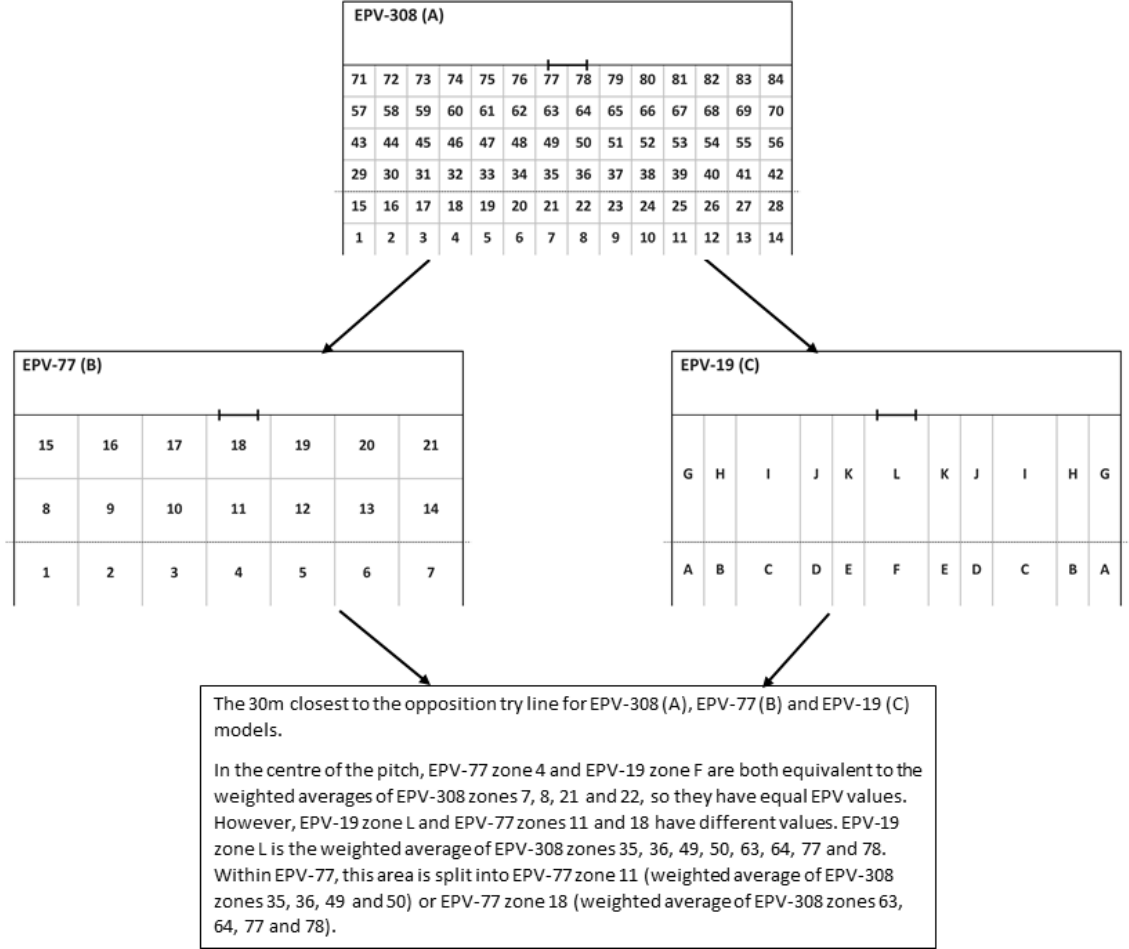


Figure 3.5: Depiction of similarities and differences between EPV-308 (A), EPV-77 (B) and EPV-19 (C) in the 30m closest to the opposition try line. Each zone from the EPV-77 and EPV-19 is a weighted average of the EPV-308 zones they are composed of. Dotted line represents the opposition 20m line.

was used to assess the number of matches required before a team's future attacking performances (i.e. their next fixture's performances) were reproduced by previous fixtures.

Definition 3.2.3. The KL Divergence D_{KL} calculates the similarity between a true distribution P and an approximating distribution Q :

$$D_{KL}(P || Q) = \sum_{s \in S} P(s) \log \frac{P(s)}{Q(s)}, \quad (3.4)$$

where s refers to the finite elements of the distribution, in this case the zones within the EPV models

The KL Divergence is a measure used in information theory and provides an understanding of the similarity between two distributions of values. It is an unbounded measure, where a value of 0 indicates two distributions are perfectly matched, but a value of infinity indicates that there is no relationship between the two distributions. A value of infinity typically occurs when an element of the approximating distribution has no value (i.e. has not been visited), but has been visited by the true distribution.

In this study, the subsequent fixture was used as the true distribution. The 1 to 10 previous fixtures were separately used as the approximating distribution to identify how many previous fixtures were required to establish an understanding of the team's future attacking performances. Prior to evaluation by the KL Divergence, all EPV's were normalised to produce an EPV distribution (EPVD). The EPV distribution for each zone s in k matches prior to match i was calculated as

$$EPV_M(s) = \frac{\sum_{k=1}^{k=i-1} \sum_{s=1}^S G_{mk}(s)}{\sum_{k=1}^{k=i-1} \sum_{s=1}^S G_m(s)}, \quad (3.5)$$

where M is the set of k matches and S is the set of all zones in the model.

The percentage of non-infinity values was used to provide an understanding of how many of the subsequent match's zones were visited in the previous matches. The KL Divergence value was used as a measure of similarity between the two EPV distributions' values. All results are provided as a mean and standard deviation values across the twelve Super League clubs.

3.2.5 Z-score Analysis

After identifying the most suitable zone sizes for use within rugby league (EPV-19 model), z-score analysis was proposed as a novel method of quantifying individual team attacking performances. Each team's EPV distribution across the whole 2019 Super League season was calculated for the EPV-19 via Equation 3.5. Z-score analysis of the EPV distributions was used to calculate a standardised value evaluating how the proportion of match EPV a team obtained from a zone compared to the average across all teams in the Super League.

Values of +1 and +2 z-scores were chosen to represent a greater and much greater proportion of match EPV generated by the zones relative to the average team (i.e. the zone was more valuable to the team), values of -1 and -2 were used to represent a lower and much lower proportion of match EPV generated (i.e. the zone was less valuable to the team).

3.3 Results

The aim of this study was to apply, adapt and extend previous work in rugby league (Kempton et al., 2016) and football (Singh, n.d.), which used EPV models with arbitrarily selected zone sizes. To achieve this aim, 180 xml files provided by Opta were preprocessed into a table of 59,233 rows and 7 columns (Table 3.4). Using this data, an MRP, similar to Kempton et al. (2016), produced the EPV-308 model from 308 $\sim 5\text{m} \times 5\text{m}$ zones. To adapt and improve previous work (Kempton et al., 2016), five further EPV models were produced from the EPV-308 zones. In total, six models were produced - two with fixed zone sizes: EPV-308 and EPV-77; and four with zones aggregated or differentiated based on statistical evaluation of their column and row match EPVs: EPV-37, EPV-19, EPV-13 and EPV-9. The KL Divergence was used to empirically evaluate the zone configurations' ability to reproduce future team attacking performances. The EPV-19 was validated as the most suitable model for use within rugby league. Z-score analysis was proposed as a novel method of quantifying the differences in team attacking performances across the 2019 Super League season using the EPV-19.

3.3.1 EPV models

Six EPV models were developed in this study, two with fixed zone sizes: EPV-308 ($\sim 5\text{m} \times 5\text{m}$ zones), EPV-77 ($\sim 10\text{m} \times 10\text{m}$ zones, equivalent to Kempton et al. (2016)); and four with statistically calculated variable zone sizes: EPV-37, EPV-19, EPV-13 and EPV-9. Figure 3.6 illustrates the zone values for all six EPV models. Tables 3.5 to 3.16 outline the results of the statistical analyses conducted to obtain the four variable zone size models (EPV-37, EPV-19, EPV-13, EPV-9). Tables 3.5 and 3.6 provide the column match EPV and row match EPV respectively for the EPV-308. Tables 3.7 and 3.8 outline the column and row match EPVs, aggregated after visual inspection, from which the statistical tests for the EPV-37, EPV-19, EPV-13, EPV-9 began.




Figure 3.6: EPV-308 (A), EPV-77 (B), EPV-37 (C), EPV-19 (D), EPV-13 (E), EPV-9 (F). These are the six EPV-models produced in this study, where the number represents the number of zones present within the model. Lines represent the 20m line and 50m line. All zone values are shaded to the same scale.

The statistical tests used to identify the columns and rows of EPV-37, where the SESOI was 0.5, are shown in Tables 3.9 and 3.10. The tests used to identify the columns and rows of EPV-19, where the SESOI was 1.0, are shown in Tables 3.11 and 3.12. The tests used to identify the columns and rows of EPV-13, where the SESOI was 1.5, are shown in Tables 3.13 and 3.14. The tests used to identify the columns and rows of EPV-9, where the SESOI was 2.0 are shown in Tables 3.15 and 3.16.

Across all models, there is a general trend that the closer the zone is to the opposition try line, the more valuable it is. Similarly, central zones are more valuable than wider

Table 3.5: Linear mixed model derived average column match EPV for EPV-308 columns. Bracketed values are metres included for column and 95% confidence intervals for column match EPV.

Column	Column match EPV	Column	Column match EPV
1 (0-3m)	1.41 (1.04 - 1.78)	8 (34-38m)	12.27 (11.92 – 12.61)
2 (4-8m)	2.72 (2.38 - 3.06)	9 (39-43m)	9.72 (9.38 – 10.07)
3 (9-13m)	4.52 (4.18 - 4.87)	10 (44-48m)	7.76 (7.41 – 8.10)
4 (14-18m)	6.06 (5.72 – 6.40)	11 (49-53m)	5.71 (5.37 – 6.05)
5 (19-23m)	7.10 (6.75 – 7.44)	12 (54-58m)	5.43 (5.09 – 5.78)
6 (24-28m)	8.57 (8.23 – 8.92)	13 (59-63m)	2.55 (2.20 – 2.89)
7 (29-33m)	12.78 (12.44 – 13.12)	14 (64-68m)	1.13 (0.76 – 1.50)

Table 3.6: Linear mixed model derived average row match EPV for EPV-308 rows. Bracketed values are metres included for row and 95% confidence intervals for row match EPV.

Row	Row match EPV	Row	Row match EPV
1 (-10 to -5m)	0.19 (-0.38 – 0.76)	12 (46-50m)	3.03 (2.77 – 3.29)
2 (-4 to 0m)	0.45 (0.09 – 0.81)	13 (51-55m)	3.59 (3.33 – 3.85)
3 (1-5m)	1.12 (0.86 – 1.38)	14 (56-60m)	4.16 (3.90 – 4.42)
4 (6-10m)	2.12 (1.86 – 2.38)	15 (61-65m)	3.67 (3.42 – 3.93)
5 (11-15m)	3.31 (3.05 – 3.57)	16 (66-70m)	4.03 (3.77 – 4.28)
6 (16-20m)	3.21 (2.95 – 3.47)	17 (71-75m)	4.91 (4.65 – 5.17)
7 (21-25m)	3.15 (2.89 – 3.41)	18 (76-80m)	5.29 (5.03 – 5.55)
8 (26-30m)	3.22 (2.97 – 3.48)	19 (81-85m)	6.72 (6.46 – 6.98)
9 (31-35m)	2.88 (2.62 – 3.14)	20 (86-90m)	9.12 (8.86 – 9.38)
10 (36-40m)	3.00 (2.75 – 3.26)	21 (91-95m)	11.43 (11.17 – 11.68)
11 (41-45m)	3.38 (3.12 – 3.64)	22 (96-100m)	6.52 (6.25 – 6.79)

zones as indicated by the darker colours in these areas. These findings are congruent with those of [Kempton et al.\(2016\)](#). Similarly, in all six models much greater value is generated within 20-30m of the opposition try line, compared with more than 30m away from the try line. This finding is similar to previous research within football, which shows that the chance of scoring is significantly reduced to below 7% when shots are taken from outside the 18 yard box ([Spearman,2018](#)). The identification of these zones of value in

Table 3.7: EPV-308 symmetrically aggregated column match EPVs. Bracketed values are metres included for column and 95% confidence intervals for column EPV from linear mixed model.

Column	Column match EPV per fixture
1 (1,14)	1.27 (0.97 – 1.57)
2 (2,13)	2.63 (2.35 – 2.92)
3 (3,12)	4.98 (4.70 – 5.26)
4 (4,11)	5.88 (5.60 – 6.17)
5 (5,10)	7.43 (7.14 – 7.71)
6 (6,9)	9.15 (8.87 – 9.43)
7 (7,8)	12.52 (12.24 – 12.81)

Table 3.8: EPV-308 sequentially aggregated row match EPVs. Bracketed values are metres included for row and 95% confidence intervals for row EPV from linear mixed model.

Row	Row match EPV per fixture
1 (1,2)	0.32 (-0.02 – 0.67)
2 (3,4)	1.62 (1.42 – 1.82)
3 (5,6)	3.26 (3.06 – 3.46)
4 (7,8)	3.19 (2.99 – 3.39)
5 (9,10)	2.94 (2.74 – 3.14)
6 (11,12)	3.20 (3.00 – 3.41)
7 (13,14)	3.87 (3.67 – 4.08)
8 (15,16)	3.85 (3.65 – 4.05)
9 (17,18)	5.10 (4.90 – 5.30)
10 (19,20)	7.92 (7.72 – 8.12)
11 (21,22)	8.97 (8.77 – 9.18)

all six models provides a new method through which attacking possessions can be valued. Furthermore, they provide a valuable methodology through which the zones visited in tactical set plays could be measured to establish which play may be most advantageous against a given team.

Table 3.9: Statistical comparison between columns at an SESOI of 0.5 for EPV-37. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 0.5. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally.

Columns included	Difference	Significance	Combine?
2 - 1	1.36 (1.06 – 1.66)	$P < 0.0001$	No
3 - 2	2.35 (2.06 – 2.63)	$P < 0.0001$	No
4 - 3	0.91 (0.62 – 1.18)	$P = 0.0025$	No
5 - 4	1.54 (1.26 – 1.82)	$P < 0.0001$	No
6 - 5	1.72 (1.44 – 2.00)	$P < 0.0001$	No
7 - 6	3.38 (3.09 – 3.66)	$P < 0.0001$	No

Table 3.10: Statistical comparison between rows at an SESOI of 0.5 for EPV-37. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 0.5. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area.

Rows included	Difference	Significance	Combine?
2 - 1	1.30 (0.93 – 1.67)	$P < 0.0001$	No
3 - 2	1.64 (1.41 – 1.87)	$P < 0.0001$	No
4 - 3	-0.07 (-0.30 – 0.16)	$P = 0.5542$	Yes
5 - 3,4	-0.28 (-0.48 – -0.08)	$P = 1.000$	Yes
6 - 3,4,5	0.07 (-0.12 – 0.26)	$P = 1.000$	Yes
7 - 3,4,5,6	0.73 (0.54 – 0.91)	$P = 0.0079$	No
8 - 7	-0.02 (-0.26 – 0.21)	$P = 1.000$	Yes
9 - 7,8	1.24 (1.04 – 1.44)	$P < 0.0001$	No
10 - 9	2.82 (2.59 – 3.05)	$P < 0.0001$	No
11-10	1.06 (0.82 – 1.29)	$P < 0.0001$	No

Table 3.11: Statistical comparison between columns at an SESOI of 1.0 for EPV-19. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 1.0. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally.

Columns included	Difference	Significance	Combine?
2 - 1	1.36 (1.06 – 1.66)	P = 0.0096	No
3 - 2	2.35 (2.06 – 2.63)	P < 0.0001	No
4 - 3	0.91 (0.62 – 1.18)	P = 0.7441	Yes
5 - 3,4	2.00 (1.75 – 2.24)	P < 0.0001	No
6 - 5	1.72 (1.44 – 2.00)	P < 0.0001	No
7 - 6	3.38 (3.09 – 3.66)	P < 0.0001	No

Table 3.12: Statistical comparison between rows at an SESOI of 1.0 for EPV-19. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 1.0. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area.

Rows included	Difference	Significance	Combine?
2 - 1	1.30 (0.93 – 1.67)	P = 0.0559	Yes
3 - 1,2	2.29 (2.04 – 2.53)	P < 0.0001	No
4 - 3	-0.07 (-0.30 – 0.16)	P = 1.000	Yes
5 - 3,4	-0.28 (-0.48 – -0.08)	P = 1.000	Yes
6 - 3,4,5	0.07 (-0.12 – 0.26)	P = 1.000	Yes
7 - 3,4,5,6	0.73 (0.54 – 0.91)	P = 0.9983	Yes
8 - 3,4,5,6,7	0.56 (0.38 – 0.74)	P = 1.000	Yes
9 - 3,4,5,6,7,8	1.72 (1.54 – 1.89)	P < 0.0001	No
10 - 9	2.82 (2.59 – 3.05)	P < 0.0001	No
11-10	1.06 (0.82 – 1.29)	P = 0.3333	Yes

Table 3.13: Statistical comparison between columns at an SESOI of 1.5 for EPV-13. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 1.5. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally.

Columns included	Difference	Significance	Combine?
2 - 1	1.36 (1.06 – 1.66)	P = 0.8165	Yes
3 - 1,2	3.03 (2.78 – 3.28)	P < 0.0001	No
4 - 3	0.91 (0.62 – 1.18)	P = 1.000	Yes
5 - 3,4	2.00 (1.75 – 2.24)	P < 0.0001	No
6 - 5	1.72 (1.44 – 2.00)	P = 0.0614	Yes
7 - 5,6	4.24 (4.00 – 4.48)	P < 0.0001	No

Table 3.14: Statistical comparison between rows at an SESOI of 1.5 for EPV-13. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 1.5. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area.

Rows included	Difference	Significance	Combine?
2 - 1	1.30 (0.93 – 1.67)	P = 0.8581	Yes
3 - 1,2	2.29 (2.04 – 2.53)	P < 0.0001	No
4 - 3	-0.07 (-0.30 – 0.16)	P = 1.000	Yes
5 - 3,4	-0.28 (-0.48 – -0.08)	P = 1.000	Yes
6 - 3,4,5	0.07 (-0.12 – 0.26)	P = 1.000	Yes
7 - 3,4,5,6	0.73 (0.54 – 0.91)	P = 1.000	Yes
8 - 3,4,5,6,7	0.56 (0.38 – 0.74)	P = 1.000	Yes
9 - 3,4,5,6,7,8	1.72 (1.54 – 1.89)	P = 0.0086	No
10 - 9	2.82 (2.59 – 3.05)	P < 0.0001	No
11-10	1.06 (0.82 – 1.29)	P = 0.9999	Yes

Table 3.15: Statistical comparison between columns at an SESOI of 2.0 for EPV-9. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the columns and whether this was significant according to an SESOI of 2.0. Combine? Indicates whether the column values were averaged before comparing to the next column. Comparisons began at the outermost columns before progressing more centrally.

Columns included	Difference	Significance	Combine?
2 - 1	1.36 (1.06 – 1.66)	P = 1.000	Yes
3 - 1,2	3.03 (2.78 – 3.28)	P < 0.0001	No
4 - 3	0.91 (0.62 – 1.18)	P = 1.000	No
5 - 3,4	1.99 (1.75 – 2.24)	P = 0.518	Yes
6 - 3,4,5	3.06 (2.83 – 3.29)	P < 0.0001	No
7 - 6	3.38 (3.09 – 3.66)	P < 0.0001	No

Table 3.16: Statistical comparison between rows at an SESOI of 2.0 for EPV-9. Where numbers are separated by commas, they have been averaged within the linear mixed model. Difference and significance provide the mean difference (95% confidence intervals) between the rows and whether this was significant according to an SESOI of 2.0. Combine? Indicates whether the row values were averaged before comparing to the next row. Comparisons began inside the attacking teams try area before progressing up the pitch to the opposition try area.

Rows included	Difference	Significance	Combine?
2 - 1	1.30 (0.93 – 1.67)	P = 0.9999	Yes
3 - 1,2	2.29 (2.04 – 2.53)	P = 0.0112	No
4 - 3	-0.07 (-0.30 – 0.16)	P = 1.000	Yes
5 - 3,4	-0.28 (-0.48 – -0.08)	P = 1.000	Yes
6 - 3,4,5	0.07 (-0.12 – 0.26)	P = 1.000	Yes
7 - 3,4,5,6	0.73 (0.54 – 0.91)	P = 1.000	Yes
8 - 3,4,5,6,7	0.56 (0.38 – 0.74)	P = 1.000	Yes
9 - 3,4,5,6,7,8	1.72 (1.54 – 1.89)	P = 0.9991	Yes
10 - 3,4,5,6,7,8,9	4.31 (4.13 – 4.48)	P < 0.0001	No
11-10	1.06 (0.82 – 1.29)	P = 1.000	Yes

3.3.2 Reproducibility of Future Attacking Performances

To empirically evaluate the suitability of different zone configurations, the EPV models' ability to reproduce future attacking performances was measured via the KL Divergence. Table 3.17 shows the percentage of non-infinity values for all six models after 1-10 previous matches (i.e. the percentage of fixtures where all subsequent match's zones had been visited in the previous matches). For EPV-308, there were only three occasions where this was greater than 0% (8, 9 and 10 previous matches). There was a consistent increase in the percentage of non-infinity values as the number of previous fixtures increased for EPV-77 and EPV-37, peaking at $77 \pm 8 \%$ and $97 \pm 4 \%$ respectively after 10 previous fixtures. For EPV-19, there was a large increase in the percentage of non-infinity values before 6 previous fixtures were considered, after which limited change was observed (95-98 % from 6 to 10 fixtures). A similar trend was present for EPV-13 before 3 previous fixtures were considered (96-100 % from 3 to 10 fixtures). In EPV-9, all values were not infinity after only 3 previous fixtures.

Figure 3.7 shows the KL Divergence for EPV-77, EPV-37, EPV-19, EPV-13 and EPV-9. After 8 (KL Divergence = 1.50 ± 0.19), 9 (KL Divergence = 1.41 ± 0.15) and 10 (KL Divergence = 1.44 ± 0.15) previous matches, the KL Divergence for EPV-308 was still greater than any other model after any number of previous matches, rendering it considerably less useful than all other models. The KL Divergence reduced as more previous matches were considered in all EPV models in Figure 3. For EPV-37, EPV-19, EPV-13 and EPV-9, the majority of this reduction occurred between 1 and 3 previous matches before the values stabilised. For EPV-77, the values stabilised after six previous matches.

Although commonly identified as a key element of any model quantifying team attacking performances, few studies have attempted to evaluate the reproducibility of future attacking performances in their models (Sarmiento et al., 2014). The results showed that although the EPV-308, EPV-77 and EPV-37 provide significantly more variability than either the EPV-19, EPV-13 and EPV-9 with regards to the values of different zones, they had poor reproducibility between fixtures. This was noticeable in both the percentage of subsequent match zones visited and in the similarity in reward distributions between the previous and subsequent matches. The EPV-308, EPV-77 and EPV-37 therefore have limited application in practice when evaluating team attacking performances. By contrast, the

Table 3.17: Percentage of non-infinity zones for each model, providing the percentage of matches where the complete set of the subsequent match' s zones were visited in the previous n matches. Values are mean (standard deviation) percentage across all clubs.

Model	Number of previous matches									
	1	2	3	4	5	6	7	8	9	10
EPV-308	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	2(2)	4(4)	5(5)
EPV-77	0(0)	2(2)	13(6)	26(7)	40(9)	51(12)	63(9)	69(8)	72(8)	77(8)
EPV-37	1(2)	19(7)	48(8)	67(11)	77(8)	85(9)	89(9)	92(6)	95(5)	97(4)
EPV-19	23(8)	59(6)	76(6)	88(8)	91(5)	95(4)	96(3)	96(3)	98(3)	98(3)
EPV-13	61(9)	89(3)	96(4)	98(4)	99(2)	100(0)	100(0)	100(0)	100(0)	100(0)
EPV-9	86(5)	99(2)	100(0)	100(0)	100(0)	100(0)	100(0)	100(0)	100(0)	100(0)




Figure 3.7: KL Divergence values for EPV-77, EPV-37, EPV-19, EPV-13 and EPV-9. A lower value indicates greater similarity in reward distributions between previous matches and the subsequent match. EPV-308 not included as values could not be calculated for the first seven matches due to no non-infinity values being present. Line provides mean value, shaded area indicates standard deviation.

EPV-19, EPV-13 and EPV-9 all showed excellent reproducibility between fixtures. When six previous matches were considered, these three models were able to visit all zones in the subsequent match on 95-100% of occasions. Furthermore, the EPV distributions had low KL Divergence values indicating that proportion of points obtained from each zone was also very similar to the subsequent match.

Six matches is a relatively small number of matches to consider given the excellent reproducibility shown, suggesting that any of the EPV-19, EPV-13 or EPV-9 models could be used to evaluate team attacking performance in rugby league. However, it is the usefulness of the zones generated that should define which model is used in practice. The EPV-19 and EPV-13 both contain four rows (-10 to 10m, 10-70m, 70-80m, 80-100m), whereas the EPV-9 only contains three rows (-10m to 10m, 10-80m, 80-100m). As five of the six models produced suggest that the value of zones in the 70-80m row can be

differentiated from those around it, it is possible that the EPV-9 has oversmoothed the data, reducing its usefulness in practice. The EPV-19 and EPV-13 models only differ in the manner through which they split group the columns along the x-axis. The EPV-19 has more columns (6), separating out the widest and second most central areas of EPV-13. This results in the EPV-13 having a smoother progression of zone values from wide to central. However, it also results in the value of the zones just outside the posts being much smaller relative to EPV-19 and EPV-37. Given the value of the central zones is of upmost importance for conversions of tries, the EPV-19 may be considered the more useful set of zones, but either model could be used in practice.

3.3.3 Quantification of Teams' Attacking Performances

To quantify team attacking performances across the 2019 Super League season, z-score analysis of the EPV-19 model values was used. Figure 3.8 provides a numbered zone breakdown for the EPV-19. Figure 3.9 depicts the z-score analysis of each team's attacking performances across the 2019 Super League season using the EPV-19 model. Using z-score analysis of the match EPV, it is clear that Hull generates a greater proportion of match EPV from different zones to Leeds when 10m-70m from its own try line. Hull also gains greater match EPV from wide areas (zones 3 and 4), whereas Leeds gains more match EPV centrally (zones 5-7) relative to other zones. The identification of these zones pre-match could assist teams in their tactical preparations. Furthermore, Figure 3.9 shows those teams who spread their attack more evenly. For example, from 80-100m, St Helens obtained a small proportion of its match EPV from the widest zone (14) compared to other Super League teams, but they generated similar proportions of match EPV across the rest of the zones. It is possible that this ability to generate value close to the average team across the majority of the pitch made the team difficult to defend against and could explain why they were one of the top points scorers across the season. The use of z-score analysis has strong potential as a method through which the areas on a pitch where an opposition team may attack can be highlighted quickly and efficiently, regardless of the EPV zone configuration used, enabling tactical preparations for future matches to be tailored to the opposition.

The results of this study were presented to the England Performance Unit, where the work was widely praised as significantly more advanced than anything previously pro-




Figure 3.8: EPV-19 zones numbered so they can be distinguished from each other. Where numbers are repeated, both sides of the pitch make up the same zone (e.g. zone 2 is comprised of the widest $\sim 5\text{m}$ on both sides of pitch, between 10m and 70m from the team's own try line).

posed within the sport. However, reservations were made with respect to its immediate usability within the sport as the location of play-the-balls is heavily dependent on both the actions taken by the team in possession and the strategies of the defensive team. It was suggested that it could definitely be used as a starting point for further analysis or to confirm the anecdotal observations of performance analysts/coaches. To this end, the Performance Unit validated the reliability of the model by agreeing that the z-score analysis supported their anecdotal opinions of teams' attacking performances in the 2019 Super League season.

All analyses were conducted using bespoke Python scripts (Python 3.7, Python Software Foundation, Delaware, USA) or via Proc Mixed (SAS University Edition, SAS In-




Figure 3.9: Z-score analysis of teams' attacking performances for the whole 2019 Super League season. Numbers 1-19 reflect the zone numbers in Figure 3.8. A greater value indicates a greater proportion of EPV was obtained from the zone than the average Super League 2019 team. Distances are measured from the teams own try line.

stitute, Cary, NC). The analyses were completed using a HP-Pavilion Laptop with an Intel Core i5-8250U 1.60 GHz processor and 8GB RAM. The Monte Carlo every visit algorithm used for the calculation of EPV-308 state values was timed at 17.24s.

3.4 Summary

This study applied, adapted and extended the work of [Kempton et al.\(2016\)](#) in rugby league and [Singh\(n.d.\)](#) in football. Two EPV models with fixed zone sizes (EPV-308 and EPV-77) were produced using an MRP similar to a previous study ([Kempton et al., 2016](#)). In contrast to previous studies using Markovian approaches ([Kempton et al.,2016](#); [Singh,n.d.](#)), four further variations of the EPV model (EPV-37, EPV-19, EPV-13, EPV-9) were produced using statistical analysis of the match EPV produced by columns and rows of zones from the EPV-308. Empirical evaluation of all six models, comparing their ability to reproduce future team attacking performances, identified the EPV-19 as the most suitable model for use within rugby league. A novel method of quantifying team attacking performances was proposed. It was applied to the empirically selected EPV-19 model to understand its validity, but could have been applied to any of the EPV models shown in this chapter. Practitioners can use the model to produce a high level understanding of the areas on the pitch that different teams generate value. The results indicate that only six previous matches need to be considered to provide this understanding.

To validate the work, the results were published in PLOS One and presented to the England Rugby League Performance Unit. The England Performance Unit coaches provided partial validation for the reliability of the model by agreeing that it provided an accurate representation of the attacking performances of different teams across the 2019 Super League season. However, they questioned the long-term usability of the model due to its use of play by play, rather than action by action data, which limited the insights that could be generated.

Although the work provides an excellent high level understanding of the areas on the pitch that different teams generate value, it is only a starting point for this type of analysis and is subject to limitations. The first of these, as indicated by the England Performance Unit, is that it only uses the first action of each play. This means it does not consider any action (e.g. pass, kick or run) individually so there is no possibility of adapting the model to evaluate player performances or understand team sequences of play. The second is that it does not directly predict future team attacking performances, nor does it provide any information as to whether being aware of future opposition team attacking performances can help a team to win fixtures.

In this study, it was not possible to conduct action by action analysis due to the quality

of data required for such analyses being unavailable in the 2019 season. However, in 2020 this data became available. Therefore, the next study will build upon the work completed in this chapter and consider individual actions within the EPV framework to evaluate player and team performances.

Markov Decision Processes for the Evaluation of Player and Team Performances

4.1 Introduction

A key limitation of the analysis conducted in Chapter 3 was that it could only consider play by play data (i.e. the first action of every play). Analysing action by action data (i.e. the location of each individual action) would allow the performances of teams and players to be evaluated in more detail. However, in 2019, such data was not available. In 2020, Opta improved their data collection processes and location data became available for every action in the Super League, allowing action by action analyses to be conducted for the first time. Consequently, this study builds on the work described in Chapter 3 by proposing a model which evaluates player and team performances in rugby league using a Markov Decision Process (MDP) framework, applying and adapting ideas previously used in ice hockey (Routley and Schulte, 2015) and football (Singh, n.d.). Action impact ratings are adapted from ice hockey (Routley and Schulte, 2015) to analyse player and team performances using action values from the MDP for the first time in rugby league. Furthermore, context nodes are adapted from ice hockey (Routley and Schulte, 2015) and football (Decroos et al., 2019) to allow further insight into player and team performances to be gleaned. The study was presented at the UKCI 2021 conference (Thomas Sawczuk,

Anna Palczewska, Ben Jones. "Markov Decision Processes with contextual nodes as a method of assessing attacking player performance in rugby league" in *20th UK Workshop on Computational Intelligence*, Aberystwyth, 2021), published in *Advances in Computational Intelligence Systems* (Thomas Sawczuk, Anna Palczewska, Ben Jones. Markov Decision Processes with contextual nodes as a method of assessing attacking player performance in rugby league. In: Jansen, T., Jensen, R., Mac Parthalain, N., Lin, CM. (eds) *Advances in Computational Intelligence Systems. UKCI 2021*. Advances in Intelligent Systems and Computing, vol 1409. Springer, Cham.) and presented internally to the Sports Science department at Leeds Beckett University, whose feedback provided validation for the model.

4.2 Methodology

In this section, the methodology for this study is presented. It describes the data used and the preprocessing steps required to prepare the data for analysis. The Markov Reward Process framework is recalled and extended to describe the Markov Decision Process and state, action values used in this study. Context nodes and action impact ratings are adapted for rugby league from ice hockey (Routley and Schulte, 2015) and football (Decroos et al., 2019) to provide insights into player and team performances.

4.2.1 Data

In this study, event level match-play data were obtained from Opta (Stats Perform, London, UK) for the 2020 Super League season. The data was produced via human annotation of the actions taking place and was downloaded from www.optaprorugby.com. In addition to the 200 actions coded by Opta in the 2019 season, 33 additional unique actions were coded in 2020. These actions and their definitions are provided in Appendix B. As with the 2019 season, data for each of the 63 matches in the 2020 Super League season was provided in separate xml files, which each contained four nodes (Table 3.1).

The data received from Opta was formatted as described in Section 3.2.1. As with Chapter 3, missing data issues ensured that the same 13 variables from the match data node (Table 3.2) were deemed to be useful for this study. However, whereas in the 2019 data, location data was only fully available on a play by play basis (i.e. for the first

action of each play), in 2020 it became available on an action by action basis (i.e. for all actions). This increased data availability considerably improved the quality of analysis that could be completed, allowing the MRP approach of Chapter 3 to be extended to an MDP approach. However, only a limited amount of this higher quality data was available as the 2020 season was curtailed by the COVID-19 pandemic.

In the 2020 Super League season, 63 matches were contested by 11 teams and 352 players. A total of 146,493 match events were recorded (median 2331 events per match, interquartile range 2239-2405). Across the season, 481 tries were scored (381 successful conversion kicks, 100 unsuccessful conversion kicks), 50 penalty goals were attempted (46 successful, 4 unsuccessful) and 33 drop goals were attempted (16 successful, 17 unsuccessful).

4.2.2 Data Preprocessing

Figure 4.1 provides an overview of the data preprocessing completed within this study. The preprocessing steps converted the 63 raw xml files described above to a single file of 77,045 rows and 9 columns. Full details of each preprocessing step are provided below.

The match action data, provided in 63 xml files, were collected and integrated into a single dataset of 146,493 rows and 28 columns. To do this, the match data node variables were extracted from each xml file into separate tables and then concatenated together. The team data nodes were also extracted into tables for each fixture. These were used to map team unique identifiers to the real team name. Similarly, Opta's action definitions and unique identifiers (Appendix B) were mapped to the real name of the coded actions.

As previously stated, in 2020, location data was available for every action on the pitch. However, with only 146,493 actions across the season, it was not possible to consider a model which included both attacking and defensive actions due to the size of transition matrix this would require. Instead, all defensive and auxiliary actions were removed, providing a total of 77,045 attacking actions. For the purposes of this study, the attacking actions were organised into the seven grouped actions identified in Table 4.1. These grouped actions represented the most common attacking actions, which when grouped together in sequences described every attacking possession within the dataset.

Although location data was available for all actions, the coding protocol used by Opta was still geared towards considering accumulated action count measures, rather than ac-

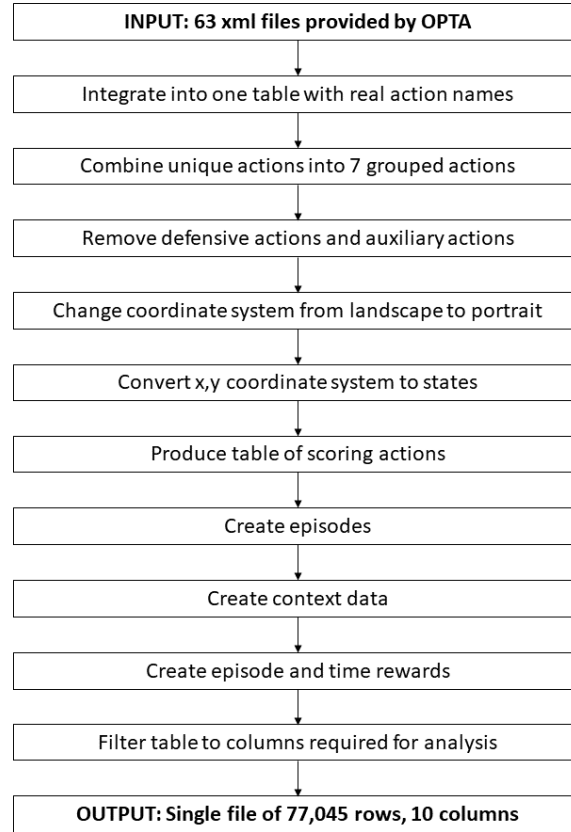


Figure 4.1: Data preprocessing workflow

tion by action analysis. Consequently, there were three action ordering issues that needed to be resolved to ensure that the sequence of actions considered by the MDP accurately represented true sequences of match play. These were: runs being coded in the action before a player had caught the ball; penalties being won before either a run or catch (or both) had occurred for the same player; and positional kicks occurring after a player had won a penalty, when the positional kick happened beforehand. All three ordering issues were resolved at this point.

As with the 2019 data, each action in the 2020 dataset was marked with an x, y location, denoting the position of the pitch relative to a standard 100m x 68m rugby league pitch. For the purposes of this study, these locations were binned into the EPV-19 zones produced in Chapter 3. The EPV-19 zones were chosen as these were empirically eval-

Table 4.1: Definitions of grouped actions used in this study.

Action	Definition	Count
Catch	A successful receipt of the ball, either from a kick or a loose ball	3198
Pass	A thrown pass aimed at a player on the same team	29192
Positional Kick	A kick aimed to advance the position of the team or for a teammate to attempt to catch, but not aimed specifically at goal	2955
Run	A run made by a player, whilst in possession of the ball	23543
PTB	The restart action performed by a player after he has been tackled, indicates the beginning of the next play	17689
Penalty Won	The act of winning a penalty	385
Goal Kick	A kick at goal, either from a penalty or a drop goal	83

uated as the most suitable set of zones for use in rugby league in Chapter 3. Figure 4.2 plots the EPV-19 zones against a standard rugby league pitch. Actions taking place in the opposition try area were included in the dataset to allow the rewards achieved to be backpropagated through the sequence of actions, but were not included or discussed when rating players because all attacking players will attempt to ground the ball for a try in this situation regardless of the previous action taken.

Three variables were prepared for use as context nodes: period; match score; and team. The period refers to which of the two 40 minute halves was being played at the time of the action (1st or 2nd) and the match score refers to whether the team in possession is winning, drawing or losing at the beginning of the possession. These contextual factors were used to identify differences in players' performances as both could affect the team's strategy and therefore individual players' decision making, either due to fatigue (second half actions are performed under greater fatigue than first half actions) or the urgency with which points need to be scored (losing is likely to be more urgent than winning). Team was chosen as a context node to allow the estimation of individual state, action values for different teams.

As with Chapter 3, each time-step in the possession was assigned a value based on the true points scored due to the action taken. Therefore, +6 was rewarded for a converted try,




Figure 4.2: Standard rugby league pitch (A). Lines represent try line, 20m from the try line and half way line. Rugby league pitch split into the 19 states used within this analysis (B).

+4 for an unconverted try, +2 for a successful penalty goal and +1 for a successful drop goal. Errors were not penalised, so all other actions received a value of 0.

The data was prepared for use within an MDP by separating it into unique attacking possessions. An attacking possession could start with any of the actions listed (except a goal kick), and ended due to handover, loss of possession due to error/foul play, points being scored, or a goal kick attempt. The possessions were only deemed to have ended if the opposing team successfully caught or passed the ball. Therefore, if a positional kick was attempted by the team in possession, and was dropped by the opposition, the resulting set of 6 tackles for the team who kicked the ball was considered part of the same attacking possession as previously. If the positional kick was successfully caught by the opposition, it began the process of a new unique attacking possession by the opposition team. Table 4.2 outlines a typical sequence for an episode used within this study. In the table, St Helens are shown to progress towards the opposition try line over the duration of

Table 4.2: Example episode sequence. FxID refers to fixture identification number, PLID refers to player identification number.

FxID	Team	PLID	Episode	Time	Perio	Score	Zone	Action	Reward
31	St	51	294	7	d 2	W	7	n Pass	0
31	Helens	73	294	8	2	W	7	Run	0
31	St	73	294	9	2	W	5	PTB	0
31	Helens	51	294	10	2	W	5	Pass	0
31	St	74	294	11	2	W	5	Run	0
31	Helens	74	294	12	2	W	7	PTB	0
31	St	51	294	13	2	W	7	Pass	0
31	Helens	28	294	14	2	W	7	Run	0
31	St	28	294	15	2	W	19	Pass	0
31	Helens	84	294	16	2	W	19	Run	6
31	St	19	295	1	2	W	4	Catc	0

3 plays (each play starts after a play-the-ball (PTB) but forms part of the same attacking possession), before a try is scored in the 16th action of their possession. Player 19 starts the next attacking possession for St Helens by catching the opposition restart kick, which is not included in the episode sequences.

4.2.3 Markov Decision Processes

The aim of this study was to extend the method presented in Chapter 3 by applying and adapting ideas from ice hockey (Routley and Schulte, 2015) and football (Decroos et al., 2019; Singh, n.d.). Drawing upon these ideas, methods of analysing player and team performances in rugby league using the action values from an MDP were proposed. Recall from Section 3.2.3 the definition of a Markov Reward Process (MRP) as a stochastic process, which extends a Markov chain by adding a reward to each state (Howard, 1971). An MDP extends an MRP by adding an action to the framework (Definition 4.2.1). MDP's are mathematical tools for sequential decision making in a stochastic environment (White and White, 1989) and form the underlying basis of the reinforcement learning framework, whereby an agent learns an optimal solution to a problem by interacting with an environment through the use of actions (Sutton and Barto, 2018).

Definition 4.2.1. A *Markov Decision Process* is a tuple (S, A, P_a, R_a, γ) where:

- S is a finite set of states
- A is a finite set of actions
- P is a transition probability matrix
- R is a reward function
- γ is a discount factor ($\gamma \in [0, 1]$)

In this study, a set of actions is added to the model. Within an MDP framework, an action defines the choice made by an agent given the state (or environment) it observes (Sutton and Barto, 2018). In rugby league, the actions are mostly movements of the ball performed by the players (e.g. pass, kick, run), but can also be defined as key moments in the game (e.g. winning a penalty, which can result in a longer attacking possession). The finite set of actions used in this study are described in Table 4.1.

The states, transition probability matrix and rewards were defined in Section 3.2.3. Their representation within rugby league remains the same in this study: the location (or zone) on the pitch can still be considered the state; the transition probability matrix still represents the probability of moving between states, although it is now conditional on the

state, action tuple rather than just the state; and the reward can be defined as the points outcome achieved by the action in the team's attacking possession.

Recall from Definition 3.2.2, the value function for a state. With the inclusion of actions within the MDP framework, the value function for each state, action pair must be updated as shown in Definition 4.2.2.

Definition 4.2.2. The state, action value function $Q_{(s,a)}$ for action a performed in state s can be defined as

$$Q_{(s,a)} = \mathbf{E}[G_t | S_t = s, A_t = a], \quad (4.1)$$

where Q is the value function for choosing action a in state s .

Alongside the inclusion of actions in the MDP, this study also includes context nodes, building on work previously conducted in ice hockey (Routley and Schulte, 2015) and football (Decroos et al., 2019). Context nodes provide an additional level of detail surrounding the results of the analysis. In rugby league, they can provide information surrounding player ratings when a team is winning, drawing or losing, or in the first or second half of the match. Furthermore, by including the team as a context node, individual state, action values can be calculated for each team within the MDP framework. The updated value function for context c , state s , action a is provided in Definition 4.2.3.

Definition 4.2.3. The context, state, action value function $Q_{(c,s,a)}$ for action a performed in state s , in context c can be defined as

$$Q_{(c,s,a)} = \mathbf{E}[G_t | C_t = c, S_t = s, A_t = a], \quad (4.2)$$

where Q is the value function for choosing action a in state s and context c .

To solve the MDP, the Monte Carlo every visit algorithm was implemented. This algorithm was also used in Chapter 3 and justification is provided in Section 3.2.3. The Monte Carlo every visit algorithm calculates the empirical mean of each state, action value by summing the discounted rewards accumulated and dividing by the total number of times the state, action pair is used. The algorithm allows every visit to every state, action pair to be valued rather than just the first occurrence within an episode. This is particularly important in rugby league as multiple actions regularly occur in similar locations or states as the defensive team attempts to stop the attacking team from progressing towards their try line.

Algorithm 2 Monte Carlo Algorithm

Input : Episode dataset, including state-action pairs and rewards, γ , empty returns list (Returns), empty counter list (Counter), empty Q-values list (Q)

Output: Q, populated with results of MDP

```

for episode do
  for timestep, t do
     $G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$ 
    Returns(s, a) = Returns(s, a) +  $G_t$ 
    Counter(s, a) = Counter(s, a) + 1
  end
end
 $Q_{(s,a)} = \frac{\text{Returns}(s,a)}{\text{Counter}(s,a)}$ 
  
```

Recall Algorithm 1, which provided the Monte Carlo every visit algorithm for state values. Pseudocode for Algorithm 2 shows the algorithm used to estimate the empirical mean return for each state, action tuple. Algorithm 2 extends the state value algorithm by ensuring that the episode returns and counter increments are added to the relevant state, action tuple in lines 4 and 5. The action value for each state, action pair $Q_{(s,a)}$ is calculated on line 8 using the relevant state, action data.

The inclusion of context nodes in this study requires a further alteration to Algorithm 2. Pseudocode for the Monte Carlo every visit algorithm simulating the empirical mean return for the context, state, action tuples is provided in Algorithm 3. Algorithm 3 differs from Algorithm 2 at the first loop. Rather than looping through each episode irrespective of any contextual factors, the algorithm now loops through each episode in each context (lines 1 and 2). The modification in lines 5 and 6 ensures that there is an individual entry for each (s, a) pair in each context, which sums the returns, and counts the visits. Similarly, the addition of c to line 9 ensures that $Q_{(c,s,a)}$ is calculated using only the state, action pairs experienced within context c.

In this study, the location of each action, binned into EPV-19 zones (Section 3.3.1), was used as the state space for the MDP. The grouped actions identified in Table 4.1 were used for the action space. The rewards were calculated based on the points outcome of the action within the attacking possession: converted try scored (+6); unconverted try scored (+4); penalty goal scored (+2); drop goal scored (+1); loss of possession or missed goal

Algorithm 3 Monte Carlo Algorithm

Input : Episode dataset, including context (c), state-action (s, a) pairs and rewards, γ , empty returns list (Returns), empty counter list (Counter), empty Q-values list (Q)

Output: Q, populated with results of MDP

```

for context c do
  for episode do
    for timestep, t do
       $G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$ 
      Returns(c, s, a) = Returns(c, s, a) +  $G_t$ 
      Counter(c, s, a) = Counter(c, s, a) + 1
    end
  end
   $Q_{(c,s,a)} = \frac{\text{Returns}(c,s,a)}{\text{Counter}(c,s,a)}$ 
end

```

attempt (0). For actions where none of these point scoring events occurred, a reward of 0 was assigned. Each time-step within the attacking possession sequence was assigned a reward, so it was possible for a state, action tuple to receive multiple rewards if more than one action was completed in a given location within the same attacking possession. However, a point scoring reward was only obtained once per attacking possession.

Four models were run to evaluate the impact of context nodes on $Q_{(s,a)}$: MDP_{nc} has no context nodes; MDP_{period} has period (1st half or 2nd half) as a context node; and MDP_{score} has match score (winning, drawing or losing) as a context node. These three models were used to evaluate player performances. A fourth model MDP_{team} used the attacking possession team as a context node to assess whether further insight into team attacking performances could be gleaned using the context nodes. Algorithm 2 was used to produce MDP_{nc} , Algorithm 3 was used for MDP_{period} , MDP_{score} and MDP_{team} . All the rewards obtained within the MDPs were directly related to the actual points scored in a rugby league match, so γ was calculated to minimise the difference between the sum of rewards obtained per team per match and the sum of Q-values obtained per team per average number of match possessions. This value was calculated for MDP_{nc} and used for all four MDPs, resulting in a γ value of 0.63, which considers actions to be beneficial to the final reward in the medium term.

4.2.4 Action Impact Ratings

To evaluate player performances in rugby league, action impact ratings were adapted from previous research in ice hockey (Routley and Schulte, 2015). Action impact ratings were calculated via Equation 4.3. They provide a method of evaluating every action completed against the state value (the weighted average of all actions completed within the state). Ratings were provided using MDP_{nc} , MDP_{period} and MDP_{score} to evaluate whether further insight into player performances could be provided through the use of context nodes. Equation 4.3 is defined as

$$\text{impact}_{(c,s,a)} = Q_{(c,s,a)} - V_{(c,s)}, \quad (4.3)$$

where $Q_{(c,s,a)}$ is the action value for action a in state s in context c and $V_{(c,s)}$ is the value of state s in context c .

The action impact rating was calculated twice: once including goal kicks as part of the weighted average of actions within the state and once without them. The action impact ratings were summed per player to create a season total and normalised to the average number of actions completed within that context per match. By calculating the action impact ratings in this way, they represent the positive or negative impact a player's actions have on the team's expected points return above or below the average player's contribution per match.

4.2.5 Team State, Action Values

To evaluate team performances, the team ID context node in MDP_{team} was used. Utilising context nodes in this manner allowed individual team state, action values to be produced. Comparing individual teams' state action values provides a method through which team performances could be evaluated by identifying those teams who generated greater value from different actions in different states. To visualise the differences and identify tactical insights, the state, action values for five actions (catch, pass, run, positional kick and play-the-ball) were plotted. These five actions were chosen as at least one action was performed in every state for every team. Multiple missing states were present for the goal kick and penalty won actions. In the plots, all team state, action values were compared to the league average state, action value. Therefore, the plots detail whether the team was better or worse at performing a given action in a given state than the league average team.

4.3 Results

The aim of this study was to evaluate player and team performances in rugby league using an MDP framework, which extended the MRP framework produced in Chapter 3 by adapting ideas previously considered in ice hockey (Routley and Schulte, 2015). To achieve these aims, the 63 xml files provided by Opta from the 2020 Super League season were preprocessed into a single file of 77,045 rows and 9 columns (Table 4.2). Using this data, four MDPs were produced: MDP_{nc} , which had no context nodes; MDP_{period} , which used period as a context node; MDP_{score} , which used match score as a context node; and MDP_{team} , which used team ID as a context node. Algorithm 2 was used for MDP_{nc} ; Algorithm 3 was used for MDP_{period} , MDP_{score} and MDP_{team} models. The MDP_{nc} , MDP_{period} and MDP_{score} models were used to evaluate player performances via action impact ratings, adapted from Routley and Schulte (2015), with and without context nodes. MDP_{team} used context nodes to evaluate individual team attacking performances.

4.3.1 Evaluation of Player Performances

To evaluate player performances in rugby league, the MRP approach used in Chapter 3 was extended to form an MDP framework. Three MDPs were produced: MDP_{nc} (no context nodes), MDP_{period} (period as a context node) and MDP_{score} (match score as a context node). Figure 4.3 depicts the variation in action values for each action across all states and contexts, including outlying values (i.e. those more than 1.5*interquartile range beyond the upper quartile). With the exception of the play-the-ball (PTB) outlier in Figure 4.3B, all outlying state, action pairs were visited on less than 100 occasions, indicating that the models may not have fully converged with the limited dataset available.

It is clear from comparing the box plots in Figure 4.3 that goal kicks are generally much more valuable than any other action. This is not surprising as a goal kick is almost always the last action of a possession, whereas catches, passes and runs happen on a much more frequent basis within possessions and so are likely to receive less value when the rewards obtained are averaged out. However, the greater value provided by goal kicks relative to other actions could skew the player ratings unfairly in favour of the limited number of players who kick at goal. Therefore, two separate action impact ratings are calculated: one including goal kicks and one removing them.




Figure 4.3: Boxplot showing variation in action values for each action across the 19 states with no context nodes included (A), period as context (B) and match score as context (C). Circles depict outliers, calculated as those values more than $1.5 \times$ interquartile range beyond the upper quartile.

4.3.1.1 Action Impact Ratings

To evaluate player performances, action impact ratings adapted from [Routley and Schulte \(2015\)](#) were used. These ratings were calculated using Equation 4.3 and compare the value of the actions taken by a player to the weighted average of all possible actions in the state (i.e. the state value). Table 4.3 provides the top 20 rated players in the Super League 2020 season as rated by MDP_{nc} with goal kicks included. Table 4.5 provides the ratings from the same model with goal kick actions excluded. Both tables provide statistics for tries, try assists, metres gained and goals scored as references for the data currently available to value players in rugby league.

The difference in players' rankings in the rating tables when goal kicks are included or removed shows the disproportionate impact of goal kicks on the results. For example, players 10457, 10625, 10107 and 6335 all fall outside of the top 100 best players when the value of their goal kicks (3 of the 4 players were in the top 5 goal kickers in 2019) is

Table 4.3: Top 20 normalised player action impact ratings for the 2020 Super League season, including goal kicks. Tries, Try Assists, Metres and Goals are provided as references of statistics currently provided for player performances. To protect anonymity, reference statistics are provided based on the whole season as: T-5 (top 5); T-10 (top 10), T-20 (top 20) and 20+ (outside top 20). Players completed a minimum of 150 actions across the season. Player ID brackets provide position in the table when goal kicks are not included in the action impact rating.

Player ID	Position	Action Impact	Tries	Try Assists	Metres	Goals
21721 (1)	Winger	1.259	T-5	20+	20+	20+
2737 (8)	Centre	1.126	T-10	20+	20+	T-10
23232 (2)	Winger	0.980	20+	20+	20+	20+
21082 (3)	Prop	0.824	20+	20+	20+	20+
11352 (4)	Winger	0.794	20+	20+	20+	20+
10665 (5)	Winger	0.655	20+	20+	T-10	20+
10167 (7)	Winger	0.609	T-20	20+	T-10	20+
24400 (6)	Centre	0.608	20+	20+	20+	20+
10457 (110)	Full Back	0.597	20+	20+	20+	20+
21286 (9)	Winger	0.584	T-20	20+	20+	20+
22557 (10)	Winger	0.584	20+	20+	20+	20+
10625 (118)	Scrum Half	0.560	20+	20+	20+	T-5
10107 (125)	Full Back	0.540	T-20	T-5	20+	T-5
20858 (11)	Winger	0.538	T-20	20+	20+	20+
21479 (12)	Winger	0.533	T-5	20+	T-5	20+
21425 (14)	Centre	0.518	20+	20+	20+	20+
3246 (16)	Winger	0.516	20+	20+	T-10	T-20
6335 (116)	Scrum Half	0.495	20+	20+	20+	T-5
20990 (13)	Second Row	0.479	20+	20+	20+	20+
10123 (15)	Winger	0.473	20+	20+	20+	20+

removed from their ratings. However, it should be noted that Player 10107 was nominated for the best player of the year award so it's reasonable to argue that he should be in the top 20 players with or without goal kicks included.

It is notable that the majority of highly rated players when goal kicks were removed were wingers, centres or props. In rugby league, wingers are typically known for scoring tries. They are expected to occupy wide areas to receive the ball and run with it (Table

2.3). Both of these elements are highly valued by the model (Figure 4.3) so it is unsurprising that these players have high action impact ratings. The three centres included within the top 15 players also provide significant support for the model's usefulness as two were included in the Super League 2020 Team of the Year and one was voted as the Young Player of the Year. This is despite two of the centres having relatively poor try, try assist, metre and goal statistics. The high rating of player 21082 is somewhat surprising as player 20567 is widely regarded to be a better player, as shown by his better reference statistics, but can possibly be explained by player 21082's season consisting of 156 actions (only just above the 150 action threshold for inclusion), compared to player 20567's 395 actions from the same position.

4.3.1.2 Contextual Insights

Alongside MDP_{nc} , which evaluated player performances using all data from the 2020 Super League season, MDP_{period} and MDP_{score} were used to assess the context nodes' ability to provide further insights into player performances. The benefit of using these nodes is shown by evaluating the five players nominated for the Man of Steel award, given to the best player in the Super League in 2020. It is noted that none of these players are included in the top 20 best players according to the action impact ratings, which is a limitation of the model. Table 4.7 provides the action impact ratings for the five nominated players across the contextual factors used within this study. It clearly shows differences in the action impact ratings between and within players as the context changes. For example, all players except player 2805 show a reduction in action impact rating in the second half compared to the first. Such a rating provides tactical insight that player 2805 is particularly dangerous in the second half of matches. Similarly, player 2228 and player 10107 are better performers when their team is losing, than when their team is winning. Opposition teams should therefore be aware of this when facing these players' respective teams.

4.3.2 Evaluation of Team Performances

To evaluate team performances, the use of context nodes was extended with MDP_{team} , which used the team ID as the context node to provide individualised team state, action values. The individualised state, action values can be used to gain insight into the value

Table 4.5: Top 20 normalised player action impact ratings for the 2020 Super League season, not including goal kicks. Tries, Try Assists, Metres and Goals are provided as references of statistics currently provided for player performances. To protect anonymity, reference statistics are provided based on the whole season as: T-5 (top 5); T-10 (top 10), T-20 (top 20) and 20+ (outside top 20). Players completed a minimum of 150 actions across the season. Player names are anonymised, brackets provide position in the table when goal kicks are included in the action impact rating

Player ID	Position	Action Impact	Tries	Try Assists	Metres	Goals
21721 (1)	Winger	1.280	T-5	20+	20+	20+
23232 (3)	Winger	0.994	20+	20+	20+	20+
21082 (4)	Prop	0.883	20+	20+	20+	20+
11352 (5)	Winger	0.813	20+	20+	20+	20+
10665 (6)	Winger	0.671	20+	20+	T-10	20+
24400 (8)	Centre	0.645	20+	20+	20+	20+
10167 (7)	Winger	0.625	T-20	20+	T-10	20+
2737 (2)	Centre	0.611	T-10	20+	20+	T-10
21286 (10)	Winger	0.606	T-20	20+	20+	20+
22557 (11)	Winger	0.600	20+	20+	20+	20+
20858 (14)	Winger	0.556	T-20	20+	20+	20+
21479 (15)	Winger	0.552	T-5	20+	T-5	20+
20990 (19)	Second Row	0.550	20+	20+	20+	20+
21425 (16)	Centre	0.542	20+	20+	20+	20+
10123 (20)	Winger	0.502	20+	20+	20+	20+
3246 (17)	Winger	0.499	20+	20+	T-10	T-20
20567 (27)	Prop	0.495	T-10	20+	T-5	20+
20940 (21)	Winger	0.490	20+	20+	20+	20+
2163 (26)	Second Row	0.486	20+	20+	20+	20+
20666 (22)	Winger	0.477	20+	20+	20+	20+

generated by different teams. Figures 4.4(Hull) and 4.5(St Helens) provide two examples of this as the plots provide extremely interesting and contrasting insights regarding the value generated by the two teams. For example, Hull show a clear area of increased value relative to the league average team when performing positional kicks in central areas close to the opposition try line. Furthermore, they generate more value when catching in areas away from the centre of the pitch. Combining these two results suggests that

Table 4.7: Contextual action impact ratings for players nominated for Man of Steel award. Player names are anonymised.

Player Name	1 st Half	2 nd Half	Winning	Drawing	Losing
20441	1.860	1.430	1.330	0.800	1.020
21560	1.260	1.080	1.110	0.434	0.714
10107	1.220	1.040	0.799	0.435	0.939
2228	1.100	0.806	0.430	0.616	0.867
2805	0.875	1.030	0.784	0.591	0.554

Hull may have been particularly adept at kicking from central areas to the outer part of the opposition try line. Conversely, St Helens were generally poor at kicking and catching centrally within 30m of the opposition try line. Instead, they were much better than the league average team at passing and running in areas close to the opposition try line, particularly out wide. That the model can differentiate between these two different styles of play is extremely valuable for a coach devising tactical strategies against future opponents.

The results of the study were presented to the Leeds Beckett University Sports Science Department for feedback. The results of the individualised team, state action values received significant praise in terms of usability and reliability based on the information that was provided and the accuracy and simplicity with which it was depicted. The action impact ratings were treated more sceptically. Although there were some success stories and some of the younger players in the top 20 ratings went on to further success the following year, the feedback from the department was that the action impact ratings may have rewarded attempting to perform a high value action irrespective of its success. It was suggested a better method may be to reward the outcome of the action, rather than its expected value. This may provide a better indication of those players who are better at performing these high value actions successfully.

All analyses for this study were conducted using bespoke Python scripts (Python 3.7, Python Software Foundation, Delaware, USA). The analyses were completed using a HP-Pavilion Laptop with an Intel Core i5-8250U 1.60 GHz processor and 8GB RAM. The Monte Carlo every visit algorithms used to estimate action values were timed at: 15.23s for MDP_{nc} ; 14.06s for $period$; 13.33s for MDP_{score} ; and 11.29s for MDP_{team} .

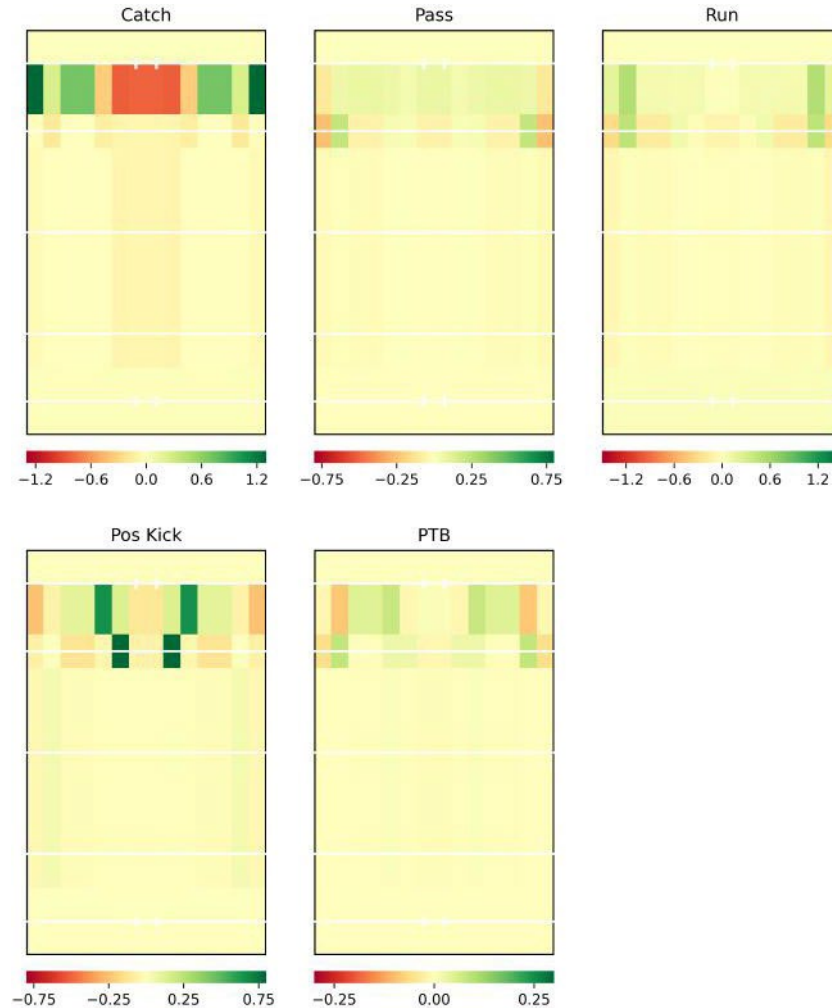


Figure 4.4: Team state, action value plot for Hull. All state, action values are centred by the league average value. Green values indicate team is better at the action in a specific state than the league average team; red indicate the team is worse.

4.4 Summary

This study extended the work conducted in Chapter 3, adapting ideas previously used in ice hockey (Routley and Schulte, 2015) and football (Decroos et al., 2019; Singh, n.d.) to produce a framework through which player and team performances could be evaluated in rugby league using MDPs. It was possible to develop an MDP framework in this study because action by action data became available, rather than the play by play data




Figure 4.5: Team state, action value plot for St Helens. All state, action values are centred by the league average value. Green values indicate team is better at the action in a specific state than the league average team; red indicate the team is worse.

available in Chapter 3. Therefore, it was the first study in rugby league to attempt to evaluate player and team performances using action by action data. Four MDPs were produced: MDP_{nc} , MDP_{period} , MDP_{score} and MDP_{team} . Action impact ratings, previously developed in ice hockey (Routley and Schulte, 2015), were shown to be reasonably successful at rating player performances compared to traditional statistics. Context nodes were shown to provide good insight into player performances when different scenarios (e.g. winning/losing and 1st/2nd half) were considered. The context nodes were also used

to provide individual state, action values for each team providing excellent insights into tactical differences between teams. Practitioners can use the player performance insights for recruitment purposes and the team performance insights to develop tactical strategies against future opponents.

The work was presented at the UKCI conference and subsequently published in *Advances in Computational Intelligence Systems*. It was presented internally to the Sports Science department at Leeds Beckett University, where it was positively received. The team performance insights were highly praised in terms of both usability and reliability. However, the player player performance insights' reliability was questioned as it was suggested that the action impact rating provided slightly indifferent results because it rewarded attempting high value actions rather than considering the outcome of the actions. To remedy this, the department suggested that a future performance rating should consider some form of actual vs expected possession outcome measure.

Given its success at a team level, future studies may now wish to apply these methods to other sports. Doing so would require a two-step process: first, the most appropriate zone size should be chosen based on a criterion measure appropriate for the sport, similar to Chapter3; second, a useful set of actions should be included so that their values can be considered. The methods considered in both chapters should be useful in any sport which takes place on a pitch or court. For example, similar methods have already been considered in football (Singh,n.d.) and ice hockey (Routley and Schulte,2015). There is also scope for the model to be used in individual sports, such as tennis, where the probability of winning a point based on performing different shots from different locations could be considered.

In addition to the action impact rating critique, there are a number of reasons why the player performance ratings may have provided mixed results. The first is that all defensive actions were removed from the model, which may have stopped players who were good defenders from being rated highly. The second is that the model values point scoring actions (goal kicks and runs) disproportionately higher than all other actions, which unduly rewards those players who attempt these actions more regularly. Finally, it could be that the zones used in this study are too large for meaningful differences to be identified. This limitation could also relate to the team performances as no major differences were identified between any team and the league average model when they were more than 30m away from the opposition try line. A likely reason for this is the size

of the zones - the middle zone is 60m long, with every location in the zone receiving the same value, which means a player or team can progress 59m closer to the opposition try line and receive no value for it under the current framework. Such a situation is obviously not realistic and so a method which produces a smoother surface than the dichotomous zoning approach currently used in rugby league is required.

As a result of these limitations, it is necessary to move away from the zonal approaches considered in the previous two chapters as it is clear that it would prove difficult to provide results that could be fully validated as usable and reliable by practitioners. Instead, the next study will attempt to provide a smooth surface. Doing so will mean that an MDP approach is no longer viable. Therefore, as a prototype for further analysis, the research question will be simplified to consider the probability a team will control the ball in a given location, rather than attempting to estimate the expected value conditional on the location or location, action tuple.

Kernel Density Estimation and Wasserstein Distance Evaluation of the Spatial Trends of Team Attacking Performances

5.1 Introduction

In the previous study (Chapter4), the limitations of using a zonal approach to evaluate player and team performances in rugby league were highlighted. The most prominent of these issues was the size of the zones, which meant a player could run more than 50m towards the opposition try line but not receive a positive rating if he remained within the same zone. This issue was not prevalent in basketball or football, as their zones were much smaller (Cervone et al.,2016;Singh,n.d.), but significantly reduces the validity of the results in rugby league. A logical solution to the zone size problem is to try and calculate a smooth surface, as has been achieved in previous deep learning (Fernández et al.,2021) and deep reinforcement learning models (Liu et al.,2020). Unfortunately, it is not possible to produce such complicated models in rugby league due to data availability issues. Therefore, this study takes the first step towards analysing player and team performances in a smooth manner by simplifying the problem of evaluating team performances. Rather than considering the value of an action conditional on its location, this study considers

the probability of a team controlling the ball in a given location. The spatial trends of team attacking performances are quantified by adapting the methods of Mallepalle et al. (2020) in rugby league. The work of Mallepalle et al. (2020) is then extended through the proposition of novel metrics, which evaluate differences in spatial trends of attacking performances between and within teams. The study was presented at the UKCI 2022 conference (Thomas Sawczuk, Anna Palczewska, Ben Jones, Jan Palczewski. "Use of Kernel Density Estimation to understand the spatial trends of attacking possessions in rugby league" in *21st UK Workshop on Computational Intelligence*, Sheffield, 2022) and will be published in *Advances in Computational Intelligence Systems*. It was also presented to coaches at Leeds Rhinos, a professional rugby league team competing in the Super League, who validated the reliability and usability of the results.

5.2 Methodology

In this section, the methodology for the study is described. It describes the data and data preprocessing steps required to prepare the data for analysis. Kernel Density Estimation and the Wasserstein Distance are presented and their adaptation to rugby league is outlined. Novel metrics evaluating the spatial trends of team attacking performances in this study (normalised axis Wasserstein distance and directional Wasserstein distance) are formalised.

5.2.1 Data

In this study, event level match-play data were obtained from Opta (Stats Perform, London, UK) for the 2021 Super League season. The data was produced via human annotation of the actions taking place and was downloaded from www.optaprorugby.com. In 2021, Opta again improved their coding protocols. Deviating from their previous methods, the company provided 20 categories of actions, 96 actions and 446 additional action descriptors, describing the type of action and its outcome. Table 5.1 provides an example of the actions and descriptors used for the passing category. Appendix C is the user manual provided by Opta, where definitions of the actions and their coding criteria can be viewed. The data for each match was provided in separate xml files, but the play-the-ball node was dropped in 2021, so the data was provided using three nodes: match data; team data;

Table 5.1: Passing example to show the action definitions utilised by Opta in 2021

Level	Options
Category	Passing
Action	Short; Out the back; Face ball; Offload; Dummy Pass
Action Descriptor	Effective; Ineffective; Off Target; Turnover; Left; Right

and fixture data. The definitions of the three nodes remained the same as previous years (Table 3.1).

In Chapter 3 the data provided by Opta were introduced. As with Chapters 3 and 4, 13 of the 28 variables provided by Opta were considered within this study. These were: ID, FXID, PLID, team id, MatchTime, x coord, y coord, action, ActionType, Actionresult, Metres, PlayNum and SetNum. Table 3.2 defines these variables. Similar to the 2020 dataset used in Chapter 4, location data was available on an action by action basis. Contrary to the 2020 dataset, a much greater number of actions were coded across the season, including supporting runs, a greater variety of defensive actions and more miscellaneous data.

In the 2021 Super League season, 138 matches were contested by 12 teams and 373 players. A total of 557,050 match actions were completed (median 4003 events per match, interquartile range 3857-4200). Across the season, 1001 tries were scored (768 successful conversion kicks, 233 unsuccessful conversion kicks), 175 penalty goals were attempted (158 successful, 17 unsuccessful) and 83 drop goals were attempted (37 successful, 46 unsuccessful).

5.2.2 Data Preprocessing

Figure 5.1 provides an overview of the data preprocessing completed within this study. The preprocessing steps converted the 138 raw xml files described above to a single file of 99,966 rows and 4 columns ready for analysis. Full details of each preprocessing step are provided below.

The match action data, provided in 138 xml files, were collected and integrated into a single dataset of 557,050 rows and 28 columns. To do this, the match data node variables were extracted from each xml file into separate tables and then concatenated together.

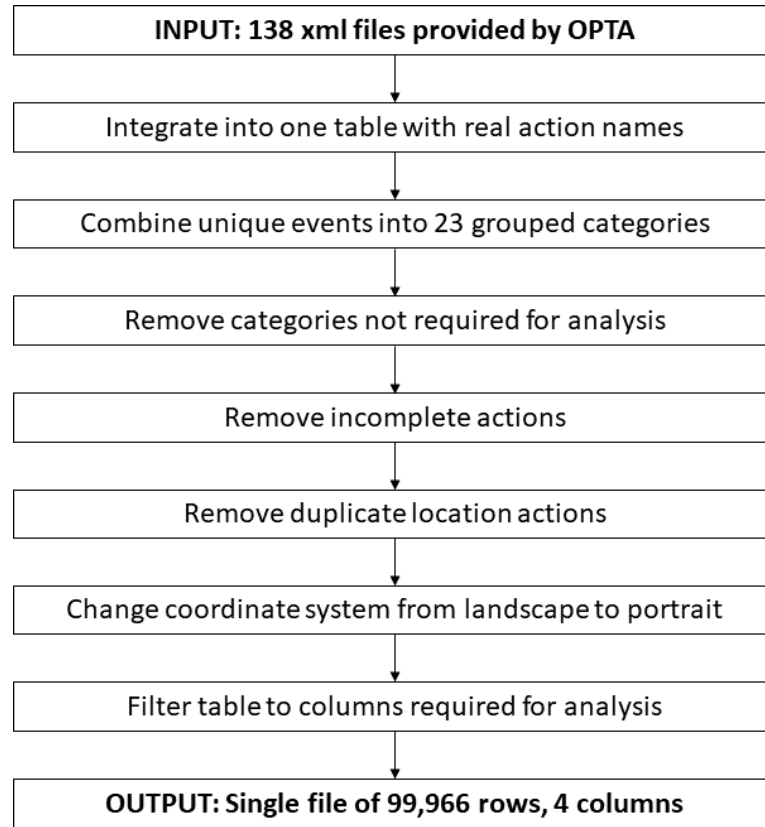


Figure 5.1: Data preprocessing workflow

The team data nodes were also extracted into tables for each fixture. These were used to map team unique identifiers to the real team name. Similarly, Opta's action definitions and unique identifiers (Appendix C) were mapped to the real name of the coded actions.

Within the 2021 dataset, 126 unique actions (i.e. different combinations of actions and action descriptors) were coded. For the purposes of this study, these actions were regrouped into 23 preprocessing categories. These preprocessing categories were used to assist with the preprocessing steps, but were not required for the location based analysis described in Section 5.2.3. A full list of the actions included in each preprocessing category is provided in Table 5.2.2.

Table 5.2: Preprocessing categories and actions included within them in this study

Category	Events
Auxiliary Information	Front Marker, Back Marker, Video Ref, Interchange, HIA, Stoppage
Generic Descriptor	Other Error, Try cause, Line Break Involvement, Line Break Assist, Break Cause, Tackle Break, Opp Error, Passing Move, Close Range, Error, Try Involvement, Long Range, Individual Effort, Other, Sin Bin Out, Yellow, Sin Bin Return To Field, Grounding, Contest, Sent Off Out, Red, Touchline/Deadball, Onside, On Report
Restart Actions	50m Restart, Goal Line Drop Out, 20m Restart
Move Self	Restart Run, Evasion, Hitup, Kick Return, Line Engaged, Ruck Run, Dummy Half, Run, Line Not Engaged
Move Team	Complete, Short - Crossfield, Short - Grubber, Own Player, Break, Short - Banana, To Ground, Short - Bomb, Short - Chip, Try
Kick Goal	Conversion, Penalty Goal, Field Goal
Kick Position	Long - To Opposition, Good, Long - To Open, Long - 40-20, Long - Touch
Catch Pass	Simple Receipt, Jump Catch
Catch Kick	Kick Receipt, Restart Receipt
Loose Ball	Defensive Cleanup, Attacking Cleanup, Attempted Intercept, Contestable Cleanup, Interception
Tackle	Made, Dominant, Offload To Ground, Turnover Ball Split, Forced Within In Goal, Stolen, Turnover Into Touch, Offload
Missed Tackle	Bumped Off, Stepped, Positional, Outpaced, Try Conceded
Run Action	Dummy Pass, Half Break, Line Break, Carried Dead Ball, Forced Into In Goal, Carried In Touch

Play-The-Ball	Lost, Won, Interrupted
Attacking Descriptor	Kick Line Break, Try Assist, From Kick, From Penalty, From Line
Defensive Descriptor	Kick Pressure
Move Self Error	Dropped Ball Unforced, Ball Jolted, Lost Ball Forced
Move Team Error	Not Out, Failure To Find Touch, Incomplete, Off Target, Forward Pass, Forward, Kick Error, Bad Offload, To Opposition, Bad Pass, PTB Fumble, Intercepted
Catch Error	Accidental Knock On, Falcon
Penalty Conceded	Defence, Penalty, Inside 10m, Attack, Ruck Infringement, Foul Play, Double Movement, Obstruction
Defensive Play	Flop, Kick Not Defused, Kick Defused, Charge Down, Attempted Steal
Off The Ball	Decoy, Support Run, Kick Shield, Kick Shepherd, Kick Chase

Although the type of action completed was not in itself important for this study, the information it provided was important for the purposes of preprocessing. To prepare the data for the analysis of the spatial trends of team attacking performances, only actions completed by the team in possession of the ball were required. In order to prevent multiple codings being present for the same action, a subset of the action categories described in Table 5.2.2 was considered: “move team”; “move self”; “catch kick”; “kick position”; “move team error”; “move self error”; “loose ball” and “kick goal”. This filtering was necessary as a player completing a “move self” action could also be given a “run action” descriptor if they ran past an opponent, thus providing two entries for the same action. Filtering the actions in this manner immediately reduced the dataset from 557,050 observations to 128,716 observations. Incomplete actions were only included if they caused the end of a possession. For example, an unsuccessful attempt at a loose ball collection or an unsuccessful attempt to intercept a pass by the defensive team were not included as these actions were not completed and including them would have unduly affected the chain of possession. However, an unsuccessful pass, or a dropped catch from a pass by the attacking team were included as these actions resulted in a change in the chain




Figure 5.2: Location of 99,966 actions used in this study. Dots representative of location only; quantity of actions at each location is not displayed in this plot.

of possession. Incomplete actions were identified if the `team_id` variable changed for the subsequent action. Their removal reduced the size of the dataset by 1122 actions to 127,594 observations.

Similar to multiple action codings, it was important for this location based analysis that locations were not coded on more than one occasion for different actions completed within the same sequence. For example, if a player caught the ball and decided to run in the same location, Opta would code each action as an individual observation. Multiple location codings were identified when the location and player variables were identical on consecutive rows. For this study, because the action chosen was not important to the location based analysis, the second observation of each multiple location coding was removed from the dataset indiscriminately. Removing these duplicate location actions resulted in a final dataset size of 99,966 actions.

Unlike Chapters 3 and 4, which used zone based approaches, the `x, y` coordinates

Table 5.3: Sample possession within the dataset

Attacking Team	Defending Team	x	y
St Helens	Salford	9	4
St Helens	Salford	9	6
St Helens	Salford	14	11
St Helens	Salford	22	13
St Helens	Salford	12	12
St Helens	Salford	37	16
St Helens	Salford	36	24
St Helens	Salford	54	35

Table 5.4: Descriptive data for action counts in the data subsets. IQR is interquartile range; Min/Max refers to minimum/maximum number of observations observed.

Level	Subsets	Median	IQR	Min/Max
Whole league	1	99,966		
Team overall	12	8105	7596-8937	7203/10324
Team-opponent	132	732	622-858	277/1551

provided by Opta were directly modelled in this study. In 2021, the x,y location was denoted relative to a standardised 100m x 70m rugby league pitch, rather than the 100m x 68m pitch used in Chapters 3 and 4. Actions occurring outside the 0-100m pitch are located in the possession team's try area (0 to -10m) or the opposition team's try area (100 to 110m). Figure 5.2 shows the locations of the 99,966 actions used in this study.

After completing the preprocessing steps, the dataset was filtered to include only the columns required for analysis. As such, the 138 individual xml files were preprocessed into a single dataframe of 99,966 rows and 4 columns. Table 5.3 provides a sample of the data used for analysis. The inclusion of the attacking and defending teams in the final dataset was important to allow for comparisons in the spatial trends of attacking performances between and within teams to be made.

Three levels of data were used in this study. The whole league data used the complete set of data. The team overall data used a subset of data for a single team controlling the ball against any opponent (i.e. using all data for the attacking team). The team-opponent

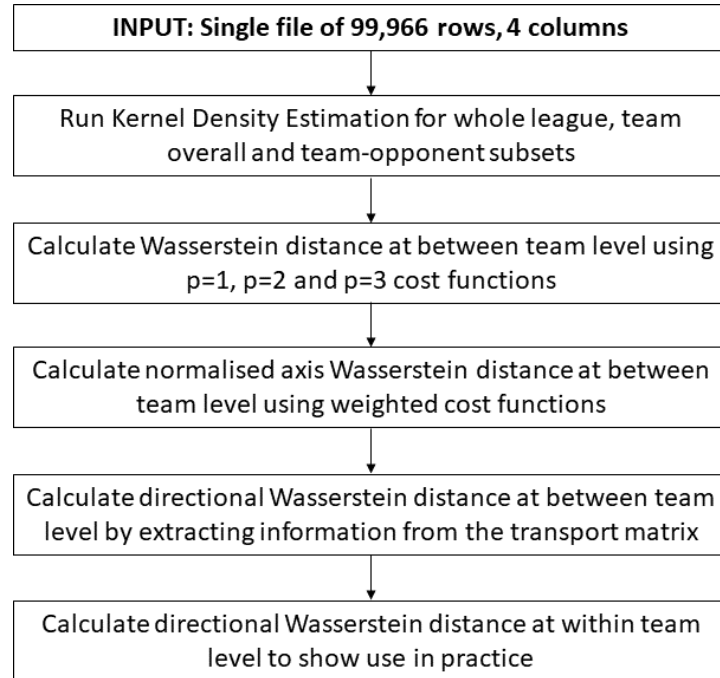


Figure 5.3: Workflow of analyses conducted within this study

data used a subset of the data, which contained actions for a single team controlling the ball (i.e. attacking) against a specific opponent (identified using the attacking team and defending team columns). There were 12 team overall data subsets (one for each team) and 132 team-opponent data subsets (one for each team-opponent combination). Table 5.4 provides descriptive data for the three levels of analysis. It shows that the greater the number of subsets produced, the smaller the number of actions included. Furthermore, there is greater variation in the size of the subsets when more are used, which can be attributed to the varied number of matches played against different opponents in the 2021 season.

5.2.3 Kernel Density Estimation

The aim of this study was to present a method of evaluating the differences in spatial trends of team attacking performances in rugby league by adapting, validating and extending the methods used by Mallepalle et al.(2020) in American football. Figure 5.3

provides a schematic of the analyses conducted in this study.

Kernel Density Estimation (KDE), described in Definition 5.2.1, provides a method through which data can be smoothed using kernel smoothing. It allows inferences about the population to be made based on a finite data sample. KDEs were used by Mallepalle et al.(2020) to quantify pass probability distributions for quarterbacks in American football. In this study, KDEs are used to quantify the spatial trends of team attacking performances.

Definition 5.2.1. Kernel Density Estimation (KDE) is a non-parametric method of estimating the unknown probability density function of a dataset (Rosenblatt,1956). The bivariate KDE (\hat{f}_H) of a sample of 2-dimensional vectors $(\vec{x}_i)_{i=1}^n$ is defined as

$$\hat{f}_H(\vec{x}) = \frac{1}{n} \sum_{i=1}^n K_H(\vec{x} - \vec{x}_i), \quad (5.1)$$

where K_H is the kernel with a 2×2 smoothing matrix H and $(\vec{x}_i)_{i=1}^n$ are vectors containing the x and y co-ordinates of action locations.

Multiple kernel options are available for KDEs including: uniform, triangular, normal, Epanechnikov and cosine Wand and Jones(1995). The standard bivariate normal kernel is used in this study:

$$K_h(\vec{x}) = (2\pi h)^{-1} e^{-\frac{\|\vec{x}\|^2}{2h}}, \quad (5.2)$$

where $\|\vec{x}\|$ is the Euclidean norm of \vec{x} . This corresponds to a diagonal smoothing matrix with the same smoothing constant h (the bandwidth) in both directions.

Previous studies have considered other approaches to the smoothing matrix, including adaptive values (Fleming et al.,2015), mixture values (Lichman and Smyth,2014) and self selected values based on neighbourhood service radii (King et al.,2016). In this exploratory study, a diagonal smoothing matrix was deemed to be most appropriate as it assumed that a location has an equal influence on all locations surrounding it. Given previous zonal approaches in sport assume that all locations within the zone are identical (Chapters 3 and 4; Cervone et al.(2016); Kempton et al.(2016)), it is a logical next step to assume that a specific location has equal influence on all surrounding locations.

The most important free parameter within KDE is the bandwidth (Chen,2015). The bandwidth parameter influences the smoothness of the KDE model and controls “over-fitting” by establishing the amount of data smoothed at every point. Smaller bandwidth

values result in a more jagged appearance with larger peaks and troughs, whereas larger values result in a much smoother appearance with smaller peaks. Typically, as the size of the dataset increases, the optimal bandwidth size reduces (Zambom and Dias, 2012). Consequently, bandwidths can be calculated individually for each subset of data to ensure that the best possible fit to each subset's data is obtained. This method was used by Mallepalle et al. (2020) when evaluating quarterback pass probability distributions. However, individually calculating bandwidths can cause issues when directly comparing differences between KDEs as the use of different bandwidths can result in artificial differences being observed (Bowman, 1984). Therefore, a single bandwidth was used across all 145 data subsets in this study so that team KDEs could be compared directly.

Given its importance to the estimates provided by KDE, numerous methods have been identified through which the optimal bandwidth can be selected, including: visual inspection; fitting to a reference distribution (Scott, 2015; Silverman, 1986); estimation to minimise the mean integrated square error (Park and Marron, 1990; Sheather and Jones, 1991); and cross validation (Bowman, 1984; Marron, 1989; Stone, 1984). Contrary to the assumptions underlying some of these approaches, the rugby league data used in this study is auto-correlated as locations of consecutive actions in a play are not independent. The effect of auto-correlation on optimal bandwidth selection is shown by Fleming et al. (2015) in their framework of minimising the mean integrated square error for a Gaussian reference distribution. Their method performs particularly well in the study of highly correlated animal location data. However, the stretches of correlated locations in this rugby league dataset are relatively short and of varying length, so the approach of Fleming et al. (2015) with a reference Gaussian distribution would not offer an improvement over a cross validation approach which is distribution-free. Furthermore, it has been suggested that a cross validation approach is more responsive to the different sample sizes encountered within this study (Table 5.4; Zambom and Dias (2012)). As such, a likelihood cross validation approach to bandwidth selection was utilised. Denoting by X_j the data left out in j -th fold, the likelihood cross-validation (CV) for a bandwidth h is calculated by

$$CV(h) = j^{-1} \sum_{j=1}^j \log \hat{f}_{h,-j}(\vec{x}), \quad (5.3)$$

where j refers to the number of folds (10 in this case) and $\hat{f}_{h,-j}$ is the kernel density

estimate using bandwidth h and without data X_j .

The bandwidth for the KDEs was chosen via a two stage process, similar to a previous study comparing the length frequencies of fish species (Langlois et al.,2012). First, the optimal bandwidths for the 132 subsets on the team-opponent level were identified using 10-fold cross validation. The team-opponent level was chosen for this stage because it represented the lowest level of data and would therefore provide a bandwidth which could be used to generate smooth results for all higher levels of data. If the top level of data (i.e., the whole league) had been used, the bandwidth obtained would have been too small for other levels to obtain usable results (Zambom and Dias,2012). Second, the geometric mean of the 132 team-opponent bandwidths was calculated; this resulted in a bandwidth equal to 4.10. A bandwidth of 4.10 suggests that locations are predominantly influenced by areas within 10m of them, which is congruent with the zone sizes initially suggested in the rugby league literature (Kempton et al.,2016). All 145 KDEs were re-run at this empirically identified bandwidth 4.10, ensuring that the under and overfitting issues which could be present due to using the same bandwidth for all analyses were minimised across all subsets of data (Bowman,1984).

5.2.4 Wasserstein Distance

To extend the work ofMallepalle et al.(2020), two novel metrics of quantifying differences in the spatial trends of attacking performances between and within teams were produced. In Section3.2.4.1, the differences in the discrete EPV distributions for a set of previous matches and the subsequent match were calculated using the KL Divergence. By considering the x, y pitch coordinates in a smooth manner in this study, it is now possible to use the Wasserstein distance (Dobrushin,1970;Panaretos and Zemel,2019) to quantify differences in the spatial trends of attacking performances between and within teams.

The Wasserstein distance is a more appropriate measure than the KL Divergence because it considers the size *and* distance of differences between the two distributions, whereas the KL Divergence only considers the size of the differences. In practice, this means that if an equal mass difference is distanced 10m or 100m apart, the KL Divergence will return the same value. Conversely, a much larger value is provided by the Wasserstein distance for the 100m movement of mass compared to the 10m movement. The difference in the Wasserstein distance for these two mass movements is dictated by

the cost function.

The p-Wasserstein distance (W_p) between two distributions μ and ν is calculated as

$$W_p(\mu, \nu) = \left(\inf_{\gamma \in \Gamma(\mu, \nu)} \int_{M \times M} \|\mathbf{z} - \mathbf{z}'\|_p^p \gamma(d\mathbf{z}, d\mathbf{z}') \right)^{\frac{1}{p}}, \quad (5.4)$$

where M is a 2-dimensional space of coordinates on the pitch, $p \geq 1$, $\|(x, y)\|_p = (|x|^p + |y|^p)^{1/p}$ is the L_p norm¹ (the cost function) and $\Gamma(\mu, \nu)$ is the set of all couplings of distributions μ and ν , i.e., the set of all joint probability distributions on $M \times M$ with marginals μ and ν .

The p-Wasserstein distance is closely related to optimal transport planning (Villani, 2003). The optimal transport plan describes the movement of mass from distribution μ to distribution ν with minimal cost, subject to a cost function, $c(\mathbf{z}, \mathbf{z}') : 1 \rightarrow [0, \infty)$. The W_p Wasserstein distance is the total cost of all mass movements described by this transport plan, subject to the L_p norm cost function. This relationship between the Wasserstein distance, cost function and transport plan means that axis (e.g. x , y) and directional (e.g. left, right and up, down) mass transport measurements could be provided by manipulating the cost function or extracting information from the transport plan; with a slight abuse of terminology, these measurements will be termed Wasserstein distances in this study. These variations of the Wasserstein distance could provide much greater insights into the spatial trends of attacking performances in rugby league than using a standard L_p norm cost function.

To evaluate differences in the spatial trends of attacking performances between teams, team overall KDEs were compared to the whole league KDE. This analysis evaluated how similar or dissimilar the team's overall distribution of actions was to the league average distribution. Larger Wasserstein distance values denoted larger differences, either due to the size of the difference, the distance involved in the difference or both factors. The within team analysis was completed by comparing team-opponent KDEs to the team overall KDEs (e.g. Wigan vs St Helens would be compared to Wigan's team overall KDE; Catalans vs Leeds would be compared to Catalans' team overall KDE). This analysis evaluated the similarity between the distribution for the attacking team against a single opponent and the distribution for the attacking team against all opponents.

As with KDE bandwidth selection in Section 5.2.3, an empirical approach was em-

¹ L_1 norm is sometimes called Manhattan norm and L_2 norm is the Euclidean norm.




Figure 5.4: Standardised, unweighted Wasserstein metrics for $p = 1$ (left), $p = 2$ (centre) and $p = 3$ (right). Distances are calculated from central point.

ployed to select the Wasserstein distance cost function. Three cost functions ($p = 1$, $p = 2$ and $p = 3$ in Equation 5.4; Figure 5.4) were evaluated at the between team level (i.e. using the team overall KDEs), providing W_1 , W_2 and W_3 values. The cost of moving one unit in the horizontal and vertical directions is identical for all three cost functions but the diagonal cost differs based on the p -norm used. The diagonal cost of mass movement, imposed by the L_p norm cost function, was greater in W_1 than W_2 (by the exponent 2) and W_3 (by the exponent 3). The between team analysis compared the mathematical properties of each cost function to the visual differences identified in the KDEs and was used to identify the most appropriate cost function for use in rugby league from those considered. From this analysis, the Minkowski distance ($p = 3$) cost function was selected and used for all further analyses in the study.

5.2.4.1 Normalised Axis Wasserstein Distance

To understand axis level differences in the spatial trends of team attacking performances, the cost function was manipulated and compared to the results of the Wasserstein distance calculated using the unweighted $p = 3$ cost function (W_3 , described above). Equal and unequal weightings of the W_3 cost function were considered (Table 5.5). Equal weightings multiplied the x and y components by an equivalent factor (2 and 3 in this study) to identify whether further insights could be gleaned by increasing the cost of movement along both axes. Unequal weightings were considered due to the non-square 70m x 120m dimensions of a rugby league pitch. With an equally weighted cost function, it was hypothesized that movements along the y -axis could have an unfairly large influence on the Wasserstein distance as the cost of moving the full distance of the y -axis (120) was greater

than the cost of moving the full distance of the x-axis (70). Using unequal weightings, the cost of moving a single unit along the x axis was manipulated to be greater than the y axis by a factor of 12/7. This normalised the cost of moving the complete length of the x and y axes. Four variations of the unequal weighting were used: 2.4x, 1.4y; 3.6x, 2.1y; 4.8x, 2.8y; and 6.0x, 3.5y. It was hypothesized that levelling the cost of movement in this manner would enhance the insights provided at an axis level.

For all weighted cost functions, the Wasserstein distance was calculated for the x and y axes independently. To achieve this, the unit cost of moving along the x or y axis was multiplied by the given factor, while the other axis maintained the standardised $p = 3$ movement cost of one unit (e.g. for the equally weighted 2x, 2y analysis, one distance measure was produced for the 2x, 1y weighted $p = 3$ cost function and one was produced for the 1x, 2y weighted $p = 3$ cost function). The increased movement cost for the factored axis ensured that the optimisation procedure only included the most important mass movements along the factored axis when calculating the optimal transport plan and allowed the influence of the weight on each axis to be considered individually. This process ensured that three Wasserstein distance values were obtained for each comparison: one using the standardised, unweighted $p = 3$ cost function; one weighted with an increased x-factor; and one weighted with an increased y-factor. The importance of mass movement along the factored axis to the Wasserstein distance was evaluated by comparing the normalised axis Wasserstein distance to the standard Wasserstein distance. The normalised axis Wasserstein distance \bar{W}_p^a was calculated as

$$\bar{W}_p^a = \frac{W_p^a - W_p}{W_p}, \quad (5.5)$$

where W_p^a is the Wasserstein distance weighted by $a \in \{x, y\}$ and W_p is the standard Wasserstein distance for norm p .

The comparison of normalised axis Wasserstein distance measures was completed at the between team level. \bar{W}_p^a distances represent the importance of density movements along axis a relative to W_p . A larger value indicates a greater cost of density movement along axis a relative to the unit density movement cost. Quadrant plots were used to evaluate the ability of \bar{W}_p^a measures to identify differences between teams.

Table 5.5: Weighted Wasserstein distance metrics considered in this study

x-factor	y-factor
2.0	2.0
3.0	3.0
2.4	1.4
3.6	2.1
4.8	2.8
6.0	3.5

5.2.4.2 Directional Wasserstein Distance

To provide directional insights with respect to the spatial trends of attacking performances between teams (i.e. understanding left/right directional differences compared to x or y axis differences), the directional Wasserstein distance was calculated using the transport matrix from the optimal transport plan. The transport matrix is computed for a fine discretisation $(x_k, y_k)_{k=1}^K$ of the field. For each pair of points (x_k, y_k) and $(x_{k'}, y_{k'})$, it specifies how much mass is moved from (x_k, y_k) to $(x_{k'}, y_{k'})$, denoted by $T((x_k, y_k); (x_{k'}, y_{k'}))$. As there is a different solution to the transport matrix for every cost function employed, the unweighted Minkowski distance ($p = 3$) cost function was used to calculate the transport matrix. The transport matrix was used to obtain directional Wasserstein distances for left (W_{left} ; Equation 5.6), right (W_{right} , Equation 5.6), up (W_{up} , Equation 5.8) and down (W_{down} , Equation 5.9):

$$W_{\text{left}} = \sum_{k=1}^K \sum_{k'=1}^K 1_{x_{k'} < x_k} T((x_k, y_k); (x_{k'}, y_{k'})) |x_k - x_{k'}|, \quad (5.6)$$

where $1_{x_{k'} < x_k}$ equals one if $x_{k'} < x_k$ and zero otherwise.

$$W_{\text{right}} = \sum_{k=1}^K \sum_{k'=1}^K 1_{x_{k'} > x_k} T((x_k, y_k); (x_{k'}, y_{k'})) |x_k - x_{k'}|, \quad (5.7)$$

where $1_{x_{k'} > x_k}$ equals one if $x_{k'} > x_k$ and zero otherwise.

$$W_{up} = \sum_{k=1}^K \sum_{k'=1}^K 1_{y_{k'} > y_k} T \left((x_k, y_k); (x_{k'}, y_{k'}) \mid y_k - y_{k'} \right), \quad (5.8)$$

where $1_{y_{k'} > y_k}$ equals one if $y_{k'} > y_k$ and zero otherwise.

$$W_{down} = \sum_{k=1}^K \sum_{k'=1}^K 1_{y_{k'} < y_k} T \left((x_k, y_k); (x_{k'}, y_{k'}) \mid y_k - y_{k'} \right), \quad (5.9)$$

where $1_{y_{k'} < y_k}$ equals one if $y_{k'} < y_k$ and zero otherwise.

The difference between W_{left} and W_{right} was calculated to provide directional insights along the x axis. Positive values were indicative of a greater mass movement towards the left side of the pitch or away from the right side of the pitch, negative values indicated greater mass movement towards the right side of the pitch or away from the left side of the pitch. The difference between W_{up} and W_{down} was calculated to provide directional insights along the y axis. Positive values were indicative of a greater mass movement towards the opponent's try area or away from the possession team's try area, negative values indicated greater mass movement towards the possession team's try area or away from the opponent's try area. Quadrant plots were used to compare the directional Wasserstein distances to the W_3 distance.

After comparing the normalised axis Wasserstein distance (Equation 5.5) and the directional Wasserstein distance (Equations 5.6 to 5.9), with respect to the insights they provide regarding differences in the spatial trends of team attacking performances at a between team level, a within team analysis was performed for two clubs. This analysis used the directional Wasserstein distances described above to show how a performance analyst may use quadrant plots to evaluate the spatial trends of attacking performances for a given team when facing different opponents. Further, it investigated the relationship between these differences and the points scored in matches against different opponents.

5.3 Results

The aim of this study was to adapt, validate and extend previous work by Mallepalle et al. (2020) by evaluating the spatial trends of team attacking performances in rugby league,

both between and within teams. To achieve this aim, 138 xml files provided by Opta were preprocessed into a single file of 99,966 rows and 4 columns (Table 5.3). From this data, 145 KDEs were produced using the whole league data (1), a subset of each team's overall data (12) and a subset of each team-opponent's data (132), using the same 4.10 empirically calculated bandwidth. Next, the differences in these KDE distributions were evaluated between and within teams using two novel Wasserstein distance metrics. At a between team level, the ability to generate insights regarding the axis (via the normalised axis Wasserstein distance) and directional (via the directional Wasserstein distance) differences in spatial trends of team attacking performances was discussed. At the within team level, quadrant plots were used to show how a performance analyst could use the directional Wasserstein distance to gain insights into the directional differences of spatial trends of attacking performances for a single team against different opponents.

5.3.1 Quantification of the Spatial Trends of Team Attacking Performances

To quantify the spatial trends of team attacking performances in this study KDEs were used, adapting previous work in American football (Mallepalle et al., 2020). KDEs were produced at the whole league, team overall and team-opponent levels. All KDEs were produced using the empirically calculated bandwidth 4.10 in the x and y dimensions as outlined in Section 5.2.3. Figure 5.5 depicts the KDE for the whole league dataset, which was the reference distribution for the between team analysis. Figure 5.6 shows the differences between each team's overall KDE and the whole league KDE. Visual inspection of Figure 5.6 shows some teams have clear areas of greater (shown in green) or lower (shown in red) densities relative to the whole league KDE. For example, Castleford (left side) and Huddersfield (central) have strong biases in the x or horizontal direction, whereas weaker horizontal biases are present for Hull Kingston Rovers (left side of pitch) and Wakefield (not left side of pitch). Leeds (opposition half), Salford (anywhere but opposition 20m) and St Helens (anywhere but either 20m) have strong biases in the y or vertical direction. Visualising these trends allows coaches and performance analysts to see clearly the locations where teams are more or less likely to perform actions across the whole season.




Figure 5.5: Kernel Density Estimation plot for whole league data at bandwidth 4.10

5.3.2 Evaluation of Between and Within Team Differences

To evaluate the between and within team differences in the spatial trends of attacking performances quantitatively, two novel Wasserstein distance metrics were produced. The analyses using the Wasserstein distance measures were conducted in two phases. First, the between team analysis was used to identify the most appropriate p-norm for use within rugby league and compare the insights gained using the two novel Wasserstein distance metrics (normalised axis Wasserstein distance and directional Wasserstein distance). Second, the results were considered at a within team level to understand how coaches and performance analysts may use them in practice.

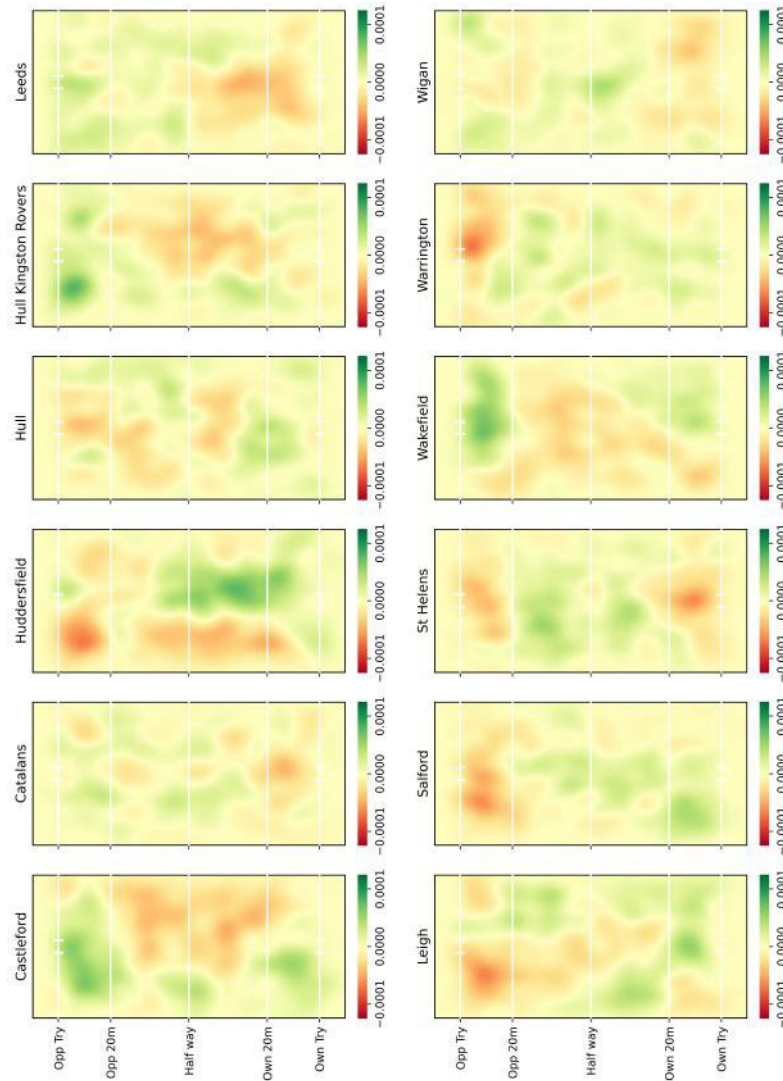


Figure 5.6: Comparison between whole league attacking KDE and each team's overall KDE. Green areas are areas represent areas where the team has a higher density than the whole league, red areas have lower densities.

5.3.2.1 Identification of Suitable p-norm

The first element of the between team analysis was to identify the most appropriate p-norm for use within rugby league (Equation 5.4). Table 5.6 shows the W_1 , W_2 and W_3 Wasserstein distances at the between team level. These represent p-norms of 1, 2 and 3 respectively. There was a general trend for the Wasserstein distance to decrease as the

Table 5.6: Wasserstein distances for W_1 , W_2 and W_3 at the between team level. Distances represent differences between the team's overall KDE and the whole league KDE. Final column ($W_1 - W_3$) is the difference between the W_1 and W_3 distances

Team	W_1	W_2	W_3	$W_1 - W_3$
Castleford	3.07	2.23	2.03	1.04
Catalans	1.08	0.96	0.94	0.14
Huddersfield	3.38	2.57	2.38	1.00
Hull	1.38	1.18	1.13	0.25
Hull KR	2.12	1.64	1.55	0.57
Leeds	2.27	1.99	1.96	0.31
Leigh	2.61	2.13	2.02	0.59
Salford	2.36	2.23	2.22	0.14
St Helens	1.60	1.32	1.28	0.32
Wakefield	1.85	1.48	1.41	0.44
Warrington	1.72	1.41	1.36	0.36
Wigan	1.34	1.17	1.14	0.20

p-norm increased. However, this difference did not appear to be proportionate between teams as Castleford (absolute difference between W_1 and $W_3 = 1.04$), Huddersfield (1.00), Leigh (0.59), Hull KR (0.57) and Wakefield (0.44) were affected to a much greater extent than Salford (0.14) and Leeds (0.31). Given Salford and Leeds both had higher W_1 distances but lower $W_1 - W_3$ differences than Hull KR and Wakefield, the variation in $W_1 - W_3$ differences is unlikely to be related to the initial W_1 distance.

Visual inspection of Figure 5.6 suggests that the difference in Wasserstein distances could be related to the location of the spatial trends of attacking performances for each team. Castleford (left), Huddersfield (central), Hull KR (left) and Wakefield (not left) showed some form of bias along the x axis, whereas Leeds (opposition half) and Salford (anywhere but opposition 20m) showed biases along the y axis. Given the equal horizontal and vertical costs of mass movement between locations for W_1 , W_2 and W_3 , the difference in Wasserstein distances must be related to the diagonal cost of mass movement defined by the cost function. The diagonal cost of mass movement was greater in W_1 than W_2 (by the exponent 2) and W_3 (by the exponent 3). These differences are shown by the lighter (greater cost) and darker (reduced cost) areas in Figure 5.4. In order to generate larger areas of increased and decreased densities (i.e. those areas spanning the length or width of

the pitch), diagonal mass movements would almost certainly have taken place. However, the extent to which these diagonal movements could influence the final Wasserstein distance along the x and y axes is partly related to the non-square 70m x 120m dimensions of the rugby league pitch. Figure 5.7 shows the difference in the size of movements in the x, y and diagonal directions when attempting to move mass from the centre of the pitch in the opposition 20m to the left side of the pitch on the team's own try line. It is clear when comparing the left plot to the right plot that the diagonal line is significantly longer relative to the horizontal line, denoting movement along the x axis, than it is to the vertical line, denoting movement along the y axis. Consequently, the diagonal cost of mass movement would have a much greater impact on differences along the horizontal x axis. Although it is extremely unlikely that the movement all occurred from a singular point of density as depicted in Figure 5.7, the difference in this diagonal to horizontal/vertical movement ratio helps to explain why Castleford, with their horizontal, left side of pitch bias (Figure 5.6) had a much greater W_1 - W_3 difference than Salford, who had a vertical, anywhere but the opposition 20m pitch bias (Figure 5.6). The overpowering influence of diagonal mass movements on the W_1 measure ensures that it is less appropriate for the analysis of differences in the spatial trends of attacking performances in rugby league than the W_2 or W_3 distances. Given the similarities between W_3 and W_2 in Table 5.6, the W_3 distance was used for later analyses as the rectangular shape it produces when x and y factors are considered is much more appropriate for this rugby league approach than the ellipsoid produced by W_2 (Figure 5.8).

5.3.2.2 Normalised Axis Wasserstein Distance Differences at the Between Team Level

To gain an understanding of x and y axis differences in the spatial trends of team attacking performances, the Wasserstein distance cost function was weighted and compared to the standard, unweighted Minkowski distance ($p = 3$) cost function. Six weighting factors were applied to the x and y dimensions of the cost functions, enabling the calculation of different Wasserstein distances. Two equal weightings were used: 2x, 2y and 3x, 3y; four unequal weightings were used: 2.4x, 1.4y; 3.6x, 2.1y; 4.8x, 2.8y; and 6.0x, 3.5y. The weightings increased the cost of moving along the stated axis by the given factor, while the cost of moving along the other axis remained unaltered. Quadrant plots of the normalised




Figure 5.7: Visualisation of the diagonal distance when moving mass from the centre of the pitch relative to the horizontal/vertical components of the movement. Left figure shows that the diagonal distance covered to move mass from the opposition 20m to the team’s own try line is significantly greater than the horizontal distance. Right figure shows that the distance covered is much more even between diagonal and vertical movements, reducing the influence of diagonal movements on the overall Wasserstein distance when mass is moved vertically relative to horizontally.

axis Wasserstein distance (Equation 5.5) were used to evaluate the influence of differences in these factors on the results obtained. Figures 5.9 and 5.10 show the equally weighted Wasserstein distance quadrant plots using factors of 2 and 3 respectively. Figures 5.11 and 5.12 show the unequally weighted Wasserstein distance quadrant plots with factors of $4.8x$, $2.8y$ and $6.0x$ and $3.5y$ respectively. In all four plots it is evident that increasing the factor increases the normalised axis Wasserstein distance. This pattern is also evident for the two weighting factors that are not plotted ($2.4x$, $1.4y$ and $3.6x$, $2.1y$).

Given the rectangular dimensions of a rugby league pitch (70m in x-direction and 120m in y-direction), unequally weighted factors (e.g. $6.0x$, $3.5y$) were calculated in




Figure 5.8: Advanced Wasserstein metrics for $p = 1$ (left), $p = 2$ (centre) and $p = 3$ (right). Plots are shown for x -factor = 6, y -factor = 1 and highlight the rectangular shape of $p = 3$ compared to the diamond for $p = 1$ and oval for $p = 2$. Distances are calculated from central point.




Figure 5.9: Quadrant plot for Wasserstein distance using factor 2 for x and y directions. y -axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x -axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.

order to assess any differences in the median normalised increase along the x and y axes. Using an x -factor of 12/7 times greater than the y -factor equalised the cost of moving the full length of the pitch in both directions, so it was possible to assess these differences in an equal manner. Interestingly, the influence of changes along the y -axis was typically still greater than those in the x -axis when the weightings were unequal. It is possible that this is because changes across the width of the pitch (i.e. those shown by St Helens and




Figure 5.10: Quadrant plot for Wasserstein distance using factor 2 for x and y directions. y-axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x-axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.




Figure 5.11: Quadrant plot for Wasserstein distance using 4.8x, 2.8y factors. y-axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x-axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.

Leeds) encompassed a greater proportion of the pitch than those along the length of the pitch. To this end, only Castleford's differences on the left side of the pitch spanned its complete length. Other teams, including Huddersfield who had the most intense area of




Figure 5.12: Quadrant plot for Wasserstein distance using 6.0x, 3.5y factors. y-axis represents the normalised increase in Wasserstein distance in x (left) or y (right) dimensions. x-axis represents standard Wasserstein distance for $p=3$. All distances represent differences between the team overall KDE and the whole league KDE. Blue lines represent median values.

density greater than the whole league KDE for 60m in the centre of the pitch, still had an area of lower density in the centre of the pitch in the opposition 20m (Figure 5.6). This area of lower density in the opposition 20m reduced the proportion of the pitch covered by Huddersfield's central bias, thus limiting the normalised axis Wasserstein distance along the x-axis. Indeed, because Huddersfield's area of lower density in the opposition 20m encompasses the whole width of the pitch, the normalised axis Wasserstein distance along the y-axis was actually greater than that along the x-axis.

By comparing the normalised axis Wasserstein distances to the standardised $p = 3$ Wasserstein distance, it is possible to gain clear axis-level insights at the between team level. These insights provide a high level understanding of the differences between a team's overall KDE and the whole league KDE. For example, the x plot of Figure 5.12 is able to highlight the importance of x axis differences to Castleford and Huddersfield (which are clearly visible in Figure 5.6), as well as those of Wakefield, which are less visible on the plot but are important relative to the size of the overall difference between Wakefield and the whole league KDE. Similarly, along the y axis, it is possible to see that Catalans and St Helens have large vertical differences relative to the size of their difference with the whole league KDE, whereas Leeds and Salford have large differences

along the y axis, as well as large differences overall. The low value of Castleford in the y axis plot isn't mirrored by Huddersfield, which highlights the discussion in the previous paragraph surrounding the greater influence of differences covering the whole length or width of the pitch on the Wasserstein distance. Coaches and performance analysts can use these quadrant plots as an easy method of understanding how they compare to other teams in the league. For example, if Castleford were winning the league, coaches could look at the plots and see clearly that they have a greater difference along the x-axis than any other team. Strategies could then be developed to try and imitate this.

5.3.2.3 Directional Wasserstein Distance Differences at the Between Team Level

Although manipulating the cost function as described in Section 5.3.2.2 can provide useful high level insights, it is limited by its inability to provide an understanding of directional differences (i.e. it is unable to show that Castleford have a large bias towards the left side of the pitch, only showing that they have a bias along the x-axis). One method through which further insight into the direction of the differences between teams could be obtained is through the directional Wasserstein distance, which is calculated by extracting information from the transport matrix. Equations 5.6(W_{left}), 5.7(W_{right}), 5.8(W_{up}) and 5.9(W_{down}) show how the directional Wasserstein distance was calculated from the transport matrix to provide information regarding directional differences in the spatial trends of team attacking performances.

Figure 5.13 shows how comparing left and right directional Wasserstein distances along the x-axis and up and down directional Wasserstein distances along the y-axis can provide insight into the directional differences present between teams. The left plot shows the left-right Wasserstein distance differences. As with the 6.0x, 3.5y weighted distances (Figure 5.12), Castleford and Huddersfield have the two greatest distances but now they are differentiable in terms of their direction, as Castleford shift predominantly to the left whereas Huddersfield shift predominantly to the right. Similarly, Hull Kingston Rovers and Warrington are noted as teams who shift left and Wakefield, Leigh and Hull are noted as teams who shift right relative to the whole league. It should be noted that the manifestation of these differences in Figure 5.6 is not uniform. For example, Castleford and Hull Kingston Rovers have higher densities on the left side of the pitch, which cause their left shift, but Warrington's left side bias is caused by lower densities on the right side of




Figure 5.13: Quadrant plot showing directional Wasserstein distance at the between team levels. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. Blue lines represent median values.

the pitch with areas of higher density spread more evenly between central and left areas. Similarly, Huddersfield's right side bias is caused by lower densities on the left side of the pitch rather than a specific right side preference. However, the general direction of movement for all teams aligns well with the KDE plots.

The right plot in Figure 5.13 shows the up-down Wasserstein distance differences. Once again, the two teams with the largest values in the 6.0x, 3.5y weighted distance plots have the greatest differences and are plotted on different ends of the graph: Leeds have a predominantly upward shift towards the opposition 20m; and Salford have a predominantly downward shift of density towards their own try line. Hull Kingston Rovers, Wakefield, Huddersfield and Warrington who all have large positive or negative densities in the opposition 20m are also detected as teams with larger upward/downward density movement using this method. Catalans, who were shown as having one of the larger Wasserstein distances using the weighted method have a much smaller Wasserstein distance difference, which shows a small upward shift. This provides a better representation of the paler colours present within their KDE (Figure 5.6).

On the whole the differences shown by the directional Wasserstein distance plots provide significant information in a simple format which coaches can use to identify in greater detail how they compare to other teams within the league. The incorporation of the directional element (as opposed to the axis element described in Section 5.3.2.2) is an

extremely valuable tool to enable these comparisons.

5.3.2.4 Evaluation of Within Team Differences

Conducting analyses at the between team level as described in the sections above allows practitioners to understand the spatial trends of attacking performances for a given team relative to other teams in the league. In this section, the results of the analysis are discussed at a within team level. At a within team level, the results of the directional Wasserstein distance comparisons provide useful insights into the spatial trends of attacking performances for a selected team against different opponents. For some teams (e.g. Castleford; Figure 5.15), these differences are easily noticeable when all the team's data is present; for others (e.g. Catalans; Figure 5.17), it is important to zoom in on the majority of the data as a single extreme value can skew the graph's axes. This section provides an example of how these plots can be used to understand teams' spatial trends of attacking performances and how they relate to match performances against different opponents for two clubs: Castleford and Catalans.

Figure 5.14 visualises the within team attacking differences for Castleford's KDEs against different oppositions; Figure 5.15 shows their within team quadrant plot. The left plot of Figure 5.15 shows left sided biases for matches against Hull Kingston Rovers, Catalans and St Helens, which are congruent with the respective plots in Figure 5.14. These biases are in addition to the left side bias that Castleford already shows at a team level ensuring that practitioners can be sure that Castleford attacked on the left side of the pitch against these opponents. The right side bias for Salford and Huddersfield is reciprocated by areas of higher density on that side of the pitch in Figure 5.14 for both teams. However, the specific locations of these biases differ, with Huddersfield having an area of green density centrally in the opposition 20m, but down the right side of the pitch in between the opposition 20m and own 20m and Salford following the reverse pattern (right side in the opposition 20m, central between own and opposition 20m). The right side of Figure 5.15 shows an upward shift against Hull and Huddersfield, likely due to the increased densities in the opposition 20m.

The within team plots provide a method through which coaches can consider the spatial trends of attacking performances for a team against a given opponent. This can provide valuable information when preparing for matches against upcoming matches, partic-

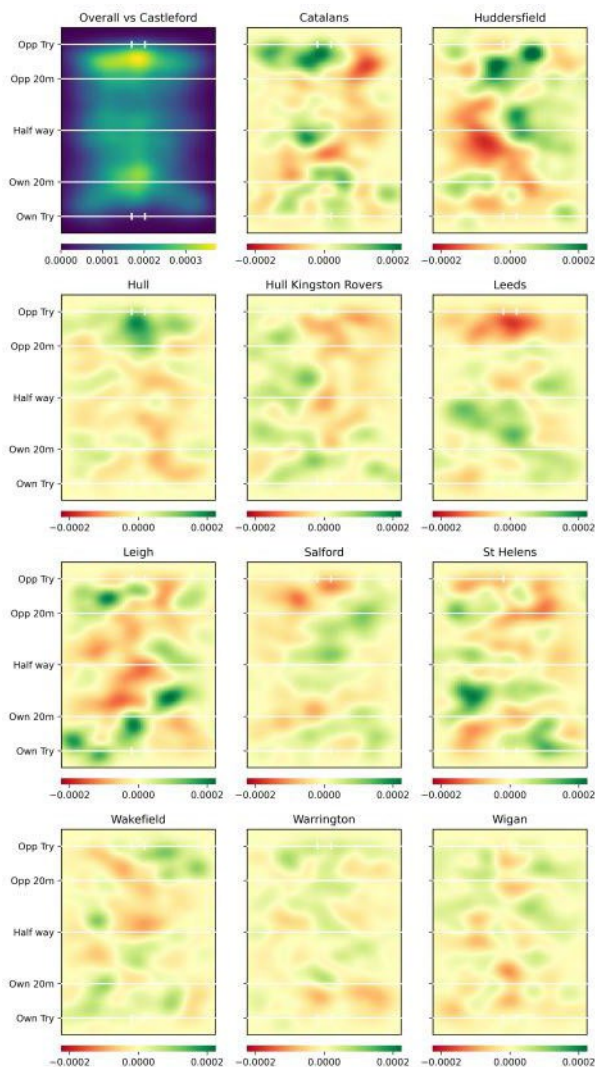


Figure 5.14: Comparison between team-opponent KDE and team's overall KDE for Castleford. Top left plot shows overall KDE for Castleford. For all other plots, green areas are areas represent areas where the team has a higher density against a specific opponent than the overall KDE, red areas have lower densities.

ularly if the result of the matches are known. For example, if Huddersfield beat Castleford, a coach could infer that part of the result may have been due to Huddersfield successfully reducing Castleford's ability to attack down their favoured left side. However, it should be noted that it is currently unclear how much information the plots provide with regards to the scoring performance of teams. For example, Castleford showed a left-sided shift




Figure 5.15: Quadrant plot showing directional Wasserstein distance for Castleford against different opponents. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. Blue lines represent median values.

against Hull Kingston Rovers, Catalans and St Helens. Across the season, they played 5 matches against these teams, scoring 26 and 19 points against Hull Kingston Rovers, 6 against Catalans and 0 and 20 against St Helens. Taken in isolation, these results have a limited relationship with scoring performance but when it is considered that Catalans and St Helens were the two best teams in the league in 2021, whereas Hull Kingston Rovers finished close to the bottom of the league, the ability to score 20 points against St Helens could actually be considered a positive outcome.

Figure 5.17 shows the quadrant plot for Catalans skewed by the outlying opponent Leeds. It is skewed by Leeds as Catalans had a much higher density in the opposition 20m than their season average (Figure 5.16). When removing Leeds as in Figure 5.18, it is easier to visualise the differences between opponents. The left plot shows that as the size of difference between the team-opponent KDE and the team overall KDE increases, so does the tendency to shift the density of actions rightwards. This is true for all opponents except St Helens, who are also one of only two teams to result in a shift downwards as the Wasserstein distance between the team opponent KDE and the team overall KDE increased.

As with Castleford, it is possible for coaches to gain insights into the performances

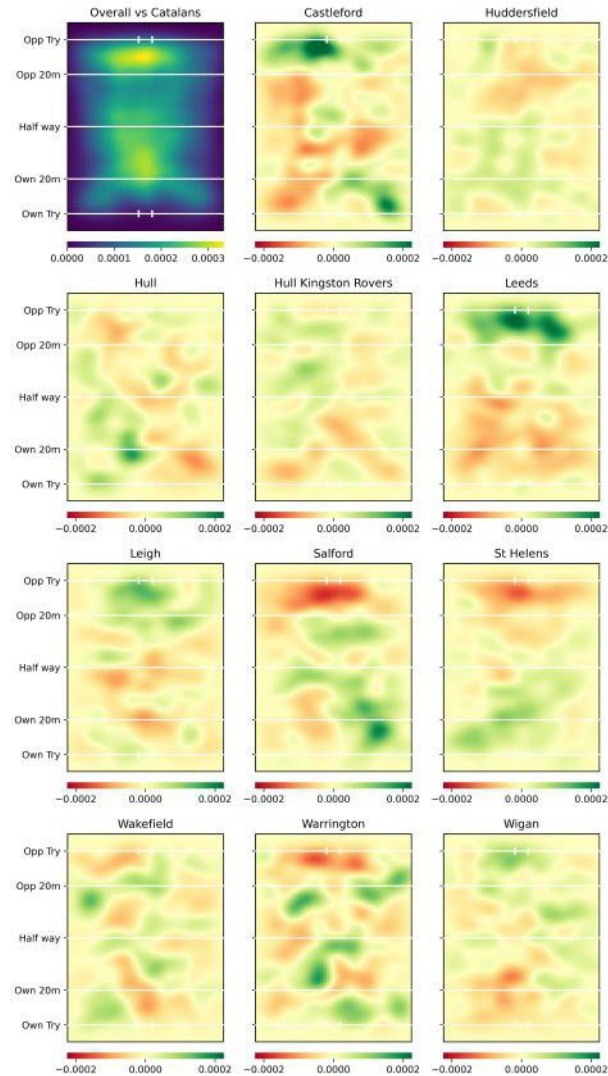


Figure 5.16: Comparison between team-opponent KDE and team's overall KDE for Catalans. Top left plot shows overall KDE for Catalans. For all other plots, green areas are areas represent areas where the team has a higher density against a specific opponent than the overall KDE, red areas have lower densities.

of Catalans on an opponent by opponent basis using the quadrant plots, which can provide important knowledge with respect to developing tactical strategies for future matches against them. Once again though, further context is required to establish the likely pattern of the matches the team played in. From Catalans perspective, the three key teams who were separate from all others within their directional plot (Figure 5.16) were Leeds




Figure 5.17: Quadrant plot showing directional Wasserstein distance for Catalans against different opponents. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. Blue lines represent median values.

(27 and 26 points scored, very large upwards shift), Salford (42 and 42 points scored, large right shift and moderate downward shift) and St Helens (20, 12, 31 and 10 points scored, large downward shift). These teams are interesting when contextual knowledge is applied because it is possible to identify three different types of matches which Catalans may have experienced based on the abilities of the opposition. St Helens won the Super League Grand Final in 2021, so it is plausible that the majority of matches against them would have been really hard fought matches in the centre of the pitch, potentially describing the large downward shift in actions relative to Catalans team overall KDE. Salford finished 11th in the league, conceding 584 points across the league suggesting their defence was weak. As such, the moderate downward shift in density may have been because Catalans been able to score points against them in fewer actions than against Leeds who finished 5th in the league and may have been much better at defending, causing Catalans to perform more actions in the opposition 20m before they were able to score points. The ability to gain these insights from a single plot of analysis (alongside widely known information for coaches) ensures KDE and directional Wasserstein distance plots are a particularly valuable tool for use within rugby league.

The work from this study was presented to coaches at Leeds Rhinos rugby league




Figure 5.18: Zoomed in quadrant plot showing directional Wasserstein distance for Castleford against different opponents. Positive numbers in the left plot represent an overall shift to the left, positive numbers in the right plot represent an overall shift upwards. The opponent Leeds is removed. The co-ordinates for Leeds would be (9.41, -0.58) and (9.41, 7.14) for the left and right plot respectively. Blue lines represent median values, including Leeds.

club, who were widely impressed. They validated the reliability of the findings regarding Huddersfield's central densities and also explained that Leeds Rhinos' increased density within the opposition half was probably due to the team's ability to progress into the opposition half particularly quickly in the 2021 season. They had not noticed the left sided bias shown by Castleford and indicated had they known this it may have influenced their tactical preparation for matches against them providing strong validation of the usability of the models. They strongly believe the plots have use from a bench marking perspective - that is, understanding how the team is performing compared to others and whether there are areas that they can improve or investigate further as a result.

The analyses for this study were conducted in Python using the scikit-learn (Pedregosa et al.,2011) and Pythonot (Flamary et al.,2021) packages.

5.4 Summary

In this study, the work of [Mallepalle et al.\(2020\)](#) was adapted, validated and extended in rugby league to provide a method through which KDEs could be used to visualise the spatial trends of attacking performances at the whole league, team overall and team-opponent levels. Two novel metrics were derived from the Wasserstein distance to evaluate between and within team differences in the spatial trends of attacking performances: the normalised axis Wasserstein distance provided insights at the axis level (e.g. x, y differences); and the directional Wasserstein distance provided directional insights (e.g. left, right, up, down differences). The insights provided by the directional Wasserstein distance were shown to be particularly valuable to coaches from a bench marking and tactical strategy development perspective.

The work was presented at the UKCI conference and will subsequently be published in *Advances in Computational Intelligence Systems*. It was also presented to coaches at Leeds Rhinos, who were impressed with its ability to provide high level spatial information in such a simple way, providing support for its usability ("I believe [the] work has the capacity to have a tangible impact on the day to day operation of an elite rugby league team") and reliability ("it passes the 'eye test' in terms of validity to expected outcomes"). [Appendix D](#) provides the full impact statement from this meeting. The coaches suggested that one method of tailoring the analysis to their needs in the future would be to perform a KDE/Wasserstein distance analysis based on actions performed against a specific team (rather than the current approach of considering all actions performed by a specific team) so that all teams' performances against an opponent could be considered more easily. They also suggested performing the analysis at a match by match level would be interesting, however this idea could run into sample size related random variation issues which may limit the usefulness of the plots.

Future work should now consider applying the KDE and Wasserstein distance methods employed in this study to other sports. There is a direct link between its usability in rugby league and other invasion sports as by understanding the areas through which a team is likely to control the ball, it can be possible to understand where they are likely to attack even if the methods here don't allow for the value of these areas to be considered. An alternative application however may be in tennis or golf. Although in these sports there is no concept of "possession", players are still likely to have tendencies in the

direction in which their shots move. Consequently, the methods proposed in this study could provide important tactical insights into shifts in serve direction in pressure points in tennis (e.g. when serving at break point or serving for the set/match). Similarly, in golf, the directional Wasserstein distance could provide an understanding of those players who are more likely to drive to the left or right, shorter or longer. Knowledge of these elements could provide training and/or scheduling insights with respect to different players' shot making tendencies as they could help to understand which type of player is more successful at different courses.

The results of this chapter provide a valuable method through which the spatial trends of attacking performances can be evaluated between and within teams in rugby league. However, they are subject to some limitations. The first of these is the difficulty the Wasserstein distance metric has with quantifying differences in the centre of the pitch due to the smaller distance required to move to any other location on the pitch. The second is the limited relationship present with the points scored by teams without additional contextual information being provided. Such a limitation is understandable as all locations on the pitch were treated equally within the KDE analyses, which is not the case within the sport where controlling the ball closer to the opposition try line is significantly more valuable. As such, combining the results of the analyses here with a model which is able to value locations on the pitch like Chapter3, but in a smooth manner could prove to be an extremely powerful performance measurement tool.

Building on the work of this study, which deliberately simplified the research question to only consider the probability of controlling the ball in a given location, the next study will produce a model which provides a smooth surface estimating the value of an action conditional on its location. This will allow context to be added to the team KDEs produced in this chapter in a systematic and quantitative manner. The model will also provide a method of rating player performances, building on the expert feedback received in Chapter4.

A Bayesian Approach to the Evaluation of Team and Player Performances

6.1 Introduction

In Chapter 5, the spatial trends of team attacking performances were evaluated in a smooth manner using KDEs. The work was well received by coaches at Leeds Rhinos but required further information regarding performances (e.g. the number of points a team scored against the opposition, or the league position of a team) to be used most effectively in practice. Furthermore, it was unsuitable for the analysis of player performances.

This study provides a framework through which team and player performances can be evaluated in a data driven manner (i.e. without needing additional contextual information) in rugby league. A novel EPV model, which uses a Bayesian Mixture Model approach to provide a smooth pitch surface is proposed, improving upon previous zonal aggregation methods (Chapters 3 and 4; Cervone et al.(2016); Kempton et al.(2016); Singh (n.d.)). The model derives the EPV from estimates of the probability of each point scoring outcome occurring at the end of the possession. By estimating individual possession outcome probabilities in a single model, this study improves all previous work, which has either aggregated the calculation of all possession outcome probabilities (Chapters 3 and 4; Cervone et al.(2016); Chan et al.(2021); Kempton et al.(2016)) or produced separate models for different possession outcomes (Decroos et al.,2019). It provides significantly greater flexibility to the model as possession outcome probabilities have the

potential to provide extensive tactical insight into player and team performances (e.g. are there areas on the pitch where a team is more likely to score a converted try or a penalty goal?). Novel performance metrics are proposed, combining the results of this study with the KDE model from Chapter 5, allowing team and player performances to be evaluated in an objective, data driven manner. The results have been presented to coaches at Leeds Rhinos (Appendix D provides an impact statement), who provided support for its validity.

6.2 Methodology

In this section, the methodology for the study is described. It describes the data and data preprocessing steps used in this study. The novel EPV model, which estimates the probability of individual possession outcomes using a Bayesian Mixture Model approach, is outlined. Novel expected points scored and actual vs expected ratings metrics are proposed to evaluate team and player performances.

6.2.1 Data

In this study, event level match-play data were obtained from Opta (Stats Perform, London, UK) for the 2021 Super League season. The data was produced via human annotation of the actions taking place and was downloaded from www.optaprourugby.com. In the 2021 Super League season, 138 matches were contested by 12 teams and 373 players. A total of 557,050 match events were completed (median 4003 events per match, interquartile range 3857-4200). Across the season, 1001 tries were scored (768 successful conversion kicks, 233 unsuccessful conversion kicks), 175 penalty goals were attempted (158 successful, 17 unsuccessful) and 83 drop goals were attempted (37 successful, 46 unsuccessful).

Chapter 3 introduces the data provided by Opta. Like all previous chapters, 13 variables were considered for this analysis. These were: ID, FXID, PLID, team id, MatchTime, x_coord, y_coord, action, ActionType, Actionresult, Metres, PlayNum and SetNum. Table 3.2 defines these variables. The dataset is identical to that used in Chapter 5, so action by action location data were available. Section 5.2.1 provides an overview of the actions in the data and its presentation in individual xml files.

6.2.2 Data Preprocessing

This study used the same dataset and initial preprocessing steps as Chapter5. Section 5.2.2 describes the steps taken to convert the 138 raw xml files described above to a single file of 99,966 rows and 4 columns. In this study, these steps are followed until the file contained 99,966 rows and 18 columns (i.e. until the penultimate step, as the last step filtered the data to the four columns used in Chapter5).

The dataset for this study was devised using different inclusion and exclusion criteria for actions compared to Chapter4, due to the two studies having subtly different aims. Chapter4 attempted to estimate the expected value of a possession conditional on its location, action tuple. This resulted in the study's possessions being defined to clearly replicate sequences of actions within match play. However, in this chapter, the aim was to estimate the probability of scoring points conditional on the location on the pitch. It was therefore necessary to reduce the number actions included in any given location so that only sequences of location changes are provided to the model. For example, after a tackle has been completed, a play-the-ball occurs, which is received by a player, who chooses to pass or run with the ball. All of these actions typically occur in the same location. In Chapter4, all three actions were included as they allowed for the value of each action to be estimated. However, in this chapter only one of the actions (the pass or run) was included as including all three actions would ensure the sequence of location changes was coded incorrectly. As a result of the different inclusion and exclusion criteria, the key actions removed from this study's dataset compared to Chapter4 were play-the-balls (39,752) and pass catches (100,952). Kick catches remained in the dataset as these would typically begin (when receiving a kick from the opposition team) or end (when catching a kick from the same team and trying to score a try) a possession.

Possessions followed the same definition as episodes in Chapters3 and 4. A possession was defined to begin when a team successfully gained possession of the ball and ended due to a handover, loss of possession due to an error/foul play, points being scored or a goal kick attempt. As with previous chapters, it was possible for an attacking possession to encompass more plays than the typical attacking set of 6 tackles if an error/foul was made by the opposition team.

Possession outcomes were treated as discrete categories in this study to enable the estimation of individual probabilities. Unlike Chapters3 and 4, where every action was

coded based on its unique possession outcome, in this study, every action in a possession was marked with the possession outcome category from the end of the possession. Five possession outcome categories were used: converted try; unconverted try; penalty goal; drop goal; and no try. These are the same as the rewards previously described in Chapters 3 and 4. However, by considering the outcomes discretely (i.e. as categories), rather than providing a numerical value (e.g. 6 points for a converted try), it was possible to estimate the probability of all five possession outcomes occurring from a given location in a single model. No study has yet attempted this method of analysing possession outcomes, but doing so creates an extremely flexible model which can provide detailed insights into team performances. For example, understanding where teams are more likely to end possessions by scoring drop goals or penalty goals is useful when preparing strategies to face opposition teams. Similarly, it allows visualisation of those areas on the pitch that are most likely to result in tries being scored by the opposition team. It would not be possible to gain these insights if the numerical value of the possession outcome was aggregated like in previous Markov models (Chapters 3 and 4; Cervone et al.(2016); Chan et al. (2021); Kempton et al.(2016)).

To enable the estimation of individual possession outcomes at multiple levels of analysis, the data for this study were organised into 25 subsets. Similar to Chapter 5, 13 of these subsets represented the whole league data (1) and team attacking data (12). After considering the feedback from Chapter 5, this study also utilised 12 subsets representing team defending data (i.e. all actions taken against that team). Table 6.1 provides descriptive statistics for the subsets of data included within the study. There is remarkable similarity in the median number of observations per team attacking and team defending subsets, but more variability in the attacking subsets. This variability suggests the number of actions a team performs when attacking is more varied than the number of actions they face when defending.

The final dataset contained 99,966 observations and 7 columns. Included within these columns were: the attacking and defending team, which acted as identifiers for data subsets; the player ID, which allowed player performances to be evaluated; the x, y locations of every action; the possession number, which identified unique possessions; and the discrete possession outcome. Table 6.2 provides a sample of the final dataset. It shows an attacking possession for St Helens against Salford, where seven players are involved in advancing the ball towards the opposition try line, but no points are scored.

Table 6.1: Descriptive data for data subsets. IQR is interquartile range; Min/Max refers to minimum and maximum number of observations.

Comparison	Subsets	Median	IQR	Min/Max
Whole league	1	99966		
Team attacking	12	8105	7596-8937	7203/10324
Team defending	12	8077	7878-8700	7579/9620

Table 6.2: Sample possession within the dataset. Data includes the teams involved in the possession, the player ID, the x, y coordinates of the action, the possession number (PosNum) and the possession outcome (PosCat; in this case no try for all rows).

Attacking Team	Defending Team	Player ID	x	y	PosNum	PosCat
St Helens	Salford	3107	9	4	1	0
St Helens	Salford	21716	9	6	1	0
St Helens	Salford	1983	14	11	1	0
St Helens	Salford	2904	22	13	1	0
St Helens	Salford	11439	12	12	1	0
St Helens	Salford	21795	37	16	1	0
St Helens	Salford	20567	36	24	1	0
St Helens	Salford	2904	54	35	1	0

6.2.3 Bayesian Mixture Model

In this study, a Bayesian Mixture Model was used to develop a novel EPV model capable of evaluating team and player performances. Bayesian analysis (Definition 6.2.1) is ideally suited to the estimation of this model in rugby league for several reasons. Firstly, it adopts an evidence based approach, which allows it to estimate certainty and uncertainty within parameter estimates. This ensures that in areas of the pitch where little evidence is available, a parameter value can be estimated, but confidence within the estimate can be tempered as appropriate. Using the Markovian approaches of Chapters 3 and 4, this was not possible. Secondly, the use of prior distributions can assist the model in understanding likely parameter values, reducing the amount of data required to generate plausible parameter estimates. This is in contrast to the machine learning (Decroos et al., 2019) and deep learning (Fernández et al., 2021; Liu et al., 2020) approaches previously used

in sport, which randomly initialise parameters, and is particularly important in this study where only 99,966 observations are available. Finally, the complexity of the output (five possession outcome probabilities for 33 centres) is well suited to a custom model formulation, which is easier to solve using an evidence based Bayesian approach than through machine learning methods.

Definition 6.2.1. Bayesian statistics is an approach to data analysis centred around Bayes Theorem (Equation 6.1). It consists of three main elements [van de Schoot et al.\(2021\)](#):

- the *prior distribution*, which outlines current knowledge about the parameter
- the *likelihood function*, which is the function through which the data influences the posterior distribution
- the *posterior distribution*, which balances the prior knowledge (as provided through the prior distribution) with the observed data to identify the most plausible parameter values given the evidence available

For use within Bayesian analysis, Bayes Theorem (Equation 6.1) is normally rewritten as Equation 6.2 allowing it to produce parameter estimates (θ) conditional on the observed data (y):

$$p(B|A) = \frac{p(B \cap A)}{p(A)}, \quad (6.1)$$

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}, \quad (6.2)$$

where $p(\theta|y)$ and $p(y|\theta)$ are conditional likelihoods of θ given y and y given θ , respectively, and $p(\theta)$ and $p(y)$ are the likelihoods of θ and y , respectively.

The novel EPV model produced in this study used a Bayesian Mixture Model approach to provide a smooth pitch surface. In this approach, the model estimates the probability of each possession outcome ($s \in S$) for a set of centres on the pitch. Each centre holds probabilities for the five possession outcomes, represented as a 5-dimensional vector. The probabilities for any x, y location on the pitch are calculated by taking a weighted average of probabilities at all centres on the pitch using a set of weights Z . Denoting by $P(s; x, y)$ the probability of possession outcome s at location (x, y) is calculated as

$$P(s; x, y) = \sum_k z_k(x, y) P_k(s) \quad (6.3)$$

where $z_k(x, y)$ is the weight corresponding to the location (x, y) and k -th centre and $P_k(s)$ is the probability of possession outcome s at the centre k .

The weights for each centre were treated as fixed for every x, y location on the pitch (Section 6.2.3.1) so Bayesian analysis was used to compute the distribution of possession outcome probabilities at all centres P_k . The prior distribution for P_k is Dirichlet with the parameter, α :

$$P_k \sim \text{Dirichlet}(\alpha).$$

The prior distributions are independent between centres (i.e. there are separate prior distributions, denoting the prior probabilities of all five possession outcomes occurring, for centres 1, 2, 3, 4 etc.). The Dirichlet distribution is a multivariate generalisation of the beta distribution, parameterised by a vector of positive reals (α). Its output is a vector of non-negative numbers summing up to one, i.e., a vector of probabilities. As it is only possible for one possession outcome to occur every possession in rugby league, this distribution is ideally suited as the prior for possession outcome probabilities.

The likelihood function for a set of data, $D = (x_i, y_i, s_i)_{i=1}^n$ is given by:

$$P(D|\alpha) = \prod_{i=1}^n \sum_k z_k(x_i, y_i) P_k(s_i|\alpha_k),$$

where $\alpha = (\alpha_k)_k$ is the vector of centre Dirichlet parameters and

$$P_k(s|\alpha_k) = \int \pi_s p_{\text{Dirichlet}}(\pi|\alpha_k) d\pi,$$

where $p_{\text{Dirichlet}}(\pi|\alpha_k)$ is the density of Dirichlet distribution with parameter α at point $\pi \in \{(\pi_1, \dots, \pi_5) \in [0, 1]^5 : \pi_1 + \dots + \pi_5 = 1\}$.

With the above notation, the prior is $\alpha_k = \alpha$ for all k ; the posterior distribution is the distribution of $p(\alpha|D)$. It is intractable as each observation influences 4 centres (in the field of play) or 2 centres (in the try area), and the weights and centres vary between observations. Therefore, Markov Chain Monte Carlo (MCMC) sampling was used to estimate the probabilities at each centre. MCMC sampling methods allow for the systematic random sampling from high dimensional probability distributions and those samples were used to draw conclusions on the posterior distribution of probabilities on the pitch.

6.2.3.1 Centre Weights

Within a Bayesian Mixture Model approach, a set of centres are defined on the pitch, through which the data is aggregated using a set of weights. Previous EPV models have aggregated x, y locations on the pitch into zones (Chapters 3 and 4; Cervone et al.(2016); Kempton et al.(2016); Singh(n.d.)). Using a zonal approach, the data for every location within the zone is aggregated equally resulting in dichotomous “boxed” values (Figure 3.6). This aggregation is necessary to allow EPV estimates to be produced but results in situations where moving 1m from the current location could lead to an unreasonably large difference in EPV if it results in moving to a different zone. In this study, the centres can be considered analogous to the corners of zones. However, by treating the centres individually, rather than together as a zone, a custom set of unequal weightings can be produced, which better reflect the proximity of the x, y location to the centre. For example, suppose a location is 20m away from centre 1 but only 2m away from centre 2. Using a zonal approach, the two centres would be equally weighted when aggregating data. In this study, using linear and bilinear interpolation techniques, centre 2 would have a much larger weight than centre 1, reflecting the spatial proximity of the location much more accurately and allowing a smooth pitch surface to be formed.

After consultation with professional experts, 33 centres were placed around the pitch. 30 centres were located in the field of play, uniformly positioned at $x \in \{0, 20, 35, 50, 70\}$ and $y \in \{-10, 20, 35, 65, 90, 100\}$. These locations were chosen to ensure that every location within the field of play would fall within a “zone” defined by four centres. Figure 6.1 plots the centre locations for the field of play. 3 centres were located in the opposition team try area ($x \in \{0, 35, 70\}$). No y coordinate was considered for centres in the try area as the actions players choose in this area are not influenced by their y coordinate. The three try area centre locations were produced to ensure that the x coordinate of every location in the try area fell between two centres. The field of play and try area centres were evaluated separately. This decision was made due to the different player behaviours that are observed in the two areas: in the field of play, players are equally likely to choose different actions dependent on the game situation; in the try area, players will attempt to ground the ball for a try as soon as possible irrespective of the game situation.

Every x, y location was assigned 33 weights, which described the aggregation of their data to the centres in the model. 30 weights were calculated in the field of play using




Figure 6.1: Location of 30 centres used in this study. Three centres (not plotted) were included in the try area at $x \in \{0, 35, 70\}$, equivalent to the left, middle and right centres in the field of play. No y coordinate was considered for these centres.

bilinear interpolation, 3 weights were calculated in the try area using linear interpolation. In line with the assumption of independence between the two areas of the pitch, any location in the field of play was automatically given weights of 0 for the try area centres; any location in the try area was automatically given weights of 0 for the field of play centres.

Locations in the field of play had a maximum of four non-zero weights. The value of the weights for each x, y location in the field of play was derived from the distance between the x, y location and the four centres surrounding it in a quadrilateral shape. These centres had coordinates (x_1, y_1) , (x_1, y_2) , (x_2, y_1) and (x_2, y_2) . The weights (z_{11} , z_{12} , z_{21} and z_{22}) of these centres for location x, y were calculated using Equations 6.4, 6.5, 6.6 and 6.7. The remaining centres were assigned a weight of 0.

$$z_{11} = \frac{(x_2 - x)(y_2 - y)}{(x_2 - x_1)(y_2 - y_1)}, \quad (6.4)$$

$$z_{12} = \frac{(x_2 - x)(y - y_1)}{(x_2 - x_1)(y_2 - y_1)}, \quad (6.5)$$

$$z_{21} = \frac{(x - x_1)(y_2 - y)}{(x_2 - x_1)(y_2 - y_1)}, \quad (6.6)$$

$$z_{22} = \frac{(x - x_1)(y - y_1)}{(x_2 - x_1)(y_2 - y_1)}. \quad (6.7)$$

Locations in the try area had a maximum of two non-zero weights. Here, only the x location of the centres ($x \in \{0, 35, 70\}$) was considered so the weights were derived from the distance between the x coordinate of the action location and the x coordinate of the centre. For an x, y location in the try area, linear interpolation between the two closest centres x_1, x_2 , from the above set of three, with $x_2 > x_1$, was used to calculate two weights z_1, z_2 ; the weight of the remaining centre was set to 0. The non-zero weights were given by:

$$z_1 = \frac{x_2 - x}{x_2 - x_1}, \quad (6.8)$$

$$z_2 = \frac{x - x_1}{x_2 - x_1}. \quad (6.9)$$

6.2.3.2 EPV Calculation

The EPV for a location was derived from the possession outcome probabilities. The EPV for a location with coordinates x, y was calculated using the probability of the five possession outcomes and their true points scoring values:

$$EPV_{(x,y)} = \sum_{s \in S} P(s; x, y) \text{Points}(s) \quad (6.10)$$

where $P(s; x, y)$ is the probability of possession outcome s in location x, y (Equation 6.3) and $\text{Points}(s)$ is the true point scoring value of possession outcome s (Table 2.1).

The EPV provides a single value which combines all possession outcome information like previous studies (Chapters 3 and 4; Cervone et al.(2016); Chan et al.(2021); Kempton et al.(2016); Singh(n.d.)). By calculating the EPV in this manner, the approach of this study is not dissimilar to the decomposed MDP approach of Fernández et al.(2021). However, rather than decomposing the EPV calculation into different actions, this study decomposes it into different possession outcomes.

6.2.3.3 Modelling Approach

The analysis for this study was conducted at two levels. First, the model was run using the whole league data, then the 24 team attacking and defending models were run, using the data subsets described in Section 6.2.2. It has previously been argued that a hierarchical approach is the most appropriate method of analysing data at multiple levels (van de Schoot et al., 2021), however adopting such a method would have only allowed a single averaged shift (either positive or negative) away from the whole league model estimates. As such, if a team had a greater probability of scoring a try from the left side of the pitch in the opposition try area, but a greater probability of scoring from the right side of the pitch in their own half than the whole league model estimates, it would not be possible to identify these differences using a hierarchical approach. In this study, completing two levels of analyses (i.e. whole league, then team level) allowed the team level models to include information already learned about the league within their parameter estimates through the use of an α prior distribution estimated from the whole league model posterior distribution.

The whole league data model provided league average possession outcome probabilities. For this proof of concept model, human-defined priors were used for P_k (Table 6.3). These priors were selected after discussion with experts and were informed by the results of Chapters 3 and 4. They loosely informed the model that there was a greater chance of points being scored by the end of the possession the closer the location was to the opposition try line. It would have been possible to provide more specific priors suggesting a greater probability of scoring more points in central areas as was shown in the results of Chapter 3. However, it was decided that providing more generic, loosely informative priors allowed the model to learn from the data with greater freedom.

24 models were produced at the team level, one for each team's attacking and defending data subsets. The team attacking models provided an understanding of the probability of a team achieving a given possession outcome at the end of their possession from any given location. The team defending model estimated the probability of any opposition team achieving a given possession outcome when controlling the ball at any given location against the defending team. Maximum likelihood estimation of the posterior distribution of the whole league model was used to calculate the α priors for these analyses (Table 6.4). The same α prior distribution was used for each team model.

Table 6.3: Prior α values for whole league Bayesian model. Centre coordinates are provided, alongside α values for each possession outcome. C.Try refers to converted try (6 points) and U.Try refers to unconverted try (4 points).

Centre	No Try	Drop Goal	Penalty	U.Try	C.Try
(0, -10)	90	1	1	4	4
(20, -10)	90	1	1	4	4
(35, -10)	90	1	1	4	4
(50, -10)	90	1	1	4	4
(70, -10)	90	1	1	4	4
(0, 20)	90	1	1	4	4
(20, 20)	90	1	1	4	4
(35, 20)	90	1	1	4	4
(50, 20)	90	1	1	4	4
(70, 20)	90	1	1	4	4
(0, 35)	85	1	3	5	6
(20, 35)	85	1	3	5	6
(35, 35)	85	1	3	5	6
(50, 35)	85	1	3	5	6
(70, 35)	85	1	3	5	6
(0, 65)	80	1	3	7	9
(20, 65)	80	1	3	7	9
(35, 65)	80	1	3	7	9
(50, 65)	80	1	3	7	9
(70, 65)	80	1	3	7	9
(0, 90)	75	1	3	9	12
(20, 90)	75	1	3	9	12
(35, 90)	75	1	3	9	12
(50, 90)	75	1	3	9	12
(70, 90)	75	1	3	9	12
(0, 100)	70	1	3	10	15
(20, 100)	70	1	3	10	15
(35, 100)	70	1	3	10	15
(50, 100)	70	1	3	10	15
(70, 100)	70	1	3	10	15
(0, TRY)	35	1	1	28	35
(35, TRY)	35	1	1	28	35
(70, TRY)	35	1	1	28	35

Table 6.4: α prior distribution values for team attacking and defending models. α values calculated via maximum likelihood estimation from posterior distribution of the whole league model. Centre coordinates are provided, alongside α values for each possession outcome. C.Try refers to converted try (6 points) and U.Try refers to unconverted try (4 points). Values rounded to 2 decimal places for brevity.

Centre	No Try	Drop Goal	Penalty	U.Try	C.Try
(0, -10)	184.72	0.72	2.01	4.97	6.43
(20, -10)	345.31	1.07	6.38	7.45	14.88
(35, -10)	280.81	1.37	4.63	5.08	16.00
(50, -10)	365.08	0.95	6.30	7.90	15.67
(70, -10)	230.12	0.77	1.28	5.40	6.11
(0, 20)	371.56	2.28	1.87	7.08	27.40
(20, 20)	1559.21	2.73	14.96	28.85	86.65
(35, 20)	2815.84	6.64	44.81	67.00	168.66
(50, 20)	1599.20	1.03	22.07	25.46	94.59
(70, 20)	393.10	0.73	1.75	6.92	35.35
(0, 35)	373.60	2.10	5.30	6.51	33.39
(20, 35)	1707.17	2.05	18.37	30.07	121.44
(35, 35)	2636.37	8.50	24.50	59.31	162.86
(50, 35)	1673.02	3.58	21.79	44.56	116.69
(70, 35)	310.63	3.68	6.61	10.50	25.85
(0, 65)	439.66	2.03	10.38	17.21	63.22
(20, 65)	2189.67	10.97	65.76	75.37	294.36
(35, 65)	2546.81	26.14	94.46	84.30	355.25
(50, 65)	1726.64	6.24	49.69	68.55	197.80
(70, 65)	360.60	1.82	8.02	18.25	51.31
(0, 90)	413.37	2.70	6.90	35.62	83.31
(20, 90)	1356.56	8.21	32.70	139.25	384.54
(35, 90)	2150.80	43.80	116.69	241.11	636.41
(50, 90)	1437.66	12.84	43.18	142.16	463.60
(70, 90)	365.55	0.91	11.60	34.77	50.19
(0, 100)	154.63	1.15	3.38	18.26	50.01
(20, 100)	370.25	5.46	8.15	42.98	116.60
(35, 100)	427.75	3.94	7.62	47.37	146.33
(50, 100)	299.46	4.83	7.87	40.45	99.93
(70, 100)	121.38	1.30	2.36	24.68	34.07
(0, TRY)	36.62	0.98	0.95	33.38	46.66
(35, TRY)	43.45	0.97	0.99	29.80	68.50
(70, TRY)	42.17	1.00	1.02	30.12	39.13

The mean of the posterior distribution of possession outcome probabilities P_k^μ for each model was substituted into Equations 6.3 and 6.10 to provide individual possession outcome probabilities and EPV for each x, y location:

$$P^\mu(s; x, y) = \sum_k z_k(x, y) P_k^\mu(s)$$

$$EP_{(x,y)}^\mu = \sum_{s \in S} P^\mu(s; x, y) \text{Points}(s)$$

The standard deviation of the posterior distribution of possession outcome probabilities P_k^σ for each model was used to provide an understanding of uncertainty within the parameter estimates for each x, y location. This required Equations 6.3 and 6.10 to be modified:

$$P^\sigma(s; x, y) = \frac{\sum_k z_k(x, y) P_k^\sigma(s)^2}{\sum_{s \in S} P^\sigma(s; x, y)^2 \text{Points}(s)}$$

A smooth pitch surface was produced by calculating these values for every x, y location on the pitch. Smooth pitch surfaces were produced for each of the five possession outcomes and EPV for the whole league model and every team attacking and defending model. For notation purposes, the complete smooth pitch surface is referred to from hereon in as $P_{\text{model}}^\phi(s)$ and EPV_{model}^ϕ ; where model refers to the model used to evaluate P_k (whole league (wl), team attacking (att) or team defending (def)) and $\phi \in \{\mu, \sigma\}$ denotes the posterior distribution mean or standard deviation measure used. Individual x, y location values are only used for the whole league model, and so are referred to as $P^\phi(s; x, y)$ and $EPV_{x,y}^\phi$, where $\phi \in \{\mu, \sigma\}$ denotes the posterior distribution mean or standard deviation measure used.

6.2.4 Expected Points Scored

To provide a novel measure of underlying team performances in rugby league, an expected points scored metric was produced which combined the results of this study with the KDE models produced in Chapter 5. The expected points scored for a given team (team)

versus an opponent team (opp) was calculated via Equation 6.11. It provides a quantitative measure of the underlying performance of the chosen team, by estimating the number of points they were expected to score based on the EPV model used (team attacking or opponent defending), the KDE model of their spatial trends of attacking performance and the number of scoring opportunities they had (i.e. the number of possessions they were involved in).

$$\text{Expected Points Scored}_{\text{team, opp}} = \text{EPV}_{\text{attde}}^{\mu} \times \text{KDE}_{\text{team, opp}} \times \text{Possession}_{\text{team, opp}} \quad (6.11)$$

where $\text{EPV}_{\text{attde}}^{\mu}$ refers to the mean attacking or defending team EPV smooth pitch surface used, $\text{KDE}_{\text{team, opp}}$ refers to the KDE pitch surface for the selected team in possession against the opponent (developed in Chapter 5) and $\text{Possessions}_{\text{team, opp}}$ refers to the number of possessions completed by the attacking team against the opponent.

The expected points scored calculation was completed using the team-opponent KDEs from Chapter 5. Two EPV measures were considered: $\text{EPV}_{\text{att}}^{\mu}$ the selected team's attacking mean EPV pitch surface; and $\text{EPV}_{\text{de}}^{\mu}$ the opposition team's mean defensive EPV pitch surface. Using the team's attacking EPV pitch surface in the calculation allows practitioners to understand how well the team was able to utilise their strengths against the opposition; using the opposition team's defensive EPV pitch surface allows practitioners to understand how well the attacking team was able to exploit the opposition team's defensive weaknesses. The expected points scored values were summed to provide a single performance figure for the team across the season.

6.2.5 Actual vs Expected Ratings

To evaluate player performances, Actual vs Expected (AE) player performance ratings were devised using the whole league model. Building upon feedback from Chapter 4, where actions were valued for being attempted, rather than for the outcome achieved, the AE ratings compared the actual return of the possession (i.e. the points value for the possession outcome) to its expected return (i.e. the EPV at the location where the player performed the action, $\text{EPV}_{x,y}^{\mu}$) for every action a player was involved in (Equation 6.12). The sum of the differences between actual and expected returns was divided by the median number of possessions per player team per fixture. This choice of the denominator ensured that players from teams who had more possessions within a match were not un-

duly favoured by the results. Only players completing more than 200 actions across the season were included in the analysis. Recalling the dataset $D = (x_i, y_i, s_i)_{i=1}^n$ Equation 6.12 was defined as:

$$\text{Player X's AE rating} = \frac{\sum_{i \text{ such that player X in possession}} (s_i - \text{EPV}_{(x_i, y_i)})}{\text{Player X's team median number of possessions per fixture}} \quad (6.12)$$

6.3 Results

The aim of this study was to provide a framework through which player and team performances could be evaluated in rugby league in a data-driven manner. To achieve this aim, 138 xml files provided by Opta were preprocessed into a single file of 99,966 rows and 7 columns (Table 6.2). Using this data, the study proposed a novel EPV model which produced a smooth pitch surface estimating the probability of individual possession outcomes. 33 centres (30 in the field of play, 3 in the try area) were used in a Bayesian Mixture Model approach, which applied linear and bilinear interpolation techniques to produce a smooth pitch surface. The EPV was derived from the individual possession outcome probabilities and their true point scoring values. Both the Bayesian Mixture Model approach used and the estimation of individual possession outcome probabilities are novel additions to the literature. EPV models were produced at the whole league, team attacking and team defending levels. Two novel performance metrics were introduced: an expected points scored measure which quantifies the underlying performances of teams; and an actual versus expected player rating which evaluates player performances.

A key component of the EPV model is the use of a novel Bayesian Mixture Model approach, which allows data to be aggregated in a more flexible manner than the zonal approaches previously employed in this thesis and the literature (Chapters 3 and 4; Cervone et al.(2016); Kempton et al.(2016); Singh(n.d.)). In this study, a maximum of four centres received non-zero weights for any x, y location. These centres formed the corners of a quadrilateral surrounding the x, y location of an action. They could therefore be considered analogous to the corners of a “zone” used in a Markov model. Table 6.5 provides a sample of weights for five actions from this study. In a Markov model, the weights for the four non-zero centres would each be equal to 0.250, irrespective of the x, y location of the action. In this study, the interpolation techniques used to calculate the

Table 6.5: Abbreviated example of weights used in this study. Only first 8 centres are shown, column titles represent centre coordinates. All other centres ($n = 25$) are assigned weights of 0 based on the x, y location of these example actions.

x	y	(0,-10)	(20,-10)	(35,-10)	(50,-10)	(70,-10)	(0,20)	(20,20)	(35,20)
9	4	0.293	0.240	0.000	0.000	0.000	0.257	0.210	0.000
9	6	0.257	0.210	0.000	0.000	0.000	0.293	0.240	0.000
14	11	0.090	0.210	0.000	0.000	0.000	0.210	0.490	0.000
22	13	0.000	0.202	0.031	0.000	0.000	0.000	0.664	0.102
12	12	0.107	0.160	0.000	0.000	0.000	0.293	0.440	0.000

weights allowed much more flexible weights to be produced based on the x, y location's proximity to the centre. This is shown by the $(20, -10)$ and $(20, 20)$ centre weights for the locations $(9, 4)$ and $(9, 6)$. The y coordinate half way between -10 and 20 is 5 . Both $(9, 4)$ and $(9, 6)$ are 1m away from this point in opposite directions and therefore have opposing weights for the centres $(20, -10)$ and $(20, 20)$. To this end, any probability or EPV estimates for location $(9, 4)$ were influenced to a slightly greater extent by the centre $(20, -10)$, whereas any probability or EPV estimates for location $(9, 6)$ were influenced to a slightly greater extent by the centre $(20, 20)$. It is this weighting system that allowed a smooth pitch surface to be produced.

Within the mixture model approach, the decision was made to treat centres in the field of play and try areas independently. Of the $99,966$ actions included in this study, only 91 occurred in the try area. The $\text{EPV}_{(x,y)}^\mu$ of these actions was much greater than those actions outside of the try area. For example, using centre values, the highest EPV in the field of play was at centre $(50, 100)$, where $\text{EPV}_{(50,100)}^\mu = 1.73$. All three try area centres had much greater values ($\text{EPV}_{(x,y)}^\mu \in \{3.52, 3.72, 3.16\}$). This difference justifies the separation between field of play centres and try area centres. For clarity of figure presentation, the try area values are removed from all plots for the rest of this chapter.

6.3.1 Estimation of Individual Possession Outcomes

Alongside the modelling approach, the estimation of individual possession outcomes probabilities in a single model is another novel addition to the literature. Previous approaches either aggregated the calculation of all possession outcome probabilities numer-




Figure 6.2: Whole league model mean pitch surface plot. 4pt Try and 6pt Try refer to unconverted and converted tries respectively. With the exception of EPV, probabilities are plotted. Smooth pitch surface for possession outcome probabilities is calculated using Equation 6.3 for each x, y location on the pitch. EPV for each location is calculated using Equation 6.10. Brighter areas represent higher values.

ically (Chapters 3 and 4; Cervone et al.(2016); Chan et al.(2021); Kempton et al.(2016)) or produced separate models for different possession outcomes (Decroos et al.,2019). By producing smooth pitch surfaces for each possession outcome within a single model, as well as deriving the EPV, it is possible to gain greater insights into areas where teams generate value.

Figure 6.2 depicts the mean pitch surface for the whole league data model ($P_w^\mu(s)$) for

all $s \in S$ and EPV_w^μ). There is a greater probability of points being scored by the end of the possession, the closer the location is to the opposition try area. The darker arc on the no points probability surface close to the opposition try line indicates that more points are likely to be scored from central locations than wider locations unless a player is extremely close to the try line. In general, a converted try was more likely to occur at the end of a possession than an unconverted try and a penalty goal was more likely to be scored than a drop goal. An interesting observation is the increased probability of an unconverted try on the very right hand side of the pitch close to the try line. This links with a prevailing theory within rugby league that right footed kickers (the majority of kickers in the Super League are right footed) find it harder to kick conversions from the right side of the pitch. The greater probability of drop goals and penalty goals being scored in the centre of the pitch is also notable and suggests that typically teams only attempt either type of kick in areas where they feel they are guaranteed to score points. The smooth pitch surface for each individual possession outcome produced by the model in this study allows much greater information to be extracted than the zonal EPV-19 models produced in Chapters 3 and 4 and the literature (Cervone et al., 2016; Chan et al., 2021; Kempton et al., 2016).

An advantage of adopting a Bayesian approach to the model proposed in this study is that it allows an understanding of the certainty within parameter estimates to be gleaned. Figure 6.3 provides standard deviation pitch surface plots for all possession outcome probabilities using the whole league data model ($P_w^\sigma(s)$ for all $s \in S$). Also plotted in Figure 6.3 is the whole league KDE produced in Chapter 5, which provides an understanding of the location densities across the pitch. In all possession outcome plots, there is greater variability in the wider areas of the pitch, which is accompanied by a lower density of actions in those areas in the KDE plot. This variability is particularly large for the penalty goal probabilities in wide areas, which could indicate that the result does not occur that often from these areas and that the probabilities estimated for those areas should be treated with a degree of caution. A similarly interesting observation is the increased variability in try/no try probabilities in both corners of the pitch on the opposition try line. This could be due to a number of factors, including fewer actions happening in these areas (suggested by the KDE plot) and more conversion kicks being missed from these areas after the try has been scored (suggested by the increased probability of unconverted tries on the right side of the pitch).




Figure 6.3: Whole league model EPV standard deviation plot. 4pt Try and 6pt Try refer to unconverted and converted tries respectively. Whole league KDE is reproduced from Chapter 5 to provide an understanding of location densities. Brighter colours indicate higher values.

6.3.2 Evaluation of Team Performances

To evaluate performances at a team level, individual EPV models were produced for each team's attacking and defensive data subsets. The team attacking models provided an understanding of the probabilities of a team achieving a given possession outcome at the end of their possession for any given location. The team defending model estimated the probabilities of any opposition team achieving a given possession outcome at the end

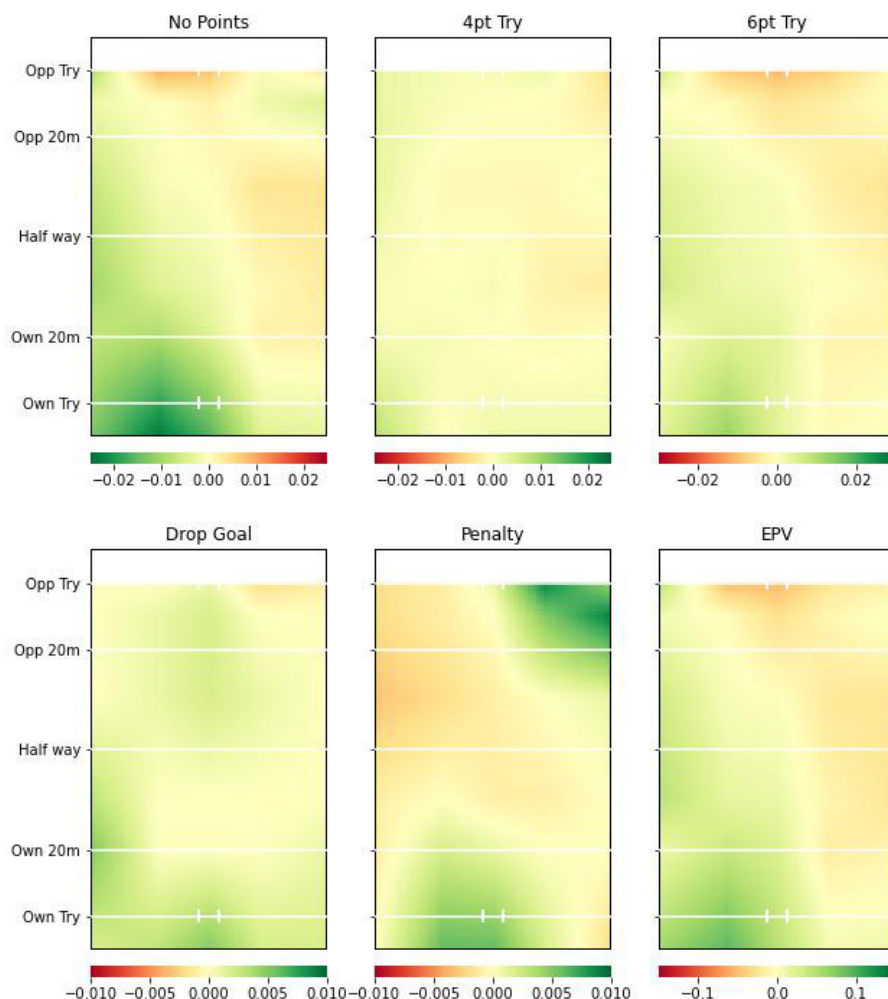


Figure 6.4: Castleford smooth pitch surface plot from team attacking model. Green areas represent higher value for more favourable event (i.e. greater probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively.

of their possession when controlling the ball at any given location against the defending team. As with the whole league model, the EPV was derived from these possession outcome probabilities (Equation 6.10). Each team's smooth EPV pitch surfaces for the attacking EPV_{att}^{μ} and defending EPV_{de}^{μ} models were used in the calculation of a novel expected points scored metric (Equation 6.11).

Similar to the whole league model, visual inspection of the smooth pitch surface plots

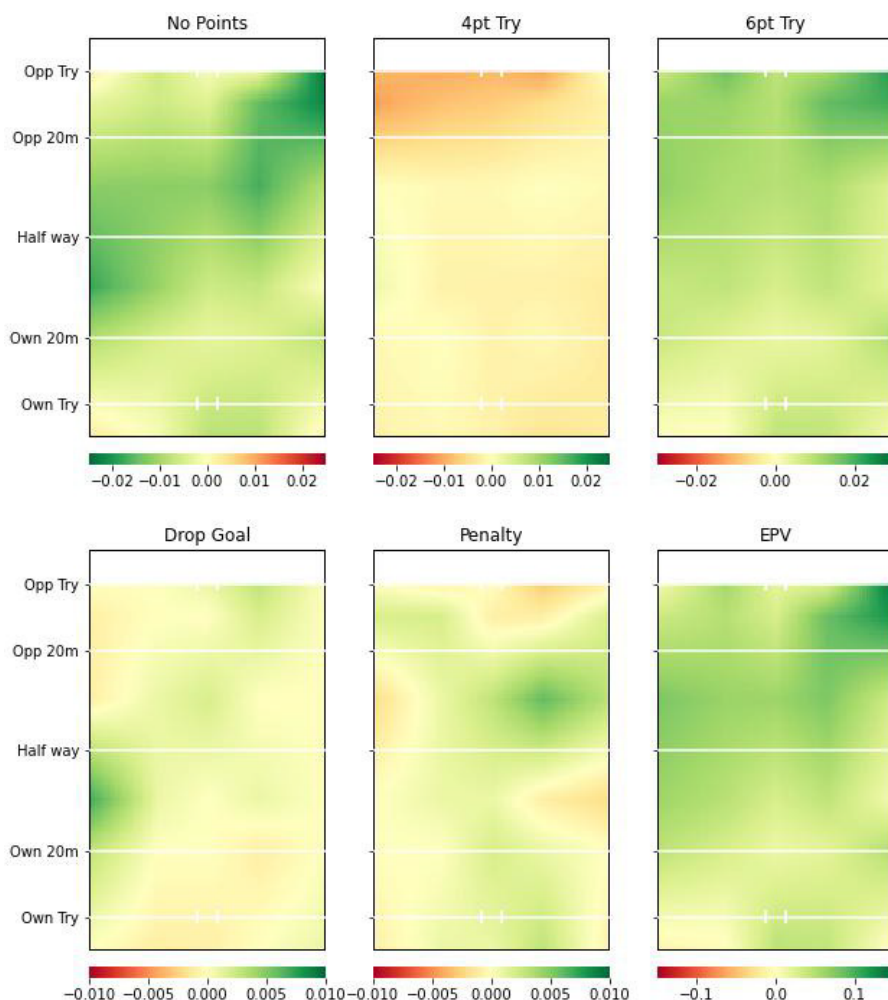


Figure 6.5: Catalans smooth pitch surface plot from team attacking model. Green areas represent higher value for more favourable event (i.e. greater probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively.

produced for each team could be used to evaluate their strengths and weaknesses across the pitch. Figures 6.4, 6.5, 6.6 show the smooth pitch surface plots for the team attacking models ($P_{att}^{\mu}(s)$ for all $s \in S$ and EPV_{att}^{μ}) of Castleford, Catalans and Wigan respectively, compared to the whole league model. Castleford (Figure 6.4) show a greater probability of scoring tries on the left side of the pitch compared to the whole league model. On the right side of the pitch, they show a much greater probability of scoring penalties, particularly in

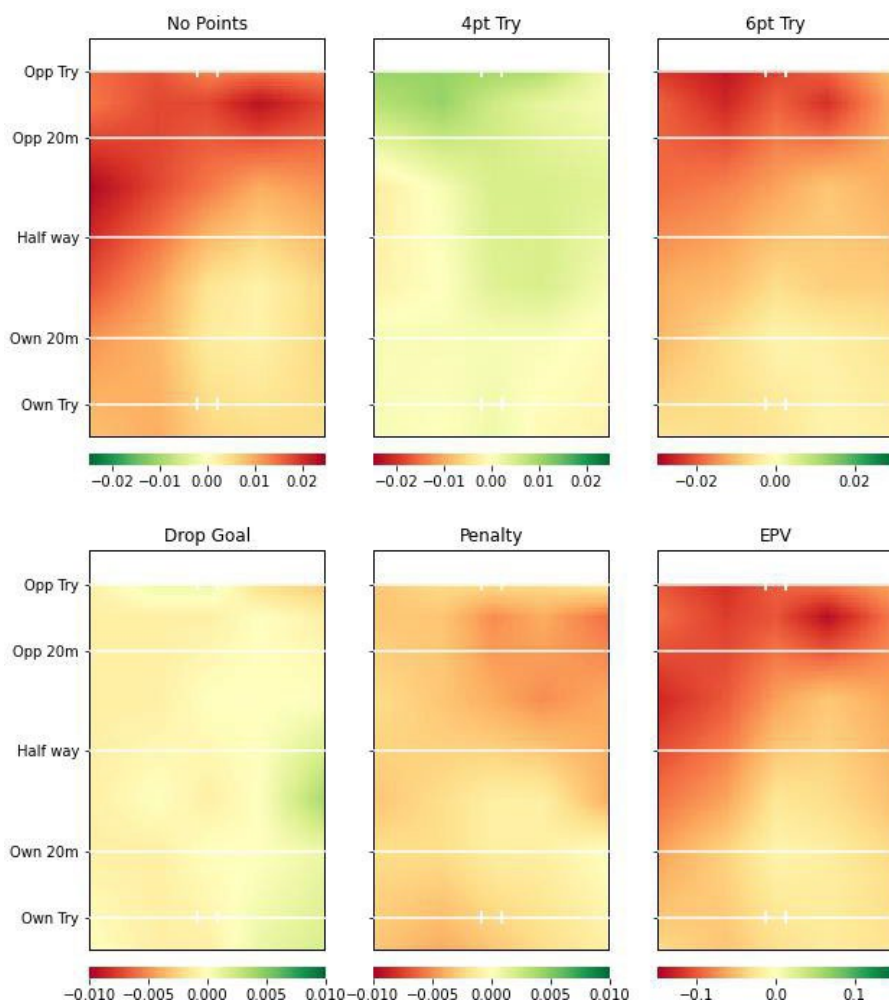


Figure 6.6: Wigan smooth pitch surface plot from team attacking model. Green areas represent higher value for more favourable event (i.e. greater probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively.

the opposition 30m, potentially providing insights into their decision making processes in different areas of the pitch. Catalans (Figure 6.5) were less likely to score an unconverted try compared to the whole league model in the opposition 20m but were more likely to score a converted try across the majority of the pitch. They were also less likely to score no points across the majority of the pitch, which is reflected in their EPV plot being greater than the whole league model in most pitch locations. Wigan (Figure 6.6) showed a greater

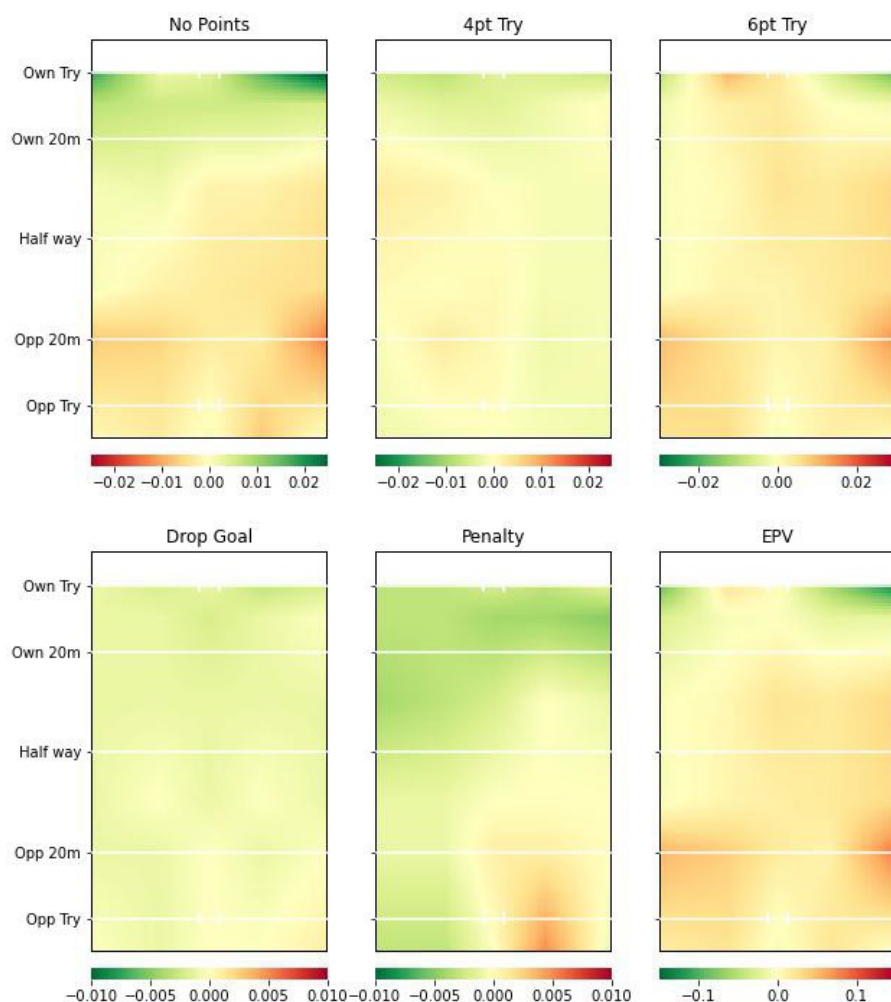


Figure 6.7: Castleford smooth pitch surface plot from team defending model. Green areas represent higher value for more favourable event (i.e. lower probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively.

probability of scoring unconverted tries or no points across the pitch in the opposition half in their model. This trend is reflected by the red (meaning reduced probability or value) areas present across the converted try, penalty and EPV plots. They were also more likely to score unconverted tries in the opposition 20m as shown by the green areas in that plot. Coaches can use this information to understand different opponents' attacking strengths and develop defensive strategies to face them.

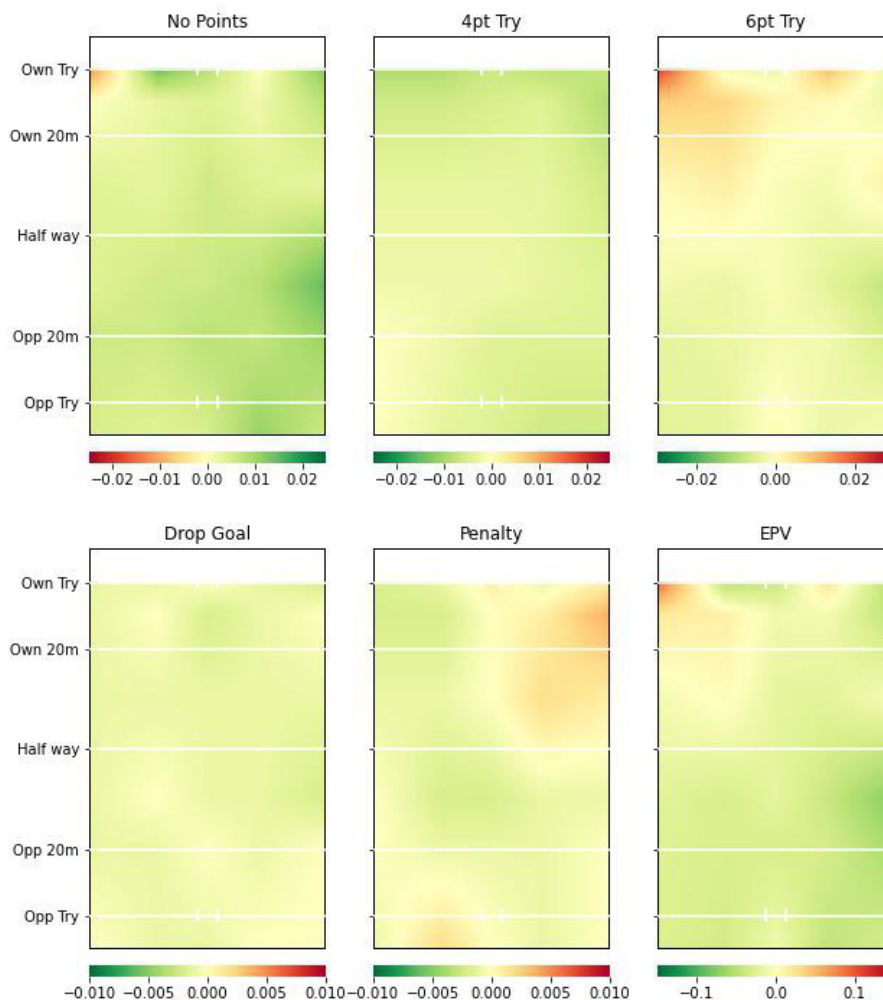


Figure 6.8: Catalans smooth pitch surface plot from team defending model. Green areas represent higher value for more favourable event (i.e. lower probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively.

Figures 6.7, 6.8, 6.9 show smooth pitch surface plots for the team defending models ($P_{\text{def}}^{\mu}(s)$ for all $s \in S$ and EPV_{def}^{μ}) for Castleford, Catalans and Wigan. These models represent the actions of all opposition teams against the nominated team. As such, the colour scheme is reversed to provide a true representation of the teams' defensive ability (i.e. whereas an increased probability of scoring no points is bad and coloured red when attacking, it is good and coloured green when defending). Castleford's plot (Figure 6.7)

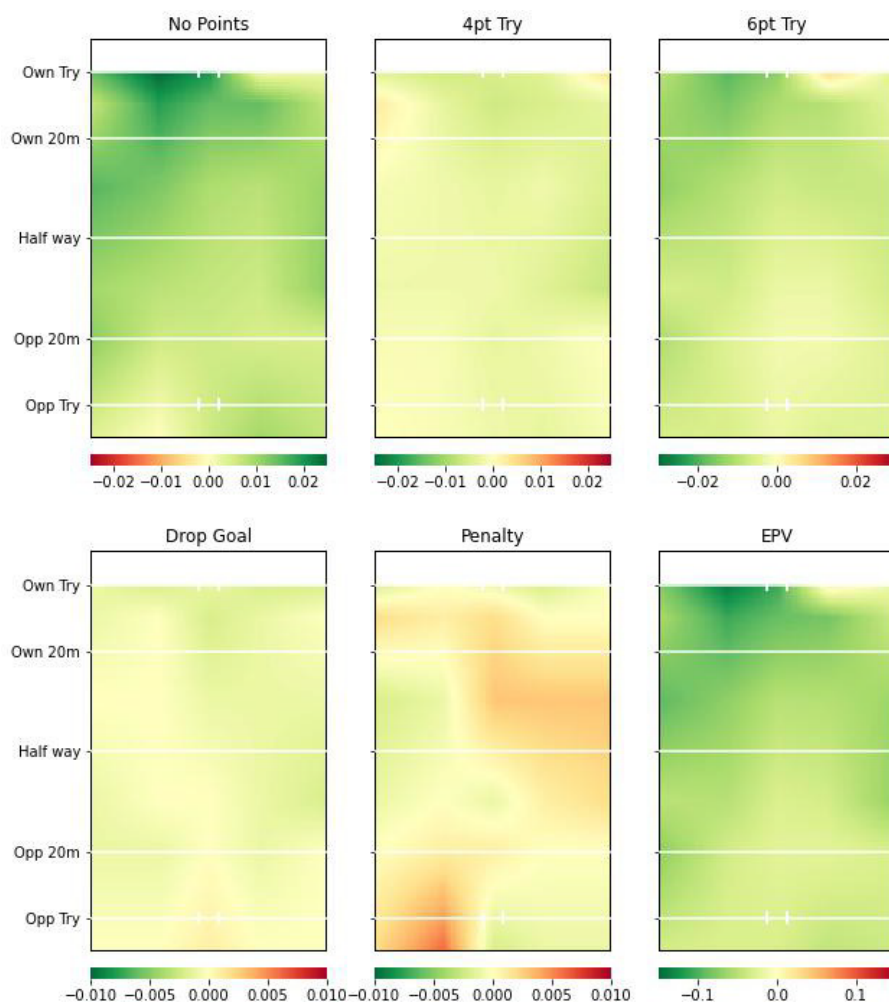


Figure 6.9: Wigan smooth pitch surface plot from team defending model. Green areas represent higher value for more favourable event (i.e. lower probability of all events occurring except No Points) compared to whole league model. 4pt Try and 6pt Try refer to unconverted and converted tries respectively.

indicated a greater probability of no points being scored with actions close to their try line and a lower chance of drop goals and penalties being scored across the pitch, but an average to greater than average chance of opposition teams scoring a converted try everywhere except on the sides of the pitch near Castleford's try line. Taken together these results potentially indicate that Castleford are susceptible to direct attacks from a greater distance as fewer actions near their try line are successful but actions outside of

their 20m are more valuable. Catalans' plot (Figure 6.8) shows large areas of better than average defence, but suggest the defence was worse than average on the left side of the pitch (their right sided defenders) in their own 20m with respect to converted tries and on the right side of the pitch where teams had a greater probability of scoring penalty goals in the Catalans half. Wigan's plot (Figure 6.9) shows above average defence with regards to both converted and unconverted tries. Furthermore, they show a much greater probability of no points being scored by opposition teams. Opposition teams have a greater than average probability of scoring penalty goals against Wigan across the majority of the pitch. In both cases, the decision taken by opposition teams to score penalty goals could be an indication that teams have been unable to penetrate the strong defensive line and score tries. As with the attacking plots, coaches can use this information to devise strategic plans to beat the opposition.

6.3.2.1 Expected Points Scored

Alongside visual inspection of attacking and defending performances, an expected points scored measure was developed to provide an understanding of the underlying attacking performances achieved by teams. Table 6.6 provides details of the expected points scored and actual points scored for each team across the 2021 Super League season. The average difference between points scored and expected points scored was 30.27 when the team's attacking EPV (EPV_{att}^{μ}) was used and 30.97 when the opposition team's defending EPV (EPV_{de}^{μ}) was used. These differences are likely due to the smoothing induced by the KDE models. Of the top 6 teams reaching the playoffs, only Wigan did so without exceeding their expected points scored by more than the average amount. Wigan were clearly the best defensive team in the league in 2021, as shown by their defensive EPV plot 6.9, and so may have been able to win matches despite their poor attacking performances as they reduced the points any opposition team was able to score. Interestingly, in the 2022 season, Wigan were the top points scorers in the league suggesting that the coaches were able to solve the problem of converting their underlying attacking performances into actual points scored.

Table 6.6: Expected points scored across the season compared to actual points scored at the team overall level. Exp Att refers to expected points scored as computed using the attacking team EPV; Exp Def refers to expected points scored as computed using the defending team EPV; Diff refers to the difference between points scored and the expected points scored (positive means more points were scored than expected). Teams are provided in league position order.

Team	Points Scored	Exp Att	Att Diff	Exp Def	Def Diff
Catalans	726	563.80	162.20	527.75	198.25
St Helens	572	492.93	79.07	480.10	91.90
Warrington	588	476.44	111.56	470.55	117.45
Wigan	387	492.82	-105.82	540.34	-153.34
Leeds	572	528.10	43.90	528.05	43.95
Hull KR	526	449.15	76.85	448.29	77.71
Castleford	413	411.72	1.28	404.35	8.65
Hull	409	412.10	-3.10	410.99	-1.99
Huddersfield	460	452.61	7.39	456.10	3.90
Wakefield	482	447.81	34.19	459.56	22.44
Salford	402	415.22	-13.22	416.78	-14.78
Leigh	356	387.10	-31.10	378.50	-22.50

6.3.3 Evaluation of Player Performances

To evaluate player performances, a novel actual vs expected (AE) player rating was proposed (Equation 6.12). The AE ratings represent the average points contribution per match for a player, above or below the expected points gained. Table 6.7 rates the top 20 players in Super League 2021. The same reference statistics are provided as in Chapter 4 (tries, try assists, metres gained and goals scored). A wider variety of positions are included in the AE ratings compared to the action impact ratings in Chapter 4, highlighting the benefit of valuing how well an action was performed, rather than valuing a player performing the action. The ratings have good face validity as one of the best players in the AE ratings was awarded the Man of Steel trophy given to the best player in the league and the Young Player of the Year for the 2021 Super League season is also in the top 20 rated players.

The model was presented to coaches at Leeds Rhinos, a team competing within the Super League. The results were widely praised. Visual inspection of the team attacking

Table 6.7: Top 20 player ratings as assessed by the AE ratings (Equation 6.12). Tries, Try Assists, Metres and Goals are provided as references of statistics currently provided for player performances. To protect anonymity, reference statistics are provided based on the whole season as: T-5 (top 5); T-10 (top 10), T-20 (top 20) and 20+ (outside top 20).

Player	Position	AE Rating	Tries	Try Assists	Metres	Goals
276	Full Back	8.21	T-20	T-5	20+	20+
19	Winger	6.67	T-5	20+	T-5	20+
6335	Stand-off	6.35	20+	20+	20+	T-5
1004	Scrum Half	6.10	20+	20+	20+	20+
433	Full Back	4.96	20+	T-10	20+	T-5
158	Winger	4.82	T-5	20+	T-10	20+
188	Centre	4.78	T-5	20+	T-5	20+
1249	Hooker	4.37	20+	20+	20+	20+
92	Scrum Half	3.91	20+	T-5	20+	20+
1281	Loose Forward	3.86	20+	20+	20+	20+
406	Loose Forward	3.82	20+	20+	20+	20+
371	Winger	3.37	T-20	20+	T-20	20+
8	Prop	3.33	20+	20+	T-10	20+
282	Second Row	3.28	20+	20+	20+	20+
20528	Winger	3.11	T-5	20+	T-5	20+
22852	Full Back	3.10	T-10	T-20	20+	20+
5	Winger	3.08	T-20	20+	20+	20+
26	Hooker	2.96	20+	20+	20+	20+
440	Loose Forward	2.88	20+	20+	20+	20+
988	Stand-off	2.81	20+	20+	20+	20+

and defending plots identified elements that they agreed with anecdotally (Wigan's poor attacking prowess and Catalan's weakness on the right side of their defence) and an element that they hadn't considered (Castleford's left sided bias). They approved of the expected points scored as a potential performance measure and pointed out Wigan's improved performance in the league in 2022 after seeing it. However, they would prefer it to be used on a match by match basis rather than across a season, which may be problematic due to the low sample sizes each match. The player ratings provided 19 names that they agreed with and one name they wouldn't have expected, but suggested was likely to be in the top 20 players at the end of the season due to the number of touches of the ball he was likely to have per possession, which was not controlled for in the rating.

Preprocessing and weights calculations were completed using bespoke Python scripts (Python 3.7, Python Software Foundation, Delaware, USA). PyMC3 v3.11.4 (Salvatier et al., 2016) was used to conduct all Bayesian analyses.

6.4 Summary

This study improved upon the work presented in Chapter 5 by proposing a novel EPV model which can be used to evaluate team and player performances in rugby league without the need for additional contextual information. A Bayesian Mixture Model provided a smooth pitch surface of possession outcome probabilities and EPV, improving upon the zonal aggregation methods previously used (Chapters 3 and 4; Cervone et al. (2016); Chan et al. (2021); Kempton et al. (2016)). Similarly, this was the first model to estimate multiple individual possession outcome probabilities in a single model rather than combining them numerically (Cervone et al., 2016; Chan et al., 2021; Kempton et al., 2016) or estimating them in separate models (Decroos et al., 2019). Two novel metrics were presented: expected points scored to evaluate underlying team performances and AE (actual vs expected) player ratings to evaluate player performances. Visual inspection of the smooth pitch surface plots was also shown to be an extremely useful method of gaining an understanding of team strengths and weaknesses, in both attacking and defensive terms. The insights from the model can be used to produce tactical strategies, monitor underlying team performances and enhance player recruitment.

The work was presented to coaches at Leeds Rhinos, who provided an impact statement for the work (Appendix D). The coaches were impressed and validated the work through their anecdotal understanding of the sport at both a team and player level. The usability of the model was praised ("the ability to do this in an efficient and reliable way, which the work provides is invaluable", "this is something I would look to embed into our daily analysis procedure"), as was the reliability - particularly with respect to unexpected findings ("there were some surprising and novel results presented, which add value to the current coach and analyst knowledge"). They suggested several future directions for the work, including valuing different actions similar to the work shown in Chapter 4 and reparameterising the model for different elements of the game, for example evaluating the metres gained by the end of a play for playmakers or the metres gained before the opposition's first tackle for kickers. It would not have been possible to make any of these

changes in the timescale of this PhD due to the low data availability, but with a further season of data, it would be possible to adapt the framework in the future.

Future work may now wish to apply the model to other sports and extend it to value actions, not just the location of the action (akin to Q-values within Markov Models). The flexibility of the model ensures it is ideally suited to invasion sports such as rugby union, football or hockey. However, whether an individual would want to use this method in sports with a single scoring possession outcome is questionable as a simpler binary model would likely perform just as well. The true benefit of using this model would be to establish three outcome measures for these sports - possession goal, no goal, opposition goal. This is similar to previous work in football ([Decroos et al., 2019](#)) but using the Bayesian Mixture Model approach described in this chapter, it would be possible to evaluate the probability of all three outcomes in a single model, rather than separate models for scoring or conceding. Similarly, the model could be extended to consider different actions and generate a form of Q-value for each action. The issue with adapting the method in this manner is the evidence based approach followed by Bayesian analytics which means that if an action hasn't been attempted in a specific location, it will take the prior value. In the case of goal kicking, this could be problematic as a penalty goal wouldn't be attempted from the possession team's 20m line (i.e. 80m away from the goal), but it could have a higher value than a penalty goal attempted from 40m out, wide on the touchline, if these penalties are missed more frequently than scored. Manipulating the priors to account for this is possible but this detracts from the methods here which have systematically allowed the models to estimate the results based on the data available.

Conclusions

In this chapter, the contributions of the four studies produced in this thesis are highlighted. The limitations of the research and its future directions are also considered. The aim of this PhD project was to develop new methodologies to understand player and team performances in rugby league using match event level data. To meet this aim, five clear objectives were defined in Section 1.3:

- O1 Investigate existing methodologies of evaluating player and team performance in team sports, focusing on their application to rugby league
- O2 Apply and adapt existing methodologies and metrics evaluating player and team performances in team sports to rugby league
- O3 Validate adapted versions of existing methodologies and metrics for usability and reliability in rugby league
- O4 Develop novel methodologies and metrics to evaluate player and team performances in rugby league
- O5 Validate proposed methodologies and metrics for usability and reliability with respect to evaluating player and team performances in rugby league

These objectives were met by different chapters within the thesis: the literature review in Chapter 2 met objective 1; Chapters 3 and 4 met objective 2 and partially met objective 3; Chapter 5 partially met objectives 2, 3, 4 and 5; and Chapter 6 met objectives 4 and 5.

The sections below outline the context surrounding the research in this thesis, the research contributions for each study and how they met the objectives described above. Limitations of the work and directions for future research are also discussed.

7.1 Rugby League

The literature review in Chapter 2 made it clear that rugby league is a unique sport with respect to the expected possession value (EPV) model literature due to its low data availability (no tracking data available, only play by play data available in 2019 season, different action by action datasets available in 2020 and 2021) and unique rules (six play attacking sets and multiple possession outcomes). These unique elements ensured it was not possible to adapt the most impressive EPV models currently available for use in rugby league. The methods which could not be used in this study included: deep learning (Fernández et al., 2021) and deep reinforcement learning approaches (Liu et al., 2020; Liu and Schulte, 2018); a Bayesian approach requiring tracking data (Cervone et al., 2016); and a CatBoost machine learning algorithm, which used 18 times more data than that available over the complete duration of this thesis (Decroos et al., 2019). However, there were some EPV models available which could be adapted and improved to evaluate player and team performances in the sport (Chan et al., 2021; Kempton et al., 2016; Singh, n.d.). These approaches used Markov Reward and Decision Processes to evaluate player and team performances but arbitrarily selected their zone sizes, which may have influenced the results they obtained.

At the onset of this PhD, no previous studies had used action by action data to evaluate player and team performances in rugby league. Two studies had considered play by play data (Holbrook et al., 2019; Kempton et al., 2016) but neither had provided a methodology or metric through which player performances could be evaluated. As such, there was no framework through which the results of this thesis could be compared. Given this background, and the limited data available in the sport, the focus of this thesis was to continually improve the methodologies used rather than optimise earlier algorithms and models. Theoretically, it would have been possible to use all three seasons of data together to optimise the work from earlier chapters, but this would have been problematic from a preprocessing perspective due to the different action definitions across the three seasons. Consequently, each model was developed using slightly different data and could only be

compared with real performances and coaches' feedback for validation purposes. The decision to develop the methodologies in this way was vindicated by the final Bayesian EPV model, which could not have been developed if the 2021 data was used to further extend the Markovian models of Chapters3and4. The Markov models of Chapters3 and4only received partial validation by experts and there was no clear method through which they could be improved within the constraints of this thesis. Conversely, the final Bayesian EPV model, proposed in Chapter6represented a significant improvement on the Markovian approach and received complete validation from professional experts.

7.2 Research Contributions

This section outlines the four research studies completed within the thesis. An overview of each study is given and the key contributions are detailed. The extent to which the study meets the objectives defined in Section1.3is also evaluated.

7.2.1 Markov Reward Processes for the Quantification of Team Attacking Performances

In Chapter3, the works ofKempton et al.(2016) in rugby league andSingh(n.d.) in football were applied, adapted and extended. Markov Reward Processes (MRP) were used to estimate the value of different locations on the pitch with respect to the points outcome at the end of the possession using play by play data. Six different state spaces were considered within the MRP framework, two with fixed zone sizes (EPV-308 and EPV-77) and four with statistically computed zone sizes (EPV-37, EPV-19, EPV-13 and EPV-9). The Kullback-Leibler Divergence was used to empirically evaluate the models' ability to reproduce future team attacking performances, identifying the EPV-19 as the most suitable model for use within rugby league. Z-score analysis was proposed as a novel method of evaluating team attacking performances across the 2019 Super League season. The study was published in PLOS One and presented to the England Rugby League Performance Unit.

The study made two key contributions to the literature. First, it applied and adapted the previous literature (Kempton et al.,2016;Singh,n.d.) to a new dataset within rugby league. The work was adapted by empirically evaluating different zone sizes with respect

to their ability to reproduce future team attacking performances. This work had never previously been completed and identified the EPV-19 as the most suitable model for use in rugby league. This is because it provided the best tradeoff between representation of future match performances and granularity of results. Second, the work of [Kempton et al. \(2016\)](#) was extended by proposing a novel method through which team attacking performances could be evaluated using z-score analysis of the EPV-19 results. This analysis allowed a heat map to be produced, identifying areas on the pitch which were valuable to different teams. It could be used as a starting point for the development of tactical strategies against different oppositions.

By applying, adapting and extending previous literature ([Kempton et al., 2016](#); [Singh, n.d.](#)), Chapter 3 met objective 2 of the project. However, after presenting the results to the England Rugby League Performance Unit though, it was clear that the model only partially met objective 3. The unit thought the spatial insights produced by the heat map were reliable and excellent, but questioned the long-term usability of the model due to its use of play by play rather than action by action data.

7.2.2 Markov Decision Processes for the Evaluation of Player and Team Performances

In Chapter 3, spatial data were only available at a play by play level, which restricted the analyses that could be conducted; for Chapter 4, action by action spatial data became available. Therefore, Chapter 4 extended the work of Chapter 3 by using Markov Decision Processes to estimate action values using the EPV-19 zones, applying a previous idea in football ([Singh, n.d.](#)). Action impact ratings were used to evaluate player performances, adapting previous research in ice hockey ([Routley and Schulte, 2015](#)). Context nodes were also adapted from ice hockey ([Routley and Schulte, 2015](#)) and football ([Decroos et al., 2019](#)), providing further detail surrounding player performances and enabling individual team MDPs to be produced. The work was presented at UKCI and subsequently published in *Advances in Intelligent Systems and Computing*. It was also presented internally to the Sports Science Department at Leeds Beckett University.

The key contribution of this study was that it extended previous work in rugby league (Chapter 3; [Kempton et al. \(2016\)](#)) by providing the first action by action analysis of player and team performances in rugby league. This allowed much greater detail to be

provided within the evaluation of player and team performances, resulting in much more usable results. Using the EPV-19 model identified in Chapter3, action values were generated for every zone on the pitch, allowing their value to incorporate spatial context for the first time in rugby league. In addition to this, the study adapted ideas from ice hockey (Routley and Schulte,2015) and football (Decroos et al.,2019;Singh,n.d.) and validated their use within rugby league. The adaptation of action impact ratings (Routley and Schulte,2015) to produce player ratings was evaluated. Furthermore, context nodes were adapted from ice hockeyRoutley and Schulte(2015) and football (Decroos et al., 2019), allowing further insights to be produced with respect to these player ratings (e.g. understanding if a player performs better when their team is losing). Context nodes were also used to produce individualised team state, action values. These values provided insight into different playing styles between teams, which could be used to produce tactical strategies against future opponents.

Although producing more advanced results than Chapter3, the work conducted in this chapter still constituted the application and adaptation of previous methodologies and metrics, so like Chapter3, it met the requirements of objective 2. After consultation with the Sports Science Department at Leeds Beckett University, it was agreed that the results of this study improved upon those of Chapter3, particularly with respect to the team performance insights, which were considered both usable and reliable. However, the player ratings' reliability was questioned. Together, the results of Chapters3and4completely fulfil objectives 2 and 3 of this thesis by applying, adapting and extending previous literature. However, both studies were limited by their zone-based analyses, which reduced their reliability and usability in practice. It was therefore clear that adopting approaches which provided smooth pitch surfaces was necessary when novel methodologies and metrics were developed, to maximise their reliability and usability in practice.

7.2.3 Kernel Density Estimation and Wasserstein Distance Evaluation of the Spatial Trends of Team Attacking Performances

The key issue with the zone-based analyses conducted in Chapters3and4was the size of the zones. This was shown by the player ratings systems, where a player could run 59m towards the opposition try line and receive the same action value as a player who was forced 10m backwards if both players started and ended in the same zone. The logical

solution to this valuation problem was to provide a smooth pitch surface, which valued every location on the pitch differently. Chapter 5 began this process by simplifying the research question to consider only the location of actions on the pitch, not the value generated by the action or location. Adapting ideas previously considered at a player level in American football (Mallepalle et al., 2020), Kernel Density Estimation was used to provide a smooth pitch surface describing the probability of a team controlling the ball in any given location. By using a standardised bandwidth at three levels of data (whole league, team overall and team-opponent subsets), it was possible to provide direct comparisons regarding the spatial trends of team attacking performances between and within teams for the first time. The work of Mallepalle et al. (2020) was extended by developing two novel metrics from the Wasserstein distance to provide these comparisons: the normalised axis Wasserstein distance and the directional Wasserstein distance. Together, these metrics were able to visualise the direction or axis along which teams differed from each other in simple quadrant plots, which could be used for benchmarking purposes. The work was presented at UKCI and to coaches at Leeds Rhinos Rugby League club who were extremely complimentary and provided an impact statement surrounding its use in practice (Appendix D).

The results of the study provide two key contributions to the literature. First, the work of Mallepalle et al. (2020) was applied and adapted to provide the first implementation of KDEs to quantify the spatial trends of team attacking performances with a smooth pitch surface in any sport. Second, the work was extended to provide novel metrics through which differences in the spatial trends could be evaluated and visualised. To do this, the Wasserstein distance was used to compare KDEs for the first time in sport. Manipulating the Wasserstein distance provided novel insights into the differences along each axis (i.e. does a team differ from the league average in terms of pitch width or height?) via the normalised axis Wasserstein distance, and directional differences (i.e. does a team perform more actions on the left side of the pitch or near their own try line?) via the directional Wasserstein distance. These insights provide valuable benchmarking tools for clubs and have significant potential to improve tactical strategy preparations, thus providing assistance to coaches in both long and short-term decision making.

By adapting and extending the work of Mallepalle et al. (2020) to quantify the spatial trends of team attacking performances, this chapter met objective 2 of the thesis; by developing novel metrics from the Wasserstein distance through which differences in these

spatial trends of attacking performances could be quantified and visualised, objective 4 of the thesis was partially met. Objectives 3 and 5 were partially met when coaches at Leeds Rhinos provided validation of both the usability and reliability of the results (Appendix D). In all cases, the objectives were partially met because, in this simplified analysis, only team performances were considered.

7.2.4 A Bayesian Approach to the Evaluation of Team and Player Performances

Building upon the success of Chapter 5, Chapter 6 produced a novel EPV model with a smooth pitch surface in rugby league. Previous approaches providing smooth surfaces (Fernández et al., 2021; Liu et al., 2020) could not be adapted for use in rugby league due to data availability issues. To overcome these issues, a novel Bayesian Mixture Model approach was proposed, whereby the weights for each centre were devised using linear and bilinear interpolation techniques. The model estimated individual possession outcome probabilities and used them to derive an EPV measure. Models were produced at the whole league, team attacking and team defending levels allowing visual inspection of team strengths and weaknesses with respect to different point scoring outcomes. Two novel metrics were produced from the model: expected points scored, which combined the results of this study with the results of Chapter 5 to evaluate team performances, and actual versus expected player ratings, which were used to evaluate player performances. The work was presented to coaches at Leeds Rhinos Rugby League club who were impressed and provided an impact statement surrounding its use in practice (Appendix D).

The key contribution of this study was the development of a novel EPV model using an approach not previously considered in any sport. The Bayesian Mixture Model approach used linear and bilinear interpolation techniques to produce a unique set of weights for 33 centres around the pitch based on their proximity to each x, y location on the pitch. These weights were used to estimate the probability of individual possession outcomes for every x, y location on the pitch producing a smooth pitch surface. The ability to estimate individual possession outcomes in the same model for the first time is also an extremely important addition to the literature. A secondary contribution of the study was the development of two novel performance metrics. Expected points scored combined the results of this study with the results of Chapter 5 to evaluate the underlying attacking

performances of teams. Actual vs expected player ratings highlighted the players whose actions contributed to their team scoring points in matches. Both the model and the results have significant ability to influence the decisions made by coaching staff regarding performance evaluation, tactical strategy development and player recruitment.

The completely novel Bayesian Mixture Model described by this study allowed for the estimation of individual possession outcomes in the same model, whilst also producing a smooth pitch surface. This novel methodology, alongside the novel metrics devised from its results ensure this chapter completely met objective 4 of this project. Furthermore, the positive feedback from coaches after the work had been presented ensured that it also fully met objective 5 of this PhD project, providing usable and reliable results for both team and player performances, which could significantly impact long and short-term decision making at rugby league clubs.

7.2.5 Summary of contributions

Chapters 2 to 6 met all objectives of this thesis. However, although Chapters 3 and 4 were shown to meet objective 3 (validation of existing methods and methodologies), it was clear that limitations within the existing methodologies and the data available within the sport limited their reliability and usability in practice. The key issue with these studies was their zonal approach to analysis so novel methodologies and metrics were developed in Chapters 5 and 6. These methodologies used a smooth pitch surface to evaluate player and team performances, providing significantly more usable and reliable results.

Of the four models produced in this thesis, the EPV model described in Chapter 6 is the most comprehensive. The Bayesian Mixture Model approach improved upon the zonal approaches previously adopted in the literature (Chapters 3 and 4; Cervone et al.(2016); Kempton et al.(2016); Singh(n.d.)) to provide a smooth pitch surface. The smooth pitch surface enables team and player performances to be evaluated in much finer detail than ever before in rugby league. Furthermore, the model estimated individual possession outcome probabilities in a single model for the first time in the literature providing significant tactical insight into locations where teams may be more likely to score different points. Previous approaches have either aggregated possession outcomes numerically, losing this information (Chapters 3 and 4; Cervone et al.(2016); Chan et al.(2021); Kempton et al.(2016)) or estimated them in separate models (Decroos et al.,2019), which is an unsuit-

able approach within rugby league given the five possible possession outcomes for each team. Finally, the novel metrics produced from the Bayesian EPV model provide significantly greater insights than the first two studies of this thesis. The expected points scored metric, which utilises the KDE results from Chapter 5, measures underlying team performances, incorporating the spatial context of the match and the attacking or defending team's abilities into a single value. Utilising this metric alongside the novel Wasserstein distance metrics from Chapter 5 provides significant insight into team performances. Similarly, the actual vs expected player ratings produced from the Bayesian EPV model were shown to be much better than those produced in Chapter 4 as they considered the outcome of the possession, rather than singling out those players who frequently attempted high value actions, regardless of how successful they were.

7.3 Limitations

Throughout the duration of this process, the limitations of the models produced were discussed and used to help determine the direction of future studies. This section outlines the key limitations of each model and discusses the issues researchers may wish to consider if applying, adapting or extending these models in the future.

The overarching limitation of the thesis was data availability, which dictated the scope of the models proposed and their usefulness. Although other studies in the literature were able to access multiple millions of data points (Cervone et al., 2016; Decroos et al., 2019; Fernández et al., 2021; Liu et al., 2020), rugby league only has small data sets (approximately 100,000 data points per season). Indeed, in Chapter 3, action by action data wasn't currently available. The gold standard for EPV is a model which can employ player tracking and action by action match event level data (Cervone et al., 2016; Fernández et al., 2021), but if that is not available a more suitable modelling method should be chosen and the limitations of the method should be understood. For example, in Chapter 3, it was possible to evaluate team performances based on the sequence of location movements for each play-the-ball, but the lack of action by action data ensured it was not possible to evaluate player performance on an action by action basis like in other chapters within the thesis. Researchers and practitioners should be aware of this limitation when deciding which model to use within their context.

The key limitation of the Markov models employed in Chapters 3 and 4 was the zones

used to establish state values. Adopting a zonal system has multiple benefits, including ease of computation and aggregation of data, which can increase confidence in the results. However, the size of the zones is important. Although the EPV-19 model provided zone sizes, which were best able to replicate future team attacking performances, this came at a cost of extremely large zone sizes. Extremely large zone sizes may be useful when trying to understand the spatial element of a team's attacking performances, but when attempting to evaluate individual player performances based on the actions taken, they did prove problematic. This is because a player needing to move 60m forwards to move to a different zone is an unrealistic expectation to obtain a small performance gain. The zone sizes used in basketball (Cervone et al., 2016) were much smaller because the court is much smaller and therefore were able to represent the value of actions much more accurately than those utilised in Chapters 3 and 4 of this thesis. Researchers should be cautious about the zone sizes they use and understand the influence this can have on their results.

In Chapter 5, the key limitation of using KDE and Wasserstein distance analysis was the lack of consideration given to the value of performing an action in a specific location. As an initial movement towards adopting a smooth pitch surface, this was an accepted limitation from the outset of the study but the method developed provided results which could be used alongside those of Chapter 6 to provide an expected points scored team performance metric. There is also significant scope to consider these probabilistic models to understand player or team tendencies where value gained is not the most important element. For example, understanding the tendency of a tennis player to hit the ball to different locations with his/her serve could also be considered with this method and would be extremely valuable for the tactical considerations of a tennis player when preparing for upcoming matches. Although this thesis is predominantly based around valuing player and team performances through the use of EPV models, the work presented in Chapter 5 is extremely valuable in practical terms. Researchers and practitioners should be aware that methods that understand the tendencies of teams or players can also provide useful, actionable insights.

The Bayesian Mixture Model considered in Chapter 6 provides a comprehensive model within the constraints of the data available, which was able to improve upon the limitations of previous models. However, by adopting an evidence based approach to parameter estimation, it could be limited by the prior distributions set. This is because in

a different analysis, if no evidence is provided, the estimate will not shift from the prior distribution, which could result in an unrealistically high or low parameter value. This wasn't an issue within the location based model proposed in Chapter 6, but if extending the model to include actions in the same hierarchical manner it could be. Using a global prior of this type for goal kicks for example could result in kicks which take place behind the half way line being valued higher than those in front of the half way line. This situation could occur if there is evidence that those in front of the half way line are missed with a greater probability than the prior distribution suggests, but there is no evidence of that for those behind the halfway line because they are not attempted. Researchers and practitioners should be aware of the important contextual factors that can influence the prior distribution in these situations.

7.4 Future Work

The work presented in this thesis provides a valuable enhancement to the current literature considering EPV models in sports, as well as being completely novel within rugby league. It has been well received by sports practitioners and within the research community (as evidenced by the publication and presentation record). However, there are still opportunities to improve it. Two key areas where future work could be focused are: valuation of actions and adaptation to other outcomes or sports.

FW1 Valuation of actions refers to evaluating the definition of a successful or useful action. This thesis has focused on the points scored by the end of the possession (whether aggregating in the Markovian approaches or estimating probabilities individually in the Bayesian approach) as the definition of successful actions. This decision was taken as points are the most important outcome of possessions within rugby league - the more points a team scores or a player contributes to, the more likely the team is to win matches and be successful. However, not all players' key role is to score points. For example, the key role of props and second rows is to gain metres early in the attacks to progress the team into positions where they can score tries (Table 2.3). The majority of actions taken by these players occur in the possession team's half meaning it is unlikely that they will produce many tries so using the actual vs expected performance ratings identified in Chapter 6 for these players is

an unfair point scoring system. A better method of evaluating these players may be to change the possession outcome probabilities to an estimation of metres gained. Similar comments could apply to the half pair and hooker, who act as playmakers and so perform an action in virtually every play of their team's possession (Table 2.3). A more appropriate action valuation for these players could be speed of play, or a combination of metres gained and EPV dependent on the location on the pitch. The flexibility of the Bayesian methodology employed in Chapter 6 would allow any of these changes to be made as desired by coaches.

FW2 Adaptation to other outcomes or sports refers to whether the model can consider other tactical elements or be adapted to other sports. The former of these considerations could be extremely interesting - for example, can the probabilities of different actions be considered based on the location of the ball and the previous action? This information would be extremely valuable to coaches and is provided in sports with much higher data availability (Fernández et al., 2021), but it's not immediately clear how the EPV model in Chapter 6 could be adapted to provide this information. The second consideration refers to whether the Bayesian model could be adapted to other sports. The flexibility of the framework means that it could, but it wouldn't always be appropriate. For example, the approach could easily be used in rugby union or American football, which follow similar episodic approaches to rugby league and have multiple point scoring outcomes. However, in sports such as football or netball, where only a binary possession outcome is present (i.e. goal or no goal), it may not be necessary to adopt the EPV model proposed in Chapter 6 unless the definition of possessions is changed to incorporate the next scoring action, rather than the end of the current attacking possession. In this situation, the model could estimate the probabilities of no score (i.e. half/full time), possession team score or opponent score in a single model. Adapting this model to football may provide a better solution to the current method of estimating probabilities for both teams, which uses individual models for scoring/conceding goals (Decroos et al., 2019). An alternative method through which the model could be applied to other sports is to gain tactical insights regarding specific possession ending actions from any given location on the pitch. For example, in its current format it could be used to evaluate the probability of a team shooting, crossing or being tackled at the end

of its possession conditional on its location on the pitch. Similarly, in tennis it could be used to evaluate the probability of players selecting a specific type of shot (e.g. top spin forehand, top spin backhand, forehand drop shot, forehand slice, backhand drop shot, backhand slice etc.) conditional on their position on the court.

References

- Barker, M. (2016). The nba is wisely opening up its sportvu data to fans and media outlets.
URL: <https://www.sporttechie.com/the-nba-is-wisely-opening-up-its-sportvu-data-to-fans-and-media-outlets/>
- Bellman, R. (1954). The theory of dynamic programming, *Bulletin of the American Mathematical Society* **60**(6): 503–516.
- Bishop, C. M. (1994). Mixture density networks.
- Bowman, A. W. (1984). An alternative method of cross-validation for the smoothing of density estimates, *Biometrika* **71**(2): 353–360.
- Bransen, L. and Haaren, J. V. (2020). Player chemistry: Striving for a perfectly balanced soccer team, *Proceedings of the MIT Sloan Sports Analytics Conference 2020*.
- Bransen, L., Robberechts, P., Haaren, J. V. and Davis, J. (2019). Choke or shine? quantifying soccer players’ abilities to perform under mental pressure, *Proceedings of the MIT Sloan Sports Analytics Conference 2019*.
- Bukiet, B., Harold, E. and Palacios, J. L. (1997). A markov chain approach to baseball, *Operations Research* **45**(1).
- Cervone, D., D’Amour, A., Bornn, L. and Goldsberry, K. (2016). A multiresolution stochastic process model for predicting basketball possession outcomes, *Journal of the*

REFERENCES

- American Statistical Association* **111**(514): 585–599.
URL: <http://dx.doi.org/10.1080/01621459.2016.1141685>
- Chan, T. C. Y., Fernandes, C. and Puterman, M. L. (2021). Points gained in football: Using markov process-based value functions to assess team performance, *Operations Research* **69**(3): 877–894.
- Chen, S. (2015). Optimal bandwidth selection for kernel density functionals estimation, *Journal of Probability and Statistics* .
- Chiu, T., Fang, D., Chen, J., Wang, Y. and Jeris, C. (2001). A robust and scalable clustering algorithm for mixed type attributes in large database environments, *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '01*, ACM Press, New York, NY, pp. 263–268.
- Cohen, B. (2016). The golden state warriors have revolutionized basketball.
URL: <https://www.wsj.com/articles/the-golden-state-warriors-have-revolutionized-basketball-1459956975>
- Croft, H., Wilcox, B. and Lamb, P. (2017). Using performance data to identify styles of play in netball: An alternative to performance indicators, *International Journal of Performance Analysis in Sport* **17**(6): 1034–1043.
- Dalton-Barron, N., Palczewska, A., Weaving, D., Rennie, G., Beggs, C., Roe, G. and Jones, B. (2022). Clustering of match running and performance indicators to assess between- and within-playing position similarity in professional rugby league, *Journal of Sports Sciences* **40**(15): 1712–1721.
- Dalton-Barron, N., Whitehead, S., Roe, G., Cummins, C., Beggs, C. and Jones, B. (2020). Time to embrace the complexity when analysing gps data? a systematic review of contextual factors on match running in rugby league, *Journal of Sports Sciences* **38**(10): 1161–1180.
- Decroos, T., Bransen, L., Haaren, J. V. and Davis, J. (2019). Actions speak louder than goals: Valuing player actions in soccer, *25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*.

REFERENCES

- Dobrushin, R. L. (1970). Prescribing a system of random variables by conditional distributions, *Theory of Probability & Its Applications* **15**(3): 458–486.
- Fernández, J. and Bornn, L. (2018). Wide open spaces: A statistical technique for measuring space creation in professional soccer, *12th MIT Sloan Sports Analytics Conference*.
- Fernández, J., Bornn, L. and Cervone, D. (2019). Decomposing the immeasurable sport: A deep learning expected possession value framework for soccer, *13th MIT Sloan Sports Analytics Conference*.
- Fernández, J., Bornn, L. and Cervone, D. (2021). A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions, *Machine Learning* **110**: 1389–1427.
- Fisher, R. A. (1925). *Statistical Methods for Research Workers*, Oliver and Boyd, Edinburgh, UK.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N. T., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D. J., Tavenard, R., Tong, A. and Vayer, T. (2021). Pot: Python optimal transport, *Journal of Machine Learning Research* **22**(78): 1–8.
- Fleming, C. H., Fagan, W. F., Mueller, T., Olson, K. A., Leimgruber, P. and Calabrese, J. (2015). Rigorous home estimation with movement data: A new autocorrelated kernel density estimator, *Ecology* **96**(5): 1182–1188.
- Gabbett, T. (2010). Gps analysis of elite women’s field hockey training and competition, *Journal of Strength and Conditioning Research* **24**(5): 1321–1324.
- Glassbrook, D. J., Doyle, T. L., Alderson, J. A. and Fuller, J. T. (2019). The demands of professional rugby league match-play: a meta-analysis, *Sports Medicine - Open* **5**(1): 1–20.
- Hardin, J. and Hilbe, J. (2003). *Generalized Estimating Equations*, Chapman and Hall, London, UK.

REFERENCES

- Harrop, K. and Nevill, A. (2014). Performance indicators that predict success in an english professional league one soccer team, *International Journal of Performance Analysis in Sport* **14**(3): 907–920.
- Hastie, T. and Tibshirani, R. (1986). Generalised additive models, *Statistical Science* **1**(3): 297–318.
- Hedeker, D. and Gibbons, R. D. (1994). A random-effects ordinal regression model for multilevel analysis, *Biometrics* **50**(4): 933–944.
- Herold, M., Kempe, M., Bauer, P. and Meyer, T. (2021). Attacking key performance indicators in soccer: Current practice and perceptions from the elite to youth academy level, *Journal of Sports Science and Medicine* **20**(1): 158–169.
- Holbrook, M., Hobbs, J. and Lucey, P. (2019). Rugby-bot: Utilizing multi-task learning amp; fine-grained features for rugby league analysis.
- Hothorn, T., Hornik, K. and Zeileis, A. (2006). Unbiased recursive partitioning: A conditional inference framework, *Journal of Computational and Graphical Statistics* **15**(3): 651–674.
- Howard, R. A. (1971). *Dynamic Probabilistic Systems, Vol II: Semi-Markov and Decision Processes*, Wiley, New York, NY.
- Hughes, C. (1990). *The Winning Formula*, Collins, London, UK.
- Hughes, M., Franks, I. and Dancs, H. (2019). *Essentials of Performance Analysis in Sport*, third edn, Routledge, New York, NY.
- Irvine, S. and Kennedy, R. (2017). Analysis of performance indicators that most significantly affect international twenty20 cricket, *International Journal of Performance Analysis in Sport* **17**(3): 350–359.
- Jobson, J. D. (1991). *Applied Multivariate Data Analysis*, Springer, New York, NY.
- Jolliffe, I. T. (2002). *Principal Component Analysis*, Springer, New York, NY.
- Kass, G. V. (1980). An exploratory technique for investigating large quantities of categorical data, *Applied Statistics* **29**(2): 119–127.

REFERENCES

- Kempton, T., Kennedy, N. and Coutts, A. (2016). The expected value of possession in professional rugby league match-play, *Journal of Sports Science* **34**: 645–650.
- King, T. L., Bentley, R. J., Thornton, L. E. and Kavanagh, A. M. (2016). Using kernel density estimation to understand the influence of neighbourhood destinations on bmi, *BMJ Open* **6**(2).
- Kullback, S. and Leibler, R. (1951). On information and sufficiency, *The Annals of Mathematical Statistics* **22**(1): 79–86.
- Langlois, T. J., Fitzpatrick, B. R., Fairclough, D. V., Wakefield, C. B., Hesp, S. A., McLean, D. L., Harvey, E. S. and Meeuwig, J. J. (2012). Similarities between line fishing and baited stereo-video estimations of length-frequency: Novel application of kernel density estimates, *PLoS One* **7**(11): 1–9.
- Lewis, M. (2003). *Moneyball: The art of winning an unfair game*, W. W. Norton and Company, New York, NY.
- Lichman, M. and Smyth, P. (2014). Modeling human location data with mixtures of kernel densities, *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, New York, NY, USA, p. 35–44.
- Liu, G., Luo, Y., Schulte, O. and Kharrat, T. (2020). Deep soccer analytics: Learning an action-value function for evaluating soccer players, *Data Mining and Knowledge Discovery* **34**: 1531–1559.
- Liu, G. and Schulte, O. (2018). Deep reinforcement learning in ice hockey for context-aware player evaluation, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, IJCAI-18, pp. 3442–3448.
URL: <https://doi.org/10.24963/ijcai.2018/478>
- Mallepalle, S., Yurko, R., Pelechrinis, K. and Ventura, S. (2020). Extracting nfl tracking data from images to evaluate quarterbacks and pass defenses, *Journal of Quantitative Analysis in Sports* **16**(2): 95–120.
- Marron, J. (1989). Comments on a data based bandwidth selector, *Computational Statistics and Data Analysis* **8**(2): 155–170.

REFERENCES

- McArdle, W. D., Katch, F. I. and Katch, V. L. (2015). *Essentials of Exercise Physiology*, fifth edn, Lippincott Williams and Wilkins, Baltimore, MD.
- Menard, S. (2009). *Logistic Regression: From introductory to advanced concepts and applications*, Sage Publications Inc, Thousand Oaks, CA.
- Merhej, C., Beal, R., Matthews, T. and Ramchurn, S. (2021). What happened next? using deep learning to value defensive actions in football event-data, *27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21)*.
- Murphy, K. and Myors, B. (1999). Testing the hypothesis that treatments have negligible effects: Minimum effect tests in the general linear model, *Journal of Applied Psychology* **84**: 234–238.
- Opta (n.d.). Super league player stats, <https://www.superleague.co.uk/stats/player-stats>. Accessed: 2022-10-29.
- Panaretos, V. M. and Zemel, Y. (2019). Statistical aspects of wasserstein distances, *Annual Review of Statistics and Its Application* **6**(1): 405–431.
- Pappalardo, L., Cintia, P., Rossi, A., Massucco, E., Ferragina, P., Pedreschi, D. and Giannotti, F. (2019). A public data set of spatio-temporal match events in soccer competitions, *Scientific Data* **6**(236).
- Park, B. and Marron, J. (1990). Comparison of data-driven bandwidth selectors, *Journal of the American Statistical Association* **85**: 66–72.
- Parmar, N., James, N., Hearne, G. and Jones, B. (2018). Using principal component analysis to develop performance indicators in professional rugby league, *International Journal of Performance Analysis in Sport* **18**: 938–949.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research* **12**: 2825–2830.

REFERENCES

- Puente, C., Del Cose, J., Salinero, J. J. and Abian-Vicen, J. (2015). Basketball performance indicators during the acb regular season from 2003 to 2013, *International Journal of Performance Analysis in Sport* **15**(3): 935–948.
- Robertson, S., Gupta, R. and McIntosh, S. (2016). A method to assess the influence of individual player performance distribution on match outcome in team sports, *Journal of Sports Sciences* **34**(19): 1893–1900.
- Rodgers, J. L. and Nicewander, W. A. (1988). Thirteen ways to look at the correlation coefficient, *The American Statistician* **42**(1): 59–66.
- Rosenblatt, M. (1956). Remarks on some nonparametric estimates of a density function, *Annals of Mathematical Statistics* **27**: 832–837.
- Routley, K. and Schulte, O. (2015). A markov game model for valuing player actions in ice hockey, *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, UAI'15, AUAI Press, Arlington, Virginia, USA, pp. 782–791.
- Salvatier, J., Wiecki, T. V. and Fonnesbeck, C. (2016). Probabilistic programming in python using PyMC3, *PeerJ Computer Science* **2**: e55.
URL: <https://doi.org/10.7717/peerj-cs.55>
- Sarmiento, H., Marcelino, R., Anguera, M., Campanico, J., Matos, N. and Leita, J. (2014). Match analysis in football: A systematic review, *Journal of Sports Science* **32**(20): 1831–1843.
- Scott, D. (2015). *Multivariate Density Estimation: Theory, Practice, and Visualization*, Wiley.
- Shapley, L. S. (1953). Stochastic games, *Proceedings of the National Academy of Sciences of the United States of America* **39**(10): 1095–1100.
- Sheather, S. J. and Jones, M. C. (1991). A reliable data-based bandwidth selection method for kernel density estimation, *Journal of the Royal Statistical Society. Series B (Methodological)* **53**(3): 683–690.
- Silverman, B. (1986). *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, London.

REFERENCES

- Singh, K. (n.d.). Introducing expected threat (xt), <https://karun.in/blog/expected-threat.html>. Accessed: 2022-08-18.
- Sirotic, A., Knowles, H., Catterick, C. and Coutts, A. (2011). Positional match demands of professional rugby league competition, *Journal of Strength and Conditioning Research* **25**: 3076–3087.
- Song, L., Zhou, L., Xu, P., Zhao, W., Li, S. and Li, Z. (2022). Research on the error of global positioning system based on time series analysis, *Sensors* **22**(3614).
- Spearman, W. (2018). Beyond expected goals, *12th MIT Sloan Sports Analytics Conference*.
- Spearman, W., Basye, A., Dick, G., Hotovy, R. and Pop, P. (2017). Physics-based modeling of pass probabilities in soccer, *11th MIT Sloan Sports Analytics Conference*.
- StatsBomb (2022). Statsbomb 360 and evolve.
URL: <https://statsbomb.com/news/statsbomb-360-and-evolve/>
- Stone, C. J. (1984). An Asymptotically Optimal Window Selection Rule for Kernel Density Estimates, *The Annals of Statistics* **12**(4): 1285–1297.
- Sutton, R. and Barto, A. (2018). *Reinforcement Learning: An Introduction*, second edn, The MIT Press.
URL: <http://incompleteideas.net/book/the-book-2nd.html>
- Tierney, G., Weaving, D., Tooby, J., Al-Dawoud, M., Hendricks, S., Phillips, G., Stokes, K. A., Till, K. and Jones, B. (2021). Quantifying head acceleration exposure via instrumented mouthguards (img): a validity and feasibility study protocol to inform img suitability for the tackle project, *BMJ Open Sport & Exercise Medicine* **7**(3).
- Tooby, J., Weaving, D., Al-Dawoud, M. and Tierney, G. (2022). Quantification of head acceleration events in rugby league: an instrumented mouthguard and video analysis pilot study, *Sensors* **22**(2): 584.
- van de Schoot, R., Depaoli, S., Gelman, A., King, R., Kramer, B., Martens, K., Tadesse, M. G., Vannucci, M., Willemsen, J. and Yau, C. (2021). Bayesian statistics and modelling, *Nature Reviews Methods Primers* **1**(3).

REFERENCES

- Vats, K., Fani, M., Walters, P., Clausi, D. A. and Zelek, J. (2020). Event detection in coarsely annotated sports videos via parallel multi-receptive field 1d convolutions.
- Villani, C. (2003). *Topics in Optimal Transportation*, The American Mathematical Society, Providence, Rhode Island.
- Wand, M. P. and Jones, M. C. (1995). *Kernel Smoothing*, Chapman and Hall, London, UK.
- Wedding, C., Woods, C., Sinclair, W., Gomez, M. and Leicht, A. (2020). Examining the evolution and classification of player position using performance indicators in the national rugby league during the 2015-2019 seasons, *Journal of Science and Medicine in Sport* **23**: 891–896.
- Whitaker, G., Silva, R., Edwards, D. and Kosmidis, I. (2021). A bayesian approach for determining player abilities in football, *Journal of the Royal Statistical Society Series C: Applied Statistics* **70**(1): 174–201.
- White, C. C. and White, D. J. (1989). Markov decision processes, *European Journal of Operational Research* **39**(1): 1–16.
URL: <https://www.sciencedirect.com/science/article/pii/0377221789903482>
- Whitehead, S., Till, K., Weaving, D. and Jones, B. (2018). The use of microtechnology to quantify the peak match demands of the football codes: A systematic review, *Sports Medicine* **48**: 2549–2575.
- Woods, C., Sinclair, W. and Robertson, S. (2017). Explaining match outcome and ladder position in the national rugby league using team performance indicators, *Journal of Science and Medicine in Sport* **20**: 1107–1111.
- Woods, C. T. (2016). The use of team performance indicator characteristics to explain ladder position at the conclusion of the australian football league home and away season, *International Journal of Performance Analysis in Sport* **16**(3): 837–847.
- Zambom, A. and Dias, R. (2012). A review of kernel density estimation with applications to econometrics.