Renewals Scheduling for High-Speed Railway Assets

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy

September, 2025

Emily Lauren Buttriss

14321184

Supervisors: John Andrews, Rasa Remenyte-Prescott

University of Nottingham

ABSTRACT

The increasing focus on reducing carbon emissions and easing traffic congestion means that the demand for rail services, particularly high-speed rail, grows ever higher. The cost of ensuring that the railway line remains in a safe and acceptable condition is reflected in the price a passenger pays for travel, hence an efficient asset management strategy should predict when various actions need to be scheduled so that appropriate funds can be provided when needed without overcharging for use of the service. One of the costliest actions is replacement, which requires an in-depth understanding of how the track degrades and what maintenance actions can be completed to improve its condition in order to discern when the useful life has been exhausted and a renewal should be completed. Modelling tools have been developed with aims to improve understanding of the life cycle of a track component, however a review of the surrounding literature revealed that the models rarely distinguished between different failures affecting the track, and the implemented renewal strategies were too simplistic.

This research aims to develop a methodology that can address these limitations by modelling the behaviour of four track assets: crossing panel, switch panel, plain line rail, and the ballast. The degradation process was studied through a reliability analysis completed on historical maintenance and condition databases of a high-speed line situated in the UK. The analysis fitted continuous probability distributions to time-to-failure datasets, to predict future degradation events. The mean time-to-failure and associated parameters of each distribution highlighted the most common failures affecting each track component and how the degradation rate of the fault was expected to change across the operational period. Railway standards were used to develop a Petri net model to represent the life cycle of the four components. Due to the complexities of asset management procedures, additional features such as place dependent transitions or inhibitor arcs were included within the structure. The methodology created sub-models for each of the inspection, degradation, maintenance, and renewal processes, with an individual degradation sub-model created for each failure type, and the ability to choose between a combination of five criteria for the implemented renewal strategy. The results from the previous reliability analysis were incorporated into the Petri net transitions, and Monte Carlo simulation was used to solve the models, to predict when a replacement is required. The models demonstrated their usefulness in

investigating the renewal process, including the importance of each criterion for triggering when a replacement is scheduled or how varying the renewal threshold influences the expected lifetime.

The final stage of the research applied an artificial neural network to understand how the different features of the track can affect degradation and ultimately when a renewal is expected to be scheduled. A variety of training parameters, input variables, and data configurations were investigated but ultimately the models were unable to find a significant relationship between the characteristics of an asset unit and the degradation rate of a fault. A review into the behaviour of the studied railway line concluded that the track was too young and in too good a condition for a clear degradation trend to be established, causing the high prediction errors from the models.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank Mat Rippin, Andy Vickerstaff, and other members of Network Rail High Speed for their invaluable knowledge and input over the course of this PhD. A huge gratitude is also owed to London St Pancras Highspeed for their financial assistance that enabled me to complete this research.

I am incredibly appreciative for my supervisors, Professor John Andrews and Dr Rasa Remenyte-Prescott, whose immense knowledge and continuous support has guided me throughout the researching and writing of this thesis. Without them, this research would not have been possible. I would like to extend my thanks to Dr Silvia Tolo for the useful feedback following my end of year reviews, and the University of Nottingham for the numerous opportunities to advance my studies.

The past seven years in Nottingham have been fun, rewarding, and ultimately life changing. I am so grateful for everyone I have met whilst studying here, from RERG to Shersoc. To my fellow researchers in engineering, thank you for listening to all my concerns and offering practical solutions. To those outside of engineering, thank you for listening to all my concerns and offering less practical solutions.

Finally, my biggest thanks go to all my friends and family back home for their support and motivation. At times, I found it hard to imagine I would ever complete this research, but you never stopped believing I could. I would especially like to thank my mum, my dad, and my sisters, Katie and Louise; your endless love, encouragement, and patience, has meant everything to me.

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION 1
1.1 Background 1
1.2 Renewal of HS1 railway assets 2
1.3 Assets to be considered 3
1.4 Maintenance processes employed to extend life 6
1.5 Development of a new renewal model 7
1.6 Aims and objectives 8
1.7 Thesis outline 9
CHAPTER 2 LITERATURE REVIEW OF MODELLING TO SUPPORT ASSET LIFE PREDICTION 11
2.1 Introduction 11
2.2 Degradation modelling 12
2.2.1 Deterministic models 12
2.2.2 Reliability based models 14
2.2.3 Artificial neural networks 21
2.2.4 Summary of the literature surrounding degradation modelling 26
2.3 Asset management modelling 27
2.3.1 Markov models 28
2.3.2 Petri net models 29
2.3.3 Summary of the literature surrounding asset management modelling 36
2.4 Conclusions 38
CHAPTER 3 SUMMARY OF THE METHODOLOGY AND CONCEPTS UTILISED 41
3.1 Introduction 41
3.2 Censored data 41
3.3 Continuous probability distributions 43
3.3.1 One-parameter exponential 43
3.3.2 Normal 43
3.3.3 Two-parameter lognormal 44
3.3.4 Two parameter Weibull 44
3.4 Parameter estimation and goodness of fit methods 45
3.4.1 Parameter estimation 45
3.4.2 Goodness of fit 46
3.5 Petri nets 47
3.5.1 Components 47
3.5.2 Firing rules 48
3.5.3 Additional features 49
3.6 Monte Carlo simulation 51
3.7 Artificial neural networks 51
3.7.1 Structure 52
3.7.2 Normalisation 53
3.7.3 Forward propagation 53
3.7.4 Calculating loss 54
3.7.5 Back propagation 55
3.7.6 Overtraining 58
3.8 Summary 59

CHAPTER 4 RELIABILITY ANALYSIS OF THE STUDIED TRACK ASSETS 61
4.1 Introduction 61
4.2 Degradation of the studied assets 62
4.3 Data processing 65
4.3.1 Rail defect history 65
4.3.2 Track geometry history 68
4.3.3 Rail profile history 73
4.3.4 No-occurrence faults 74
4.4 Fitting distributions 75
4.4.1 Mean-time-to-failure 75
4.4.2 Parameter estimation method 76
4.4.3 Goodness-of-fit tests 80
4.4.4 Best fitting distribution per failure type 81
4.5 Analysis of the fitted distributions 82
4.5.1 Switch panel 85
4.5.2 Crossing panel 89
4.5.3 Plain line rail 93
4.5.4 Ballast 97
4.6 Summary 100
CHAPTER 5 MODELLING HIGH SPEED RAILWAY TRACK RENEWAL 103
5.1 Introduction 103
5.2 Track model 103
5.2.1 Inspection 104
5.2.2 Degradation 104
5.2.3 Maintenance 107
5.2.4 Renewal 114
5.2.5 Entire asset model 121
5.2.6 Assumptions 122
5.3 Analysis of models 123
5.3.1 Construction 123
5.3.2 Transition firing times 123
5.3.3 Convergence 126
5.3.4 Renewal parameters 126
5.4 Application of models 127
5.4.1 Crossing panel 129
5.4.1.1 Validation 129
5.4.1.2 Expected lifespan of the crossing panel 131
5.4.1.3 Expected condition of the crossing panel 132
5.4.1.4 Including multiple criteria in the renewal strategy 135
5.4.2 Switch panel 137
5.4.2.1 Validation 137
5.4.2.2 Expected lifespan of the switch panel 139
5.4.2.3 Expected condition of the switch panel 140
5.4.2.4 Comparing the expected lifespan of multiple assets 143
5.4.3 Plain line rail 145
5.4.3.1 Validation 145
5.4.3.2 Expected lifespan of the plain line rail 148
5.4.3.3 Expected condition of the plain line rail 150
5.4.3.4 Investigating the importance of a renewal threshold 152
5.4.4 Ballast 154
5.4.4.1 Validation 154

5.4.4.2 Expected lifespan of the ballast 154
5.4.4.3 Expected condition of the ballast 156
5.4.4.4 Adapting the asset management strategy 158
5.5 Summary 160
CHAPTER 6 DEVELOPING ARTIFICIAL NEURAL NETWORKS TO INVESTIGATE ASSET DEGRADATION........ 163
6.1 Introduction 163
6.2 Determining track features 164
6.2.1 Geographical features 165
6.2.2 Operational features 165
6.2.3 Environmental features and opposing characteristics 166
6.3 Developing the ANN 166
6.3.1 Considered variables 166
6.3.2 Sufficient variation 167
6.3.3 Training parameters 168
6.3.4 Artificial dataset 170
6.4 Investigating the ANN's prediction accuracy 172
6.4.1 Identification of tamping interventions 172
6.4.2 Simplifying the model 173
6.4.3 Applying the ANN to the ballast 173
6.4.4 Applying the ANN to the plain line rail 176
6.5 Evaluating the model's performance given the behaviour of the HS1 line 177
6.5.1 Random behaviour within track subgroups 177
6.5.2 Omitted track features 179
6.5.3 Shared behaviour between track subgroups 180
6.6 Summary 180
CHAPTER 7 CONCLUSIONS 183
7.1 Summary 183
7.2 Reflection on the initial objectives 184
7.3 Limitations and future work 188
REFERENCES 191
APPENDIX A A-1
A.1 Switch panel A-1
A.2 Crossing panel A-4
A.3 Plain line rail A-9
A.4 Ballast A-15

ABBREVIATIONS

AC – Associated Cost

AIC – Akaike's Information Criteria

ANN – Artificial Neural Networks

ATI – Artificial Tamping Interventions

BIC – Bayesian Information Criteria

CB – Condition Based

CDF – Cumulative Density Function

CI – Confidence Interval

CP – Control Period

CWR – Continuously Welded Rail

HS1 – High Speed 1

MAC – Minimum Action Code

MAE – Mean Absolute Error

MGT – Million Gross Tonnes

ML – Maintenance Limit

MLE – Maximum Likelihood Estimation

MMA – Manual Metal Arc

MSE – Mean Squared Error

MTTF – Mean Time To Failure

NRHS – Network Rail (High Speed)

OMRC – Operations Maintenance and Renewal Costs

OR – Overall Ranking

ORR – Office of Rail and Road

PDF – Probability Density Function

PN – Petri Nets

RCF – Rolling Contact Fatigue

RRX – Rank Regression on X

RRY – Rank Regression on Y

RTI – Recorded Tamping Interventions

RUL – Remaining Useful Life

S&C – Switches and Crossing

SD – Standard Deviation

SP – Set Period

TD – Time Duration

TRV – Track Recording Vehicle

UTU – Ultrasonic Testing Unit

LIST OF FIGURES

Figure 1.1: Route of the HS1 line, including major stations ... 1
Figure 1.2: Crossing panel components ... 3
Figure 1.3: Switch panel components ... 4
Figure 1.4: Plain line rail components ... 5
Figure 2.1: Boxplots of degradation rates for different track conditions ... 14
Figure 2.2: Cause of failures on an S&C unit measured in percentage of failure minutes ... 16
Figure 2.3: Maintenance responses for geometry defects on the ballast ... 20
Figure 2.4: Fitting linear regression lines to geometry measurements ... 22
Figure 2.5: Slide chair Petri net model ... 32
Figure 2.6: Decision making process for the threshold of maintenance ... 33
Figure 2.7: Renewal strategy for the ballast section Petri net ... 35
Figure 3.1: Collecting censored data ... 42
Figure 3.2: Impact of the shape parameter of the Weibull distribution on the degradation rate ... 45
Figure 3.3: Components of a Petri net ... 47
Figure 3.4: Example of a transition firing ... 48
Figure 3.5: Inhibitor arcs ... 49
Figure 3.6: Place dependent transitions ... 50
Figure 3.7: Structure of an artificial neural network ... 52
Figure 4.1: Typical behaviour of squatting on a crossing panel unit ... 66
Figure 4.2: a) recorded and b) artificial and recorded high-output tamping interventions on a ballast unit ... 69
Figure 4.3: Segmentation between high-output tamping interventions ... 70
Figure 4.4: Detecting isolated faults for the vertical alignment measurements within a unit of ballast ... 73
Figure 4.5: SimuMatic plot of the exponential distribution with parameters estimated by a) RRX and b) MLE ... 77
Figure 4.6: SimuMatic plot of the lognormal distribution with parameters estimated by a) RRX and b) MLE ... 78
Figure 4.7: SimuMatic plot of the Weibull distribution with parameters estimated by a) RRX and b) MLE ... 79
Figure 4.8: Degradation process of cracking on the switch rail ... 88
Figure 4.9: Total number of faults experienced per year by each rail-based asset on the HS1 line ... 95
Figure 5.1: Inspection process sub-model ... 104
Figure 5.2: Degradation of a failure type ... 105
Figure 5.3: Degradation process sub-model ... 105
Figure 5.4: Implementation of opportunistic maintenance between separate failure types ... 107
Figure 5.5: Decision-making strategy for maintenance ... 109
Figure 5.6: Track quality status bands ... 110
Figure 5.7: Maintenance process sub-model ... 111

Figure 5.8: Modelling of cyclic grinding to rectify faults 113
Figure 5.9: Renewal process sub-model 115
Figure 5.10: Impact of cyclic grinding on the material removed from the rail 121
Figure 5.11: Entire asset model... 122
Figure 5.12: Span of distributions fitted to each fault 124
Figure 5.13: Petri net structure for the optimisation technique 125
Figure 5.14: Convergence of the mean predicted lifespan for a crossing panel unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation. 130
Figure 5.15: Distribution of the predicted lifetimes from simulating the crossing panel PN with the SP criterion considered in the renewal strategy 131
Figure 5.16: Average expenditure on the crossing panel per control period with the AC criterion considered in the renewal strategy... 133
Figure 5.17: Average condition of a crossing panel unit over a 40-year observation period with the AC criterion considered in the renewal strategy 134
Figure 5.18: Distribution of criteria triggering a renewal of the crossing panel, for a renewal strategy considering all five criteria 135
Figure 5.19: Distribution of criteria triggering a renewal of the crossing panel for a renewal strategy considering the SP and CB criteria 136
Figure 5.20: Distribution of the predicted lifetimes from simulating the crossing panel PN with the SP and CB criteria considered in the renewal strategy 136
Figure 5.21: Convergence of the mean predicted lifespan for a switch panel unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation. 138
Figure 5.22: Magnified plots of convergence of the mean predicted lifespan for a switch panel unit with a renewal strategy considering the a) SP b) TD c) AC criterion in isolation 140
Figure 5.23: Average condition of a switch panel unit over a 40-year observation period with the AC criterion considered in the renewal strategy 141
Figure 5.24: Expected no. faults for the shallow flange contact angle for a 40-year observation of a switch panel unit 143
Figure 5.25: Convergence of the mean predicted lifespan for a plain line rail unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation. 147
Figure 5.26: Distribution of causes for renewal on the crossing panel for a renewal strategy considering the a) ML, b) SP, c) TD, d) AC, e) CB criterion in isolation 148
Figure 5.27: Distribution of causes for renewal on the crossing panel for a renewal strategy considering the ML criterion in isolation with a threshold of 1 fault 149
Figure 5.28: Average condition of a plain line rail unit over an 80-year observation period with the AC criterion considered in the renewal strategy 151
Figure 5.29: Predicted lifespan of the plain line rail achieved by changing the thresholds for a renewal strategy considering solely the CB criteria 153
Figure 5.30: Convergence of the mean predicted lifespan for a ballast unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation.............. 155
Figure 5.31: Average condition of a ballast unit over a 100-year observation period with the AC criterion considered in the renewal strategy 157
Figure 5.32: Average condition of a ballast unit over a 100-year period with the ML criterion considered in the renewal strategy and a maintenance strategy of only maintaining once the ballast condition is poor or beyond 159

Figure 5.33: Average condition of a ballast unit over a 100-year period with the ML criterion considered in the renewal strategy and a maintenance strategy of only maintaining once the ballast condition has degraded beyond 'good' 160
Figure 6.1: A section of the HS1 schematic 165
Figure 6.2: Structure of the ANN modelling ballast degradation.............................. 169
Figure 6.3: Prediction accuracy of the ANN per epoch of the artificial training dataset 171
Figure 6.4: Residuals for the artificial testing dataset... 171
Figure 6.5: Plot of residuals for the BD2 training and testing datasets implemented in the M4 model 175
Figure 6.6: MSE of the training and testing BD2 datasets implemented in the M4 model 176
Figure 6.7: Fitted regression lines to SD vertical alignment measurements for consecutive units with identical features 178

LIST OF TABLES

Table 2.1: Lognormal parameters and goodness-of-fit results for different ballast subgroups 19
Table 2.2: Maintenance and repair strategies based on the defect density 24
Table 2.3: Definition of the five states derived from the CTR index 26
Table 2.4: Stacked probabilities for maintenance corresponding to different failure types 32
Table 2.5: Implemented renewal strategies within the considered literature 37
Table 4.1: Various inspection procedures and frequency 62
Table 4.2: Different failure types affecting the railway 64
Table 4.3: Example of a dataset containing times-to-failure for a fault 67
Table 4.4: Parameters for the linear and exponential regression lines fitted to the track geometry history 71
Table 4.5: Condition that the ballast is restored to after high-output tamping interventions 72
Table 4.6: Parameters for the linear and exponential regression lines fitted to the rail profile history *assumed parameter 74
Table 4.7: Best fitting distribution per fault 81
Table 4.8: Overall best fitting distribution per failure type 82
Table 4.9: Fitted distributions for faults on the switch panel *estimated distribution .. 86
Table 4.10: Fitted distributions for faults on the crossing panel *estimated distribution 90
Table 4.11: Fitted distributions for faults on the plain line rail *estimated distribution 94
Table 4.12: Fitted distributions for faults on the ballast *estimated distribution 98
Table 5.1: Viable maintenance actions per asset 108
Table 5.2: Historical condition of the SD of the vertical alignment when tamped 110
Table 5.3: Ratios implemented in the decision-making strategy 111
Table 5.4: Description of the places seen in Figure 5.9 116
Table 5.5: Assumed cost (in units) per maintenance action 119
Table 5.6: Impact of corrective maintenance on the condition of rail-based assets .. 120
Table 5.7: Impact of corrective maintenance on the condition of the ballast 120
Table 5.8: Implemented renewal parameters 127
Table 5.9: Average number of faults, grinding actions, and welding actions occuring on a crossing panel unit for the observed data and simulation results 129
Table 5.10: Contribution of each failure type to the faults on the crossing panel for the observed data and simulated results 129
Table 5.11: Average expected lifespan of a crossing panel unit with each renewal criterion considered in isolation for the renewal strategy 131
Table 5.12: Expected maintenance actions completed on a crossing panel unit with the AC criterion considered in the renewal strategy over a 40-year observation period... 132
Table 5.13: Average number of faults, grinding actions, and welding actions occuring on a switch panel unit for the observed data and simulation results 137
Table 5.14: Contribution of each failure type to the faults on the switch panel for the observed data and simulated results 137

Table 5.15: Average expected lifespan of a switch panel unit with each renewal criterion considered in isolation for the renewal strategy 139
Table 5.16: Expected maintenance actions completed on a switch panel unit with the AC criterion considered in the renewal strategy over a 40-year observation period........ 142
Table 5.17: Average expected lifespan of the switch and crossing panel unit for identical renewal criterion considered in isolation for the renewal strategy 144
Table 5.18: Expected maintenance actions completed on a switch and crossing panel units with the AC criterion considered in the renewal strategy over a 30-year observation period 144
Table 5.19: Average number of faults, grinding actions, and welding actions occuring on a plain line rail unit for the observed data and simulation results 145
Table 5.20: Contribution of each failure type to the faults on the plain line rail for the observed data and simulated results 146
Table 5.21: Average expected lifespan of a plain line rail unit with each renewal criterion considered in isolation for the renewal strategy 148
Table 5.22: Expected maintenance actions completed on a plain line rail unit with the AC criterion considered in the renewal strategy over a 80-year observation period........ 150
Table 5.23: Average expected lifespan of a ballast unit with each renewal criterion considered in isolation for the renewal strategy 154
Table 5.24: Expected maintenance actions completed on a ballast unit with the AC criterion considered in the renewal strategy over a 40-year observation period........ 158
Table 6.1: Variables of the SD of the vertical alignment (ballast) ANN 167
Table 6.2: Variables of the vertical wear (plain line) ANN 167
Table 6.3: Variance of the considered features 168
Table 6.4: Prediction accuracy of the ANN with a varying number of hidden neurons 170
Table 6.5: Variables included per model........ 173
Table 6.6: R^2 score given the optimal training parameters: cost function, learning rate, training epochs, and hidden neurons for models predicting the degradation rate for SD of the vertical alignment........ 174
Table 6.7: R^2 score given the optimal training parameters: cost function, learning rate, training epochs, and hidden neurons for models predicting the degradation rate for vertical wear........ 176
Table 6.8: Applying the Pearson's correlation coefficient to the track features and gradients of the fitted regression lines........ 179

Chapter 1

Introduction

1.1 Background

High-speed rail – defined as a railway line capable of operating at speeds above 200km/h (125mph) - is an increasingly popular mode of transport, with the 'bullet train' service first built in Japan in 1964. Following this, high-speed rail has been introduced across Europe, with Italy operating its first high-speed service in 1977, France in 1981, and Germany in 1991. More recently, the UK opened its first and only high-speed line in 2007, known as High Speed 1 (HS1), which runs from London St Pancras to the beginning of the Channel Tunnel (Figure 1.1).

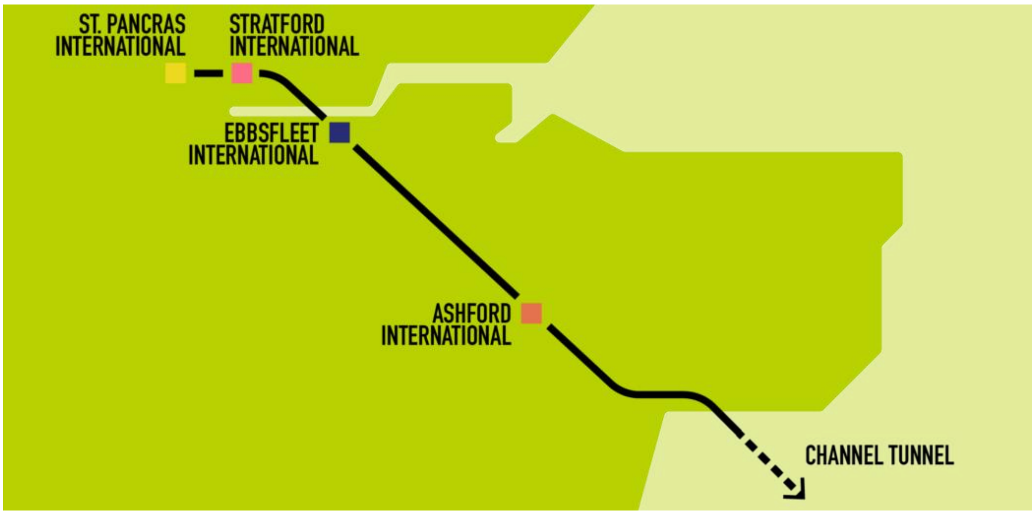

Figure 1.1: Route of the HS1 line, including major stations (London St Pancras Highspeed, 2025)

The line is owned for a 30-year concession by London St. Pancras Highspeed, who have hired the contractors Network Rail (High Speed) (NRHS) for all operation and maintenance activities. NRHS oversee all aspects of the railway line, including structures, signalling, earthworks, drainage, track, and electrification.

1.2 Renewal of HS1 railway assets

Time and use will cause the railway line to degrade. Frequent inspections and maintenance will ensure the track is still functional, but eventually the most effective action to manage degradation is a full replacement of the railway assets. Renewals are expensive - in 2023/24, £10.01 million was spent on route renewals for the HS1 line (London St. Pancras Highspeed, 2024) - and their expected frequency across an operational period will influence the costs passed onto the operators using the railway line. Therefore, it is of great importance to accurately predict when an asset has reached its end of life and a renewal is required. If an insufficient amount is charged to train and freight operators, inadequate funding for renewal will be available when needed, and if this cost is too high the operators (and subsequently the passengers) overpay for the service.

Upon opening of the HS1 line, a renewal schedule spanning 40 years was established based on fixed expected lifetimes of the track assets. Every five years - known as a control period - there is a review into the management strategy, with aims to refine the renewal schedule as more data regarding the asset’s condition becomes available. The expected cost of renewal in any control period must be approved by the economic regulator, the Office of Rail and Road (ORR), before any changes to the scheduling can be enacted, hence, NRHS must be able to justify the refinements by including evidence from supporting studies. For the most recent control period – CP3, extending from April 2020 to March 2025 – NRHS used an OMRC (Operations, Maintenance, and Renewal Costs) model written in Excel to calculate the expected price of regulating the railways, including an estimation of the renewal cost. The model allows for various traffic scenarios to be simulated, to assess the impact of the different train operating companies using the HS1 track. Within the model there are four categories of cost:

- Variable – costs from trains running on the common track
- Avoidable – long-term incremental costs caused by an operator
- Common – costs due to operating the line, such as head-office costs
- Ancillary – costs that are subject to annual change such as insurance

However, for the variable and avoidable costs (which relate to the wear and tear of the track), the planned renewals from the original 40-year schedule were used as parameters.

Therefore, even for the most up-to-date model, renewal estimations were still based off arbitrary fixed lifetimes; they do not consider actual HS1 track conditions or account for shock failure events such as a rail break. Additionally, the renewal predictions are not reflective of any implemented strategy, such as renewing when maintenance has become too frequent or when a given threshold for maintenance actions has been reached.

1.3 Assets to be considered

The focus of this research will be on predicting the lifespan of track-based assets. Within this category, the considered assets are:

I. Crossing panel

The crossing panel – also known as the 'casting' - is located at the intersection between two railway lines and exists to ensure a safe and smooth transition for the passing trains. The structure of a crossing panel is shown in Figure 1.2, consisting of two wing rails and the 'vee'. The vee and wings can be further classified into sections such as the apron, flangeway, and foot, also described in Figure 1.2.

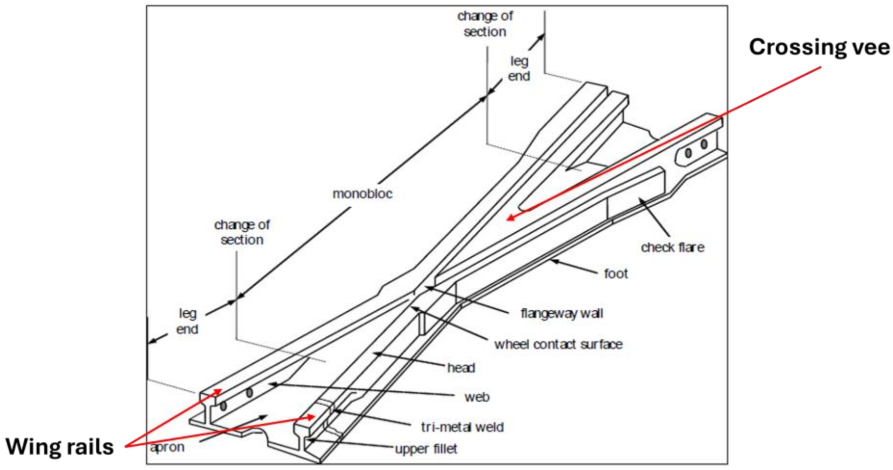

Figure 1.2: Crossing panel components (Network Rail (High Speed) Ltd, 2020a)

The crossing panel can be fixed, where the vee remains stationary, or swing, where the vee moves to reduce the gap between the wing rails. The HS1 line contains a total of 143 crossing panels, with both types used, but the swing crossings are more common on the high-speed sections of track, with fixed crossings typically present around stations and other low-speed areas.

II. Switch panel

The switch panel is the section of the railway track that moves to allow the rolling stock to change direction and cross over to a new line. On the HS1 line, the switch panel is moved via a switch motor which is operated remotely by a signaller, to ensure that the train takes the correct route. On some older railways, this process is still done manually with an engineer operating a lever. The switch panel consists only of the switch and stock rails (identified in Figure 1.3). The very end of a switch rail, and hence the first part of the rail that the rolling stock approaches, is known as the 'switch toe'.

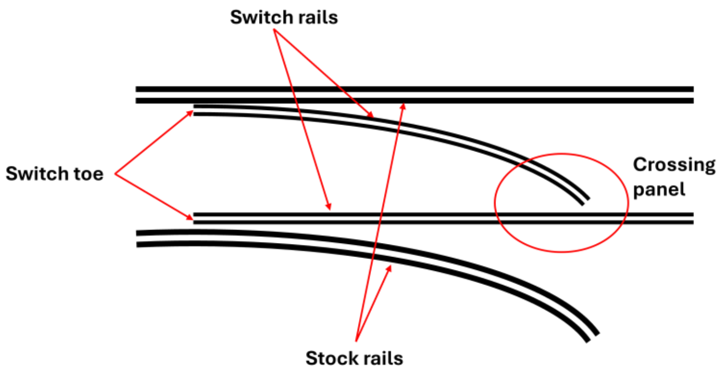

Figure 1.3: Switch panel components

There is a total of 143 switch panels across the HS1 line. A switch and crossing panel are technically components of a larger asset known as an S&C (switches and crossing) however the two components are considered as separate independent assets, with a renewal of both at the same time a coincidence rather than policy.

III. Plain line rail

The plain line rail is any section of track with rails set at a fixed distance apart that doesn't contain other assets like switches or crossings. During installation of the HS1 line, 109m segments of CEN60 type rail are welded together to create Continuously Welded Rail (CWR). This minimises the stresses placed on the rail, particularly surrounding bolt holes in the web, that occur with other rail joining methods such as fishplates (Cannon et al., 2003). Unlike the switch or crossing panel, of which there is a discrete number of units that operate independently to each other, plain line rail is several long uninterrupted segments, each spanning multiple kilometres. It is necessary to separate these segments into discrete units with homogeneous features (*maximum speed limit*, *approximate yearly tonnage* *received*, and *curvature*) before any analysis occurs. A single unit of plain line rail is defined here as a 109m section of rail with an accompanying weld join. Dividing the entire length of the rail into 109m sections and removing all sections that contain other assets or contrasting features allowed for 1739 units of plain line rail to be identified on the HS1 line. The rail itself can be divided into three areas, shown in Figure 1.4 which depicts the cross section of a rail: the head, web, and foot. The side of the rail that faces inwards, toward the other rail, is known as the 'gauge side', whereas the side of the rail that faces outwards is the 'field side'.

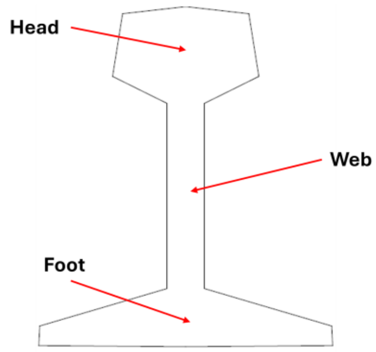

Figure 1.4: Plain line rail components

IV. Ballast

The ballast is the layer of stones (aggregate) that sit underneath the rail. They provide a foundation for the plain line rail and S&C assets, keeping them stable and aligned whilst allowing for drainage off the track. An alternative to ballast is known as slab track, where the rails are attached to a concrete slab. This is more expensive to install and replace compared to ballast, however because they are more structurally stable, slab track has a higher lifespan and is less likely to require maintenance (Michas, 2012). Given its long lifespan, and because slab track makes up less than 20% of the total track length, it is not considered in this research. A single unit of ballast is defined as a 200m long section of the aggregate beneath the left and right rails. These were identified by dividing the entire length of the HS1 line into 200m sections and removing any sections containing slab track or contrasting features, similar to the method described for the plain line rail. In total there was 849 studied units of ballast on the HS1 line.

To allow the rolling stock to travel safely to its destination, there are other assets present on the railway track such as sleepers, nuts and bolts, switch motors, and expansion joints. These often operate alongside the considered assets, however as they are renewed and maintained separately, they are not considered within the scope of this research.

1.4 Maintenance processes employed to extend life

When the track does degrade, there are several maintenance actions that can restore it back to an optimal condition. The type of maintenance utilised depends on the severity and location of the degradation as well as the asset it occurs on. The various methods are described below:

- Tamping

  The ballast is maintained via tamping machines which insert large tines into the ballast and force the aggregate to move. Measuring equipment on-board the machines operate the tines until an optimal geometry is reached. High-output tamping refers to machines that correct large sections of ballast at once, which risks unnecessarily correcting adequate ballast for small faults. This could have a negative impact on the rate of degradation as it is commonly understood that the vibration of the tines damages the ballast (Andrews, 2013), by breaking up and smoothing the aggregate thus reducing the friction between the stones. This causes the ballast to shift even more, requiring further tamping to correct. Localised tamping machines known as Sprinter tamping, that only correct small areas of ballast (<10m), were introduced as part of NRHS's maintenance fleet in 2013 to combat this.

- Grinding

  Grinding is a maintenance action that can be performed on any rail-based asset (plain line rail, switch panel, or crossing panel) and removes the top layer of material from the rail. This reshapes the rail into its optimal profile and removes any surface level defects. Grinding can be completed manually by a track engineer using a hand-held grinder to correct small faults, or via grinding trains to cover large areas at once. NRHS has implemented a grinding campaign on the plain line rail, which involves regularly removing a layer of depth 0.1-0.5mm from the head of the units via a grinding train to remove any existing surface level defects and prevent further defects from developing.

- Welding

  Welding, which is solely applied to localised faults on rail-based assets, involves adding material to a rail that has been heated until it's malleable, to cover any flaws and restore the rail profile. There are several types of welding technique employed to maintain an asset: aluminothermic, which creates a molten iron by heating aluminium powder and an iron oxide to pour between faults (Steenbergen, 2009), Manual Metal Arc (MMA) welding, which uses an electrode as a filler material that fuses to the surface of the rail when the two are connected via an arc (Weman, 2003), and flash butt welding, which applies a current to heat two segments of rail until molten and presses them together to create a join (Steenbergen, 2009).

There are further maintenance actions, such as lubricating the rail or imposing speed limits, that will not directly repair the unit but can stop or slow faults from developing further until one of the described options can be completed. The purpose of maintenance is to extend the lifespan of an asset by restoring its condition and preventing failure, but eventually the only viable option is to replace.

1.5 Development of a new renewal model

Section 1.2. has highlighted a clear need for a model that can accurately predict when a track asset needs replacing. Development of the model requires a thorough understanding of how the asset will degrade, which can be obtained from regular inspections to determine the extent of any wear, defects, or poor geometry. Analysis of historical failure records will allow for patterns and trends of the degradation to be identified, which are used to populate the model. The previously described maintenance actions must also be correctly assigned to the appropriate faults and included within the structure of the model, to reflect the current management strategy.

The model should ultimately be capable of incorporating the inspection, degradation and maintenance processes of a track asset to predict its lifespan. For investigative purposes, it should also allow for an evaluation of which renewal strategy will keep the asset in a good, economically viable condition for the longest time, without risking the safety of passengers. However, defining wear-out is complex. There are operating standards in place, such as minimum rail depth, to ensure that track assets are kept in a safe condition and reduce the risk of derailment. A violation of the standards will trigger a renewal regardless of age, but an effective renewal strategy should aim to replace an asset before the safety limits have been breached, as that implies the track is already too worn, risking the comfort and safety of passengers. The end of an assets useful life can be considered in many ways, such as:

- When a threshold has been reached regarding the total cost of operating the asset, including costs incurred from hiring labour, sourcing materials, rental of machinery, and financial penalties due to a reduced service.
- When maintenance has become ineffective at restoring the asset, indicated by frequent occurrence of failures.
- When the condition of the asset – e.g. the amount of rail covered by a weld repair, or size of the aggregate – has reached the limit that is able to be rectified by maintenance.
- When a threshold defining when effective maintenance can be performed has been exceeded.

Implementing all these criteria into a single model, with the choice to select which definition of wear is enacted as the renewal strategy, allows for it to be used as a decision-making tool for track engineers, helping to provide stronger arguments based in empirical evidence as to why the current renewal schedule should be amended.

1.6 Aims and objectives

The aim of the research presented in this thesis is to predict when a renewal of a track asset is necessary. To achieve this aim, the following objectives have been identified:

- Review the surrounding literature regarding existing models that predict track degradation and asset management
- Decide on an appropriate methodology to achieve the aim
- Understand the types of degradation and resulting maintenance actions defined by the railway standards
- Process and clean the historical data from the HS1 line, and use it to predict future degradation events
- Establish the factors that influence the wear of the considered assets with use, such as the number of train journeys passing over the asset per annum
- Develop a model that can incorporate asset management decisions, and populate with the degradation predictions to predict when a renewal is necessary

- Investigate the different renewal criteria to evaluate their impact on the lifespan and condition of an asset

1.7 Thesis outline

The layout of the thesis is as follows:

- Chapter two will provide an in-depth review of the literature completed so far modelling track asset degradation and management. Different methodologies are discussed, and their strengths and limitations are evaluated.
- Chapter three is a short methodology section explaining the processes of the relevant approaches used in this research.
- Chapter four takes the data from HS1 measurement and failure records and creates cleaned datasets available for processing. The limitations of the data are discussed. A reliability-based study of the datasets is completed, by fitting suitable continuous probability distributions to predict the occurrence of different faults.
- Chapter five describes the development of a Petri net model for a track asset. The distributions found in Chapter four are implemented and used to predict the lifespan of the asset for different renewal strategies. This chapter also demonstrates the other capabilities of the model, such as predicting the condition of the asset over its lifetime.
- Chapter six discusses the importance of different track features by using artificial neural networks to predict the degradation rate of a fault.
- Chapter seven summarises the research, highlights the contributions of this work, and suggests areas for future research.

Chapter 2

Literature review of modelling to support asset life prediction

2.1 Introduction

To best predict when a high-speed railway asset renewal is necessary, it is important to understand how other similar studies were conducted, to illustrate the methodologies other researchers have used, the results they were able to obtain, and the limitations of the study that restricts the developed model's ability for asset lifespan predictions. Evaluating the available literature can then be used to assess which modelling techniques will be the most appropriate for use in this research.

This chapter will examine the methods and analyses applied to the four considered railway components – crossing panel, switch panel, plain line rail, and ballast – in relation to degradation and asset management modelling. Here, degradation is considered as any physical alteration to the condition of the asset, including wear, localised defects like cracking or spalling, and optimal geometry alignments. The discussed literature each contribute a different perspective on the application of a methodology to achieve the aims of this research, but further comprehensive reviews into the modelling of the relevant track assets can be found in work by Hamadache et al. (2019), Rempelos et al. (2024), and Falamarzi et al. (2019). Additionally, despite this work focusing on high-speed railway, literature surrounding track assets on conventional speed lines is also considered as there are no structural differences between the operation of assets in different speed categories.

2.2 Degradation modelling

Understanding how the railway can degrade is crucial to predicting the lifespan of an asset, as the type, severity, and rate of failures experienced will affect when the useful life has been exhausted and therefore when a renewal should be scheduled. A degradation model will aim to predict the future condition of a track component, with the methodology to achieve this largely divided into three groups: deterministic models, reliability-based models, or artificial neural networks.

2.2.1 Deterministic models

Degradation can be understood through a deterministic model, which predict the future condition of a track asset by establishing a direct relationship between its age/traffic, and a failure mode, such as rail wear. If the input of a deterministic model is kept constant, the output will not change. One of the greatest strengths of this modelling technique is that they are often derived based on laboratory experiments, after which they can be applied to almost any situation including new railway lines with very little experience of failures.

Zhang et al. (2000) developed an in-depth deterministic model to investigate the degradation of the ballast, rails, sleepers and fasteners. The model uses a series of equations derived from a mechanistic approach to describe the development of different failure modes (rail wear rate, timber sleeper failure rate, concrete sleeper cracking rate, ballast settlement, subgrade settlement) on each asset. For example, the vertical wear (w_i) of the rail can be described by Equation 2.1:

$$w_i = C_{i,1} k_h k_{l,i} C_{i,2}\, \boldsymbol{W}_i \sin\psi \tag{2.1}$$

Where $C_{i,1}$ and $C_{i,2}$ are constants based on the location of the rail, k_h represents the effect of rail hardness, $k_{l,i}$ is the lubrication factor, $\boldsymbol{W}_i$ accounts for the vertical wheel load, and ψ is the angle of attack of wheel to rail.

The equations are linked via a feedback loop to emulate dependencies for the degradation of the four components, and the model is simulated for a user specified period measured in Million Gross Tonnes (MGT). To demonstrate the model's capability, it was applied to a section of the Australian Queensland Rail track, finding that increasing the axle load from 18.25 tonnes to 22.5 increased the number of sleepers prone to cracking, but not the frequency of ballast maintenance. The results can be used to predict the future behaviour of the four components; however, the current model has no consideration of localised defects on the rail, which could have a subsequent impact on the behaviour of other failure modes or components, affecting any predictions made.

Veit and Marschnig (2010) used an exponential function (Equation 2.2) to describe the deterioration of track quality - defined as the roughness of track, based on a linear combination of speed, gliding influence length, vertical track failure, horizontal track failure, and super-elevation failure - for the ballast:

$$Q(t) = Q_i \times e^{bt} \tag{2.2}$$

Where $Q(t)$ is the track quality over time, Q_i is the initial track quality, and b is the degradation rate

Equation 2.2 can be inverted to find the time at which the track quality has reached a critical level, where maintenance or renewal is required. Although the model itself is simplistic and therefore easily applicable to any railway line, the authors acknowledge that determining a value for the degradation rate is complex as it is dependent on many factors such as traffic rate, substructure quality, and speed. This is expanded upon in Guler et al.'s work (2011), which developed a linear regression model (Equation 2.3) to predict the degradation rate of a track geometry parameter (twist, gauge, alignment, cant, and level) in relation to twelve factors: traffic loads (x_1), velocity (x_2), curvature (x_3), gradient (x_4), cant (x_5), sleeper-type (x_6), rail-type (x_7), rail length (x_8), falling rocks (x_9), land-slide (x_{10}), snow (x_{11}), and flood (x_{12}).

$$\bar{\tau}_i = \alpha_{i1}x_1 + \alpha_{i2}x_2 + \alpha_{i3}x_3 + \alpha_{i4}x_4 + \alpha_{i5}x_5 + \alpha_{i6}x_6 + \alpha_{i7}x_7 + \alpha_{i8}x_8 + \alpha_{i9}x_9 + \alpha_{i10}x_{10} + \alpha_{i11}x_{11} + \alpha_{i12}x_{12} + \alpha_{i0} \tag{2.3}$$

7 years of track geometry measurements was taken from a 180km section of a railway line in Turkey, which was split into 820 homogeneous segments. The mean degradation rate for a given geometry parameter was found for each segment, and then multivariate statistical analysis – which involved finding baseline statistics, hypothesis testing, multicollinearity tests, and stepwise regression analysis – was completed to calculate the optimum values of α_{in}. When applied to the twist, alignment, and level parameters of the track geometry, the resulting models all had an R^2 value of between 0.6-0.7. This suggests that Equation 2.3 will provide a reasonable approximation of the actual degradation rate but will fail to explain the behaviour of over 30% of the segments.

Gular et al.'s research is a prime example of how deterministic models can still incorporate historical failure data, which increases the reliability in any predictions made. However, it can be demonstrated in research such as Zwanenburg's paper (2006), which found that the replacement age of switches and crossings (S&C) units had a high standard deviation, or through the interquartile range of box plots created for Goodarzi et al.'s (2023) research into the degradation rate of the standard deviation of the rail profile (*in*) for 80ft ballast segments (Figure 2.1), that in practice, the behaviour of seemingly identical railway assets are varied.

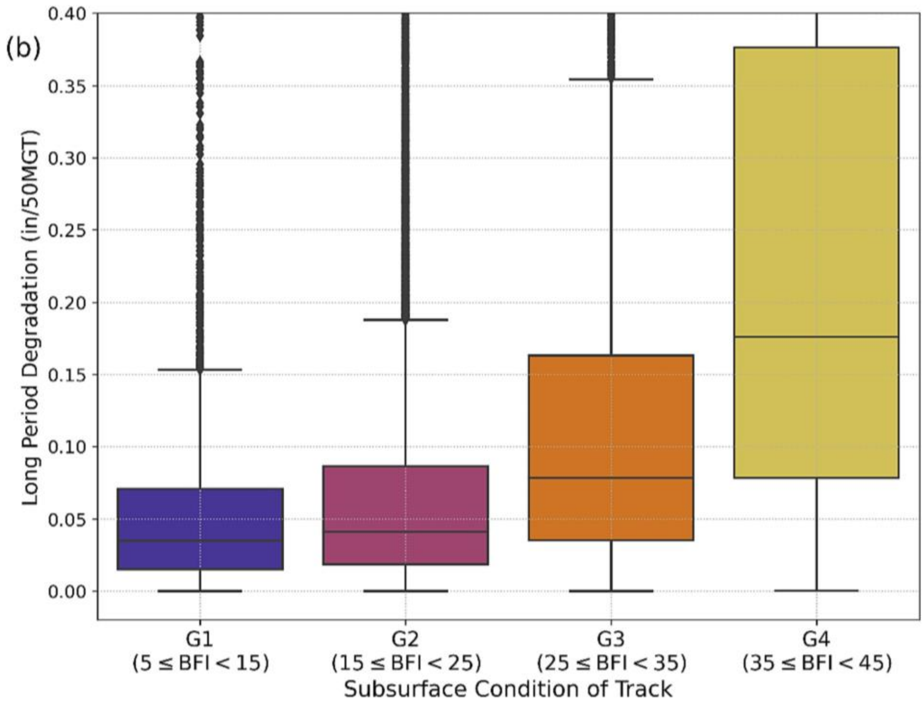

Figure 2.1: Boxplots of degradation rates for different track conditions (Goodarzi, 2023)

Although there are further examples of deterministic models applied to track assets (Sato, 1995; Shafahi et al., 2008), their very nature means they cannot incorporate the inherent randomness and uncertainty of degradation, risking oversimplification. To provide a better estimation of when the track is expected to degrade, a stochastic model – which accounts for variability – can be considered.

2.2.2 Reliability based models

A reliability-based model is derived from the analysis of operational data, finding statistics and fitting probability distributions to describe the occurrence of degradation. They are typically quick to complete and easy to interpret, making them an extremely popular modelling method in engineering. Their stochastic qualities mean they are often an improvement on deterministic models as they do not ignore the impact of isolated degradation events, and they can also be incorporated alongside a maintenance model, that describes the asset management of the track.

In 2013, Rama and Andrews (2013) used failure records spanning from 2002-2012, provided by the Fault Management System database from the UK's conventional speed rail network, to create component life models for S&C units. The considered components included the clamp lock, point heater, and fastenings, but nothing rail-based such as the crossing. This exclusion is potentially limiting to researchers wishing to evaluate what drives S&C degradation, as a study in 2011 on UK S&C components found that the switch rail alone was accountable for 45.3% of all failures (Jalili Hassankiadeh, 2011). Initially, the mean-time-to-failure (MTTF) – Equation 2.4 – was calculated for each component, by summing the total amount of complete failure lifetimes (TTF) and the total censored lifetimes (TC) and averaging them over the total number of recorded faults (NF).

$$MTTF = \frac{\sum_{i=1}^{NFT} TTF_i + \sum_{j=1}^{NCT} TC_j}{NF} \tag{2.4}$$

Equation 2.4 can be inverted to predict the failure rate (γ), the expected number of failures per day (Equation 2.5).

$$\gamma = \frac{1}{MTTF} \tag{2.5}$$

Following this, maximum likelihood estimation and linearised probability plots were used to find the best-fitting probability distribution to represent the failure data, concluding that the two-parameter Weibull distribution was the most appropriate. The corresponding confidence intervals for the shape parameter of the Weibull distributions confirmed that the S&C components in the study all have a reducing failure risk, with most faults occurring within 50 days of installation. Once this 'burn-in period' was passed, there was minimal difference between the failure rate found from the MTTF and from the Weibull distributions. The research did not discern between the type of fault affecting the component, meaning the fitted distributions cannot predict whether the next failure event will be minor or severe. This is vital for understanding when the S&C component has exhausted its useful lifetime, as a series of severe failures suggests that the asset is too worn to be functional.

In 2023, Litherland and Andrews (2023a) built on Rama and Andrew's study by using the maintenance and failure history provided by Network Rail to explore the reliability of more components of a S&C, this time including the check and wing rail, crossing, switch rail, and the stock rail. Both continuous and discrete probability distributions were used to model the likelihood of failure events, with discrete distributions applied when there were multiple identical components within the same S&C unit, making it impossible to know the exact component that the failures occurred on. This was achieved through the Poisson distribution. For components that could be modelled through continuous distributions, maximum likelihood estimation was used to find the most suitable parameters for the exponential, Weibull, normal and lognormal distributions, with the Kolmogorov-Smirnov, correlation coefficient and the likelihood value used as the goodness-of-fit assessments. The results of this study found that the occurrence of failures on both the crossing and switch rail component was again best represented through a two-parameter Weibull distribution, but the crossing was expected to have an increasing degradation rate as it ages, whereas the switch rail would experience less faults over time.

The paper explored the impact of S&C size on the expected number of failures per year, finding that for some components, the size had a positive correlation with failure frequency. The authors concluded that this is likely due to larger S&C's typically having a higher line speed and daily traffic, but these factors were not investigated further. The type of failure affecting the component was again undefined in the data available. Knowing what is driving the failures can help make more informed predictions about the type of maintenance that's needed. Figure 2.2, taken from a thesis dedicated to life-time monitoring of UK-based S&C (Cornish, 2014), shows that cracking of the rail surface (rectified by replacement or welding) was over 5 times more likely to occur on a S&C unit than plastic deformation (rectified by grinding).

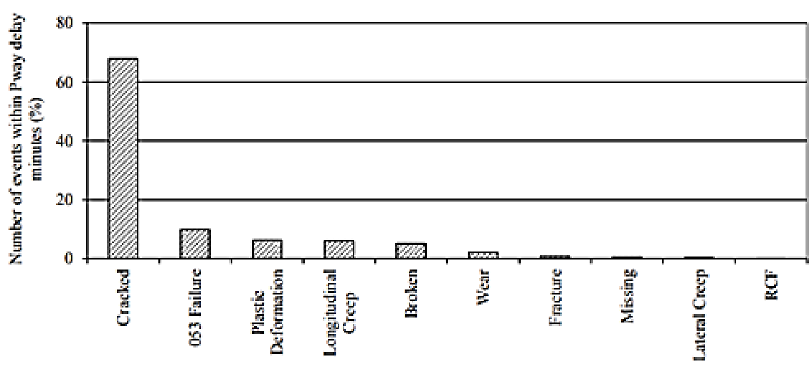

Figure 2.2: Cause of failures on an S&C unit measured in percentage of failure minutes (Cornish, 2014)

Maintenance has a significant impact on evaluating when an asset is no longer operational – for example, it is commonly understood that weld repairs can reduce the integrity of rails (Jun et al., 2016) – and a reliability-based degradation model should incorporate information from studies considering failure, similar to what was shown in Figure 2.2, when predicting the future condition of the asset and therefore what maintenance will be required.

Chattopadhyay and Kumar (2009) used data from the Swedish National Rail Administration to analyse the occurrence of rail breaks on sections of plain line rail. A software programme utilising maximum likelihood estimation found that the two-parameter Weibull and normal distribution were the best choices at explaining this behaviour, however the authors chose to use solely the Weibull distribution, given its adaptability. The obtained parameters of the Weibull distribution were then used to find the MTTF (Equation 2.6).

$$MTTF = \eta\Gamma\left(\frac{1}{\beta} + 1\right) \tag{2.6}$$

Where η and β are the scale and shape parameters of the Weibull distribution respectively, and $\Gamma()$ is the Gamma function

The research explores how the different features of the rail impacts the likelihood of a rail break occurring, finding that keeping all other parameters - curve radius, rail size, rail type and rail steel type - the same, low rail segments (the section of rail located at the inside of a curve) have an expected lifetime of 284 MGT whereas high rail segments (located at the outside of a curve) can experience an average of 382 MGT before a rail break. This suggests that low rail will experience more frequent failures, which the authors attribute to the vehicle speeds travelling over the section although this is not inspected further. In total, five features of the rail were considered but due to the low availability of data, only segments with a curve radius of 500-600 metres could be analysed, and the rail sizes were limited to 50 or 60kg, due to a lack of failures in other subgroups.

Audley and Andrews (2013a) used a reliability-based model to study the degradation behaviour of a 220-yard section of ballast. 6 years of data was provided by Network Rail, regarding the standard deviation of the vertical alignment geometry measurement over the ballast section, recorded with a 35m wavelength (σ). This was subsequently divided into intervals between recorded maintenance interventions. A linear regression line, described in Equation 2.7 where A and B are coefficients and t is a time parameter, was fitted to each interval to describe the trend of σ over time.

$$\sigma = A + Bt \tag{2.7}$$

Equation 2.7 was inverted to find the time at which a critical σ level was reached, generating a sample with around 11000 entries. Then, a variety of distributions were considered to model the behaviour of the sample, and evaluated using the coefficient of determination, finding the two-parameter Weibull distribution to be the most appropriate. The study investigated the relationship between the time taken to reach a critical σ and the maintenance history, finding that for sections with slower speeds, the shape parameter of the Weibull distribution (β) was positively correlated with the number of tamps completed, suggesting that tamping damages the ballast and increases the rate at which it degrades. However, for the top two speed bands (115-125mph, 80-110mph), there was no association between the degradation rate and maintenance history, with β typically increasing after the first tamp but decreasing after the second. This finding highlighted two issues with the study: firstly, due to this research using data from conventional-speed lines, it is unable to further explore whether high speeds can negate the impact of the maintenance history on the degradation rates. Secondly, a 6-year observation period is a relatively short time to study ballast degradation, resulting in a maximum of two tamping actions being considered in the research. Had more data been provided from sections that had experienced many tamps, the relationship between maintenance history and degradation (if any) would be clearer to establish.

In 2011, Andrade and Teixeira (2011) also studied the degradation of the ballast, investigating the impact of infrastructure features by creating four subgroups: bridges, plain track, stations, and switches. The data required for this reliability analysis was taken from 337km of the Portuguese rail line, and after further analysis, 675 homogeneous 200m ballast sections were identified, each having had 15 recorded geometry inspections. For each section, Equation 2.7 was used to model the development of the standard deviation for the short wavelength longitudinal levelling geometry parameter. The samples for coefficients A and B (labelled as c_1 and c_0 in this study) - which relate to the initial standard deviation and the degradation rate respectively - were evaluated with a Kolmogorov-Smirnov goodness-of-fit test to find the best fitting probability distribution to represent their behaviour, finding that overall, the lognormal distribution was the most appropriate. The result of this analysis is shown in Table 2.1:

	Deterioration rate, c_0 (mm/100 MGT)					Initial standard deviation, c_1 (mm)				
		Groups					Groups			
	All	Bridges	Plain track	Stations	Switches	All	Bridges	Plain track	Stations	Switches
Scale parameter	1.041	1.641	0.888	1.080	1.866	0.337	0.349	0.316	0.313	0.578
Shape parameter	0.921	1.067	0.851	1.002	0.614	0.368	0.374	0.278	0.362	0.499
N	675	72	445	99	59	675	72	445	99	59
K-S test p-value	0.0841	0.2838	0.0320	0.9357	0.9798	0.0026	0.7926	0.0166	0.5979	0.1605

Table 2.1: Lognormal parameters and goodness-of-fit results for different ballast subgroups (Andrade & Teixeira, 2011)

It was found that the bridge and switch subgroup degraded faster than ballast sections underneath stations or plain line track. However, as can be seen from Table 2.1, grouping the data will dramatically reduce the available sample sizes, with the switch subgroup containing only 59 sections (8.7% of all those considered). Whilst still significantly large enough to justify the Kolmogorov-Smirnov test, this reduced sample size may introduce bias, making any maintenance or renewal predictions resulting from this research unreliable.

The formerly described analysis of fitting regression lines to geometry measurements (Andrade & Teixeira, 2011; Audley & Andrews, 2013a) was applied to the ballast of the Main Western line in Sweden by Soleimanmeigouni et al. (2020), with Equation 2.8 used to describe the development of isolated short-wavelength vertical defects.

$$A(t) = A_o + \beta(t - t_{tamp}) + \epsilon \tag{2.8}$$

Where $A(t)$ is the absolute amplitude of the defect at time t, A_o is the absolute amplitude of the defect after the previous tamping intervention, β is the degradation rate, $t_{tamp}$is the time of the latest tamping intervention, and ϵ is the Gaussian random error term

In the previously discussed literature, only the standard deviation of the vertical alignment has been considered but as Khajehei et al. (2019) states in their research, this is typically an indicator of when to schedule preventative maintenance whereas isolated geometry defects will trigger corrective maintenance. This is demonstrated in Figure 2.3 taken from the study, where 'IL' and 'IAL' stand for 'normal corrective maintenance limit for track geometry' and 'emergency corrective maintenance limit with speed reduction or line closure' respectively. These thresholds, which cause corrective maintenance (CM) to be employed, are based solely on the severity of isolated defects, whereas preventative maintenance (PM) is dependent on the standard deviation of the geometry measurements reaching the preventive maintenance limit for track geometry 'AL'. A similar approach is applied for the HS1 line, with the consideration of Sprinter tamping, speed restrictions or line closures a sole response to the severity of isolated defects. By modelling degradation through isolated defects, the results of Soleimanmeigouni et al.'s research can assist when predicting the ballast's future condition and provide more insight into when maintenance should be scheduled or when the useful life has been exhausted. Any asset management model for the ballast excluding either isolated defects or the standard deviation of geometry parameters is therefore limited. The study, completed in 2020, used the dataset of the fitted regression lines to explore the spread of the degradation rates under different track features, but did not expand to fitting a continuous probability distribution although this could easily be calculated using a software on the created dataset.

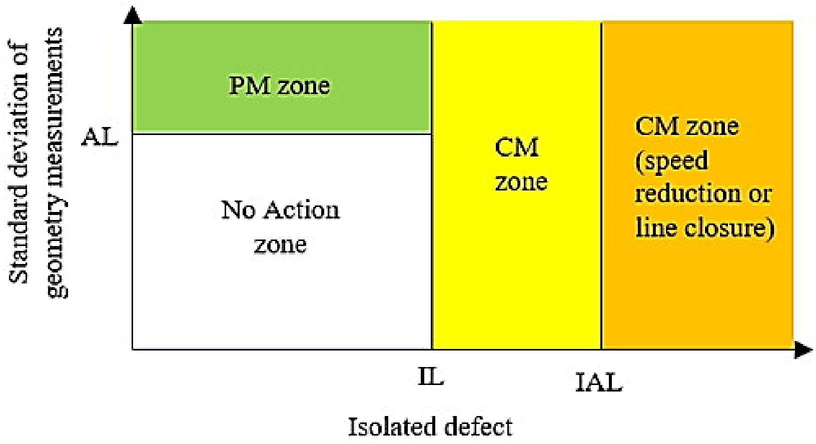

Figure 2.3: Maintenance responses for geometry defects on the ballast (Khajehei et al., 2019)

It is evident that reliability-based models can be used effectively to examine the degradation of a track-based asset and predict when a renewal is necessary. The resulting probability distributions obtained from the study are stochastic, easy to implement, and can consider changing degradation rates over an assets lifetime. However, a reliability-based model's biggest drawback is that it is completely reliant on having an adequate sample size of historical data to draw conclusions from - this may not be available when investigating a young railway such as the HS1 line, risking biased results and inaccurate predictions. Most of the discussed literature within this section has used the methodology to explore the influence of asset characteristics by dividing the data into subgroups, but not only does this risk further minimising an already small sample size, it also requires a completely new analysis for each subgroup which is time consuming.

The inflexibility of reliability-based models means that a different methodology that considers multiple parameters at once and allows the user to easily adapt them could be beneficial, especially if operating conditions were expected to change. For example, during the COVID-19 pandemic the UK's rail network ran a reduced timetable, causing train operations to drop by around 40% (Office of Rail and Road, 2021). Degradation predictions made from reliability-based models with assumed constant traffic would instantly be invalid, but a model trained to understand the association between traffic and degradation and allow for the traffic parameter to be changed would quickly be able to update any predictions. Traditionally, models based around machine learning have been used in this capacity.

2.2.3 Artificial neural networks

Artificial neural networks (ANN) are a type of machine learning model that predict the degradation of a track asset by learning the relationships between the inputs and the output. Unlike reliability-based models, once sufficiently trained, a singular ANN can be applied to any subgroup of an asset by changing the input values, so long as it is within the scope of the training data.

Yilboga et al. (2010) studied the degradation of an electro-mechanical turnout system (defined as a segment of the S&C consisting of the switches, motor, bearings, reduction gear, and drive detection rods) using readings from attached force sensors. Before application of the model, these readings were converted into discrete values representing degradation, ranging from 0-22 where 0 is a fault-free state, but the paper does not state how these values are classified based on the readings, nor what failure modes such as wear or cracking are detected through the force sensor. A time-delay ANN was applied for this research; unlike conventional ANN, time-delay ANN include time units in the input layer and use previous outputs to update the inputs in order to make time-based future predictions. 75% of the available dataset and the Levenberg-Marquardt learning algorithm was used to train the model, with the ANN using the initial or previously predicted degradation value as an input to predict the future degradation value, stopping when a predefined failure threshold was reached. The number of predictions made by the model until the threshold was met was defined as the Remaining Useful Life (RUL) of the S&C.

The predicted and actual RUL were compared, finding an average R^2 value of 0.92, demonstrating the effectiveness of the models for accurately predicting the future condition of an S&C unit. However, without clarification of how degradation is defined or disclosure of the input parameters used in the structure, the methodology cannot be replicated for this research and any found results are not applicable.

Guler (2014) developed an ANN to predict the degradation of a section of railway ballast. Recordings from track geometry inspections were collected over a 3-year period observing a 180km segment of the Turkish railways, which was split into 820 homogeneous sections. Five geometry measurements – twist, gauge, alignment, cant, and levelling – were collected for each section, and a linear regression line (Equation 2.7) was fitted to the inspection values between maintenance and renewal works, a process described in Figure 2.4, taken from the study.

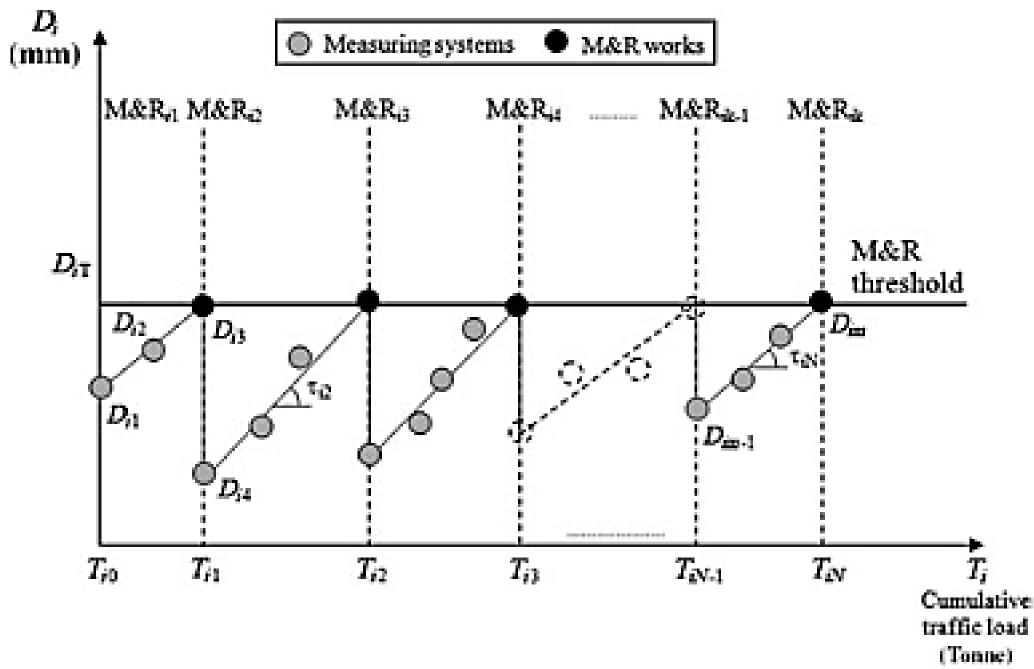

Figure 2.4: Fitting linear regression lines to geometry measurements (Guler, 2014)

The average degradation rate of a geometry measurement for a section (τ_i) was found by summing the gradients of each linear regression line, and dividing by the total amount of maintenance and renewal works that took place on that section. The ANN used 12 features of the ballast section (traffic load (MGT), speed (km/h), curvature (1/R), gradient (%), cross level (mm), sleeper type (concrete or wooden), rail type (49.43 or 49.05 kg/m), rail length (CWR or JR), falling rock (yes or no), landslide (yes or no), snow (yes or no), and flood (yes or no)) as inputs, and 2 layers of hidden neurons to predict the value of τ_i.

After training, the R^2 value of the five models ranged between 0.72-0.84, showing that the methodology can make reasonably accurate predictions.

The input values for the ANN in this research are identical to the parameters used in Guler's previous research using a deterministic model (Guler et al., 2011), and both models are applied using the same dataset. It can be seen by comparing the R^2 values that the ANN is the better degradation estimation method. This is because the ANN can find the non-linear relationships between the inputs and the output and are more robust against noisy data (Bakar & Tahir, 2009). The impact of this study means that if operating conditions were to change within a section, e.g. the wooden sleepers are renewed and replaced with concrete sleepers, the ANN can easily be used to update that sections' degradation predictions. However, by using the average degradation rate τ_i, this research neglects to consider the influence of maintenance actions (tamping or renewals), despite other studies finding that tamping can increase the degradation rate (Audley & Andrews, 2013a; Soleimanmeigouni et al., 2018). The exclusion of an important parameter could make any predictions found unreliable when applied to track asset degradation.

This was rectified in later research by Khajehei et al. (2022), who employed a similar technique on track geometry data supplied by the Swedish Transport Administration. For this study, the gradient of each linear regression line was considered separately, and the number of maintenance actions was used as an input (if a renewal occurred this value was reset to zero), alongside the ballast age, rail type, sleeper type, sleeper age, existence of a drum, existence of a level crossing, existence of a bridge, speed, traffic, degradation level following maintenance, and the curvature. It was found that a training:test ratio of 80:20 and a single hidden layer consisting of 10 neurons resulted in the best performance of the ANN, with an R^2 value of 0.85 reached. Garson's algorithm was used to identify which of the inputs had the biggest impact on the degradation rates, finding that the maintenance history, degradation level following maintenance and traffic were the three most influential factors. This demonstrates how tamping is destructive to the ballast's degradation rate and should have been considered as a parameter in Guler's work, however, unlike in (Guler, 2014), Khajehei et al. only considered the standard deviation of the longitudinal level of the ballast as a measurement of degradation.

Sadeghi and Askarinejad (2012) created an empirical formula (Equation 2.9) to calculate the defect density for the rail, sleeper, fastening and ballast assets:

$$D = 0.33LD + 0.66MD + HD \tag{2.9}$$

Where D is the overall defect density of the asset, and LD, MD, HD are the defect density of low, medium and high fault severities respectively.

LD, MD, and HD are calculated by dividing the total number of affected units (distance between two sleepers) over the total number of units in a 30m segment of the asset. The definition of a 'defect' and its corresponding severity for each asset is described in more detail within the paper but extensively covered localised faults such as scaling from 6-25mm deep, bolt hole cracks >1.5mm, or vegetation in the ballast that interferes with the movement of the train. The result from the empirical formula is then used to indicate the type of maintenance that is needed on the asset for different classes of track (Table 2.2).

Required maintenance and repair strategy	Structural defect density		
	Class B	Class C	Class D
No immediate action	< 0.013	< 0.033	< 0.066
Routine or preventive maintenance	0.013 – 0.033	0.033 – 0.079	0.067 – 0.13
Minor repair	0.034 – 0.066	0.080 – 0.13	0.14 – 0.30
Major repair	> 0.067	> 0.13	> 0.30

Table 2.2: Maintenance and repair strategies based on the defect density (Sadeghi & Askarinejad, 2012)

An ANN model utilising the backpropagation algorithm was developed to predict the value of D for an asset segment, using geometry measurements (standard deviations of the gauge, profile, alignment and twist) from the segment as input values. It was found that different assets required a different combination of inputs; the rail ANN performed best with all four considered, whereas the fastening ANN was most effective when the profile measurement was excluded. The R^2 value from the models' predictions compared to the observed data ranged between 0.74-0.93, with the authors arguing that the accuracy of the models' predictions can be used to minimise the amount of visual inspections needed of the track. The implementation of this methodology has several limitations for predicting the future condition of an asset, namely grouping useful information surrounding defects into a single artificial parameter, D. As can be inferred from Table 2.2, this could potentially lead to an asset segment classified as ‘No immediate action’ despite a high severity fault being present. It also reduces the ability of the model to distinguish what failures and at what severity are likely to occur, and additionally, the scope of the model is very short: the geometry input parameters only predict the current D value of the asset. Therefore, if wishing to estimate the defect density at some time in the future, prior knowledge of the condition of the track geometry at that point is needed.

More recently, Zhang et al. (2025) developed an ANN to model degradation of the plain line rail using seven track features – curve radius, circular curve length, transition curve length, superelevation, front rail gradient, back rail gradient, and the gradient algebraic difference - as inputs. The model was used to predict the wear (W_t) of the rail profile in mm after the rail had received a gross cumulative load of 56,064,000 tonnes, equivalent to 10 years of service. The trained ANN’s had achieved a prediction accuracy above 95% using parameters such as the sigmoid activation function and mean squared error function, and the results were used to indicate the most influential geometry features influencing wear on the rails, finding that this varied depending on whether the section of rail was on the inside or outside of a curve. The definition of wear is as described in Equation 2.10, based on a combination of lateral (W_1) and vertical (W_2) wear. However, in practice, vertical and lateral wear are usually treated as different faults as they degrade at separate rates and have separate operational safety limits. By grouping them into one, this research is unable to predict the degradation of the rail due to either consideration of wear, potentially risking the rail being left in an unsafe condition.

$$W_t = \frac{1}{2}W_1 + W_2 \tag{2.10}$$

Shafahi et al. (2008) used data from 215 sections of an Iranian railway line to create an ANN that predicts the condition of the ballast. The output of the ANN is the section’s track state value, which is a discretised form of the Combined Track Record (CTR) index, a value between 0-100 that is calculated based on track geometry recordings of four parameters – alignment, gauge, crosslevel and profile (Table 2.3). The definition of how the CTR was calculated using the geometry recordings was not included within the scope of the research.

CTR index	0–50	50–60	60–70	70–80	80–100
Track quality	Failed	Medium	Good	Very good	Excellent
Track state	1	2	3	4	5

Table 2.3: Definition of the five states derived from the CTR index (Shafahi et al., 2008)

Previous track state values, as well as the sections traffic, speed, geographical features, gradient and curvature were used as inputs, with 82% of the available dataset used for training. The research found that the ANN was able to predict ballast deterioration better than a previously implemented deterministic model and could correctly predict the track state 67% of the time. Combining failure modes into an index such as the CTR used here, or by a defect density (Sadeghi & Askarinejad, 2012), can help to instantly provide an overview for the condition of an asset and may reduce the complexity of models, but, as Andrews (2013) argues, the use of track indices is inherently flawed as a good recording could still be achieved if one measurement is in an unacceptable condition but all other measurements are good. They are also difficult to interpret – for example in Table 2.3 it is not clear whether a CTR recording of 25 is twice as bad as a recording of 50, despite both being classified as 'failed'.

The discussed examples of ANN in the literature have shown that they can be applied in a multitude of ways to predict the condition of a track asset. They are particularly useful for understanding the importance of different input parameters, allowing for predictions to be updated if operating procedures were to change. However, this is only applicable if the input parameter is changed to a value within the original scope of the training data – an ANN cannot extrapolate to unseen factors. Additionally, their prediction accuracy relies heavily on the choice of learning algorithm, structure, activation function and error function, and finding the optimal combination can be an exhaustive process.

2.2.4 Summary of the literature surrounding degradation modelling

This section has discussed literature centred around modelling degradation on the ballast, plain line rail, crossing panel, and switch panel, for three different failure modes: geometry, wear, and localised defects. The methodology used to predict degradation varied between deterministic, reliability-based, or machine learning models. Applying a similar approach with this research will provide an additional resource to aid comprehension of railway degradation, with an application to high-speed railway under UK operating conditions.

For a better understanding of when the asset has exhausted its useful lifetime, the analysis of a failure mode should be specific to the actual faults that are occurring. For example, ‘localised defects’ is a broad category as railway standards provided by Network Rail High Speed identify 24 separate failures impacting the crossing panel alone, including cracking, lipping, and pitting. In some areas of the literature this has been well researched - the geometry of the ballast has had studies conducted into both the standard deviation (Andrade & Teixeira, 2011; Audley & Andrews, 2013a; Khajehei et al., 2022) and isolated measurements of the vertical alignment (Soleimanmeigouni et al., 2020), as well as other geometry parameters such as the twist or cant (Guler et al., 2011; Guler, 2014). However, for the wear and localised defects of the plain line rail and S&C components, it was found that research has grouped failures, such as combining vertical and longitudinal wear in Zhang et al.’s study (2025) or by not distinguishing what faults occurred on the S&C component like in Litherland and Andrew’s reliability analysis (2023a). If a specific fault was named, it was considered in isolation, like rail breaks in Chattopadhyay and Kumar’s research (2009).

It has previously been demonstrated that grouping, whether by an index or considering multiple failures within one, can lead to faults being ignored, or improper maintenance scheduling as severe and minor faults are treated equally. In practice, the severe fault would be dealt with far more urgently, and with higher consequences (such as speed restrictions or line closures). Therefore, to accurately model the behaviour of a track asset and predict when a renewal is necessary, a degradation model needs to consider every possible fault separately.

2.3 Asset management modelling

The previous models can inform of the degradation process of the railway but this needs to be combined with a method of simulating the track operator’s attempts to control degradation, such as by performing inspections or completing maintenance, to understand the condition of the future asset. The resulting model can be used to determine when the end of life has been reached and be used to investigate the effects of different management strategies. This will be implemented through state-space models.

A state-space model refers to a methodology that represents the dynamic behaviour of a system through states and transitions. It is a graphical way of understanding the asset management strategy and typically allows for investigation by changing parameters such as the firing time of transitions or initial starting conditions. By observing the movement of a state-based model, statistics regarding the states visited or length of time spent in a state can be reported.

2.3.1 Markov models

One popular method of state-space modelling is Markov models (Ibe, 2013). These represent the asset management of a track component using a set of states and a transition probability matrix containing the probabilities of moving between the states. They are relatively simplistic in design and flexible enough for different management strategies to be considered and compared.

Hokstad et al. (2005) used a Markov model to model the degradation and repair of a 1-kilometre section of plain line rail, studying the occurrence of a broken rail. Two types of critical failure were considered- shock failure, which happens without warning, and failures occurring as part of the degradation process, which can be avoided through preventative maintenance. The paper does not mention a specific renewal policy, only that a critical failure (a broken rail) requires immediate repair, which is assumed to be a rail replacement. The developed Markov chain was used to optimise the inspection process by varying the frequency of inspections, finding that a shorter inspection interval lead to a longer mean time between critical failures. Several assumptions were made for the model to be functional, namely that the section of rail can never have multiple failures within an inspection interval and that the mean time to repair a fault was zero days, but these assumptions are potentially limiting as it is unrealistic to actual degradation and maintenance procedures.

Prescott and Andrews (2015) built a Markov model with 80 states to model the asset management of a 220-yard section of ballast. The condition of the ballast was characterised into four categories: good, critical, speed restriction required, and line closure required. An assumption is made that tamping - the sole maintenance action - is destructive and impacts the degradation rate, up to a limit of seven tamps, after which the performance of the ballast remains the same. The model was populated via statistical analysis of data from the UK railway network which found the mean time to deteriorate to each condition category, to be used as the transition rates. The implemented renewal strategy was a fixed lifetime of 30 years, but the model was used to investigate adapting this parameter, finding that the proportion of time that the ballast section spends in a certain condition did not significantly change even when increasing the lifetime to 40 years.

However, the research only considers tamping as a maintenance action, and the authors acknowledge that if a different option such as stoneblowing was also included, this may make the model prohibitively large and make any analysis very difficult. The asset management of a ballast section was also modelled by Lyngby et al. (2008) who created a 50-state Markov model to study the occurrence of twist faults and was subsequently used to evaluate the impact of the inspection frequency, and by Shafahi & Hakhamaneshi (2009), where a five-state Markov model was created using the track state definitions defined in Table 2.3, with the model utilised for finding the optimal maintenance action for each state over a 10-year planning horizon. The considered literature has demonstrated how Markov models can be used effectively for modelling and investigating asset management decisions (such as renewal), but they also have many limitations, namely their memoryless property. This rules that the rates of transitions between states are kept constant and do not rely on previous behaviour in the model. This is typically enacted by fitting exponential distributions to describe the development of a fault, but the results from the degradation models discussed in section 2.2.2 indicate that this may not be a suitable fit, impacting the reliability of any lifespan predictions. Additionally, asset management relies on a number of different processes all interacting at once, but Markov models are known to suffer from 'state space explosion', where the size of the model grows exponentially with every additional component that is modelled. This can risk making the models too large for efficient analysis of such complex systems.

2.3.2 Petri net models

Following the Markov models' limitations, over the past decade Petri nets (PN), whose structure are described further in Chapter three, have increasingly been used to model the asset management of components on the railway. The transition times to move between states can be sampled from distributions such as Weibull, allowing for the probability of a fault occurring to change over the asset's lifetime. Petri nets can also model processes simultaneously, unlike Markov models which typically assume events to happen successively.

In their paper, Zhang, Hu, and Roberts (2017), developed a Petri net to model the asset management of a section of plain line rail, through use of five sub-models: deterioration, inspection, lubrication, maintenance and renewal. Deterioration is defined as wear depth and isolated defects on the rail, which occur because of the independent variable MGT, not time. Within this PN, every transition has an associated firing time and probability of firing, which has been used to consider the impact of maintenance decisions, such as the effect of the periodic grinding strategy on the occurrence of surface-level defects, or the effectiveness of detecting defects. A renewal of the rail is triggered by a wear loss limit, w_{max}, being reached. The MGT that a section must receive before a renewal is necessary ($T_{renewal}$) is determined by the traffic wear rate, r_w, and the vertical grinding rate r_g (Equation 2.11):

$$T_{renewal} = \frac{w_{max}}{r_w + r_g} \tag{2.11}$$

Upon implementation of the model, using transition firing times found from previous research on degradation on the rail (including from Zhao et al. (2006) which fit Weibull distributions to describe the occurrences of three rail defects – flash butt weld defects, squats, and tache ovale defects), the values of r_w and r_g were deemed constant, reducing the stochasticity of the model and effectively imposing a fixed lifetime renewal strategy. Additionally, only the three types of rail defects were considered, although the authors argue that the structure of the developed PN allows for other defect types to be accommodated, if the corresponding maintenance actions are known.

Litherland and Andrews expanded on their previous reliability analysis (2023a) by developing a Petri net framework to model the asset management of an S&C unit by splitting it into nine components: ballast, bearers, check and wing rail, crossing, fastening, slide chair, stretcher bar, stock rail, and the switch rail (Litherland & Andrews, 2023b). Each component has its own sub-model containing the degradation and maintenance processes, which allows for a researcher to investigate the occurrence of failures between components. A full renewal of the entire S&C unit is required at the end of a fixed lifetime; however, some components (such as fastenings and bearers) have a renewal policy embedded in their maintenance strategy, due to the low cost and relative ease of replacement. A component replacement does not trigger a full unit renewal. This research is a rare example of a PN framework applied to S&C units, but it has several limitations. Rail-based components use the condition states 'working', 'minor defect', or 'safety limit exceeded' to represent the extent of degradation, instead of modelling the exact type of defect that occurred, such as a squat or lipping. Also, the assumption that every component requires a renewal at the same time (by the end of the fixed lifetime) is limiting, as components will degrade at different rates, and it may be more cost effective to replace them separately.

In later research completed by the authors (Litherland & Andrews, 2025), the developed S&C unit models were combined with models created for the plain line track, sleepers, and tunnels to create a whole system model. This time the renewal strategy was amended to a maintenance limit for the track sub-model which triggered a replacement or major refurbishment of the other assets. The model included a framework for opportunistic renewals based on the relative condition of the ballast and rails included within the track section model, but this still involved replacing every component at once even if others (such as the S&C) were not worn. Additionally, implementation of the model using a case study of a UK railway line took 6 days to run, making this renewal approach impractical for detailed analysis.

Two components of the S&C – the supplementary drive and the slide chair – were modelled using Petri nets in Sachan and Donchak's paper (2020). For both components, a complete renewal was scheduled at the end of a design lifetime, although the slide chair was able to be partially renewed as part of its maintenance policy. Unlike in Litherland and Andrews' research, as the developed PN are kept separate, the researcher can consider the cost benefits of replacing them independently. As Figure 2.5 shows, the developed model for the slide chair also does not consider the types of faults that occur but it can consider the severity of failures through transitions t2 and t3. If t3 fires before t2, implying that the slide chair is broken, a token is moved directly from place B to D, leaving place C empty so that minor repairs are ignored in favour of larger maintenance actions, such as a partial renewal. The model was then simulated to demonstrate its capabilities; however, the parameters used in the transition distributions were estimated, not obtained from historical failure data.

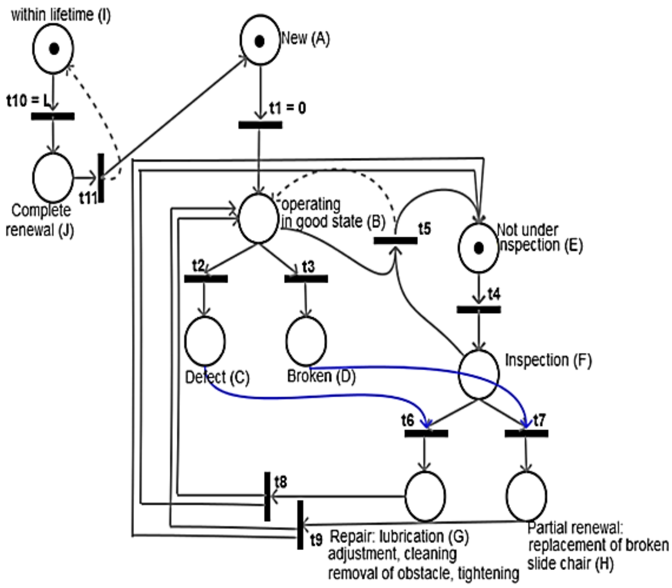

Figure 2.5: Slide chair Petri net model (Sachan & Donchak, 2020)

Saleh et al. (2023) used a PN to model the asset management of a 220-yard ‘poskey’, defined as a rail section that includes the ballast and plain line rails. The ballast degradation is represented by the SD of the vertical alignment over the length of the poskey, and is characterised into five states: excellent, very good, good, poor, and super red. Tamping, stoneblowing, and renewal are all considered maintenance actions for the ballast, depending on the level of degradation. For the plain line rail, degradation has been recognised as 12 distinct faults, that can be rectified through welding, grinding or replacement. The stacked probability of the maintenance action for each fault is given below in Table 2.4.

	squat	tache ovale	bolt hole	weld	other	RCF
Rerail	0.328	0.722	0.9	0.519	0.641	0.508
Weld	0.954	0.963	0.919	0.904	0.918	0.786
Grind or other	1	1	1	1	1	1
	wheel burn	Lipping	side wear	headwear	corrugation	unknown
Rerail	0.267	0.029	0.044	0.213	0.706	0.464
Weld	0.874	0.134	0.338	0.752	0.765	0.63
Grind or other	1	1	1	1	1	1

Table 2.4: Stacked probabilities for maintenance corresponding to different failure types (Saleh et al., 2023)

The renewal strategy implemented is if degradation of the poskey has required a renewal, or the poskey has reached a 200-year lifetime limit. Although significantly more failure types are considered for the plain line rail in this research compared to the previous PN developed for the rail by Zhang, Hu, and Roberts (2017), not every fault is considered separately and the severity of faults is grouped. For example, a squat of size 5mm and of 25mm both have a 32.8% likelihood of causing the plain line rail to be rerailed, which is unrealistic for modern practises.

Research by Audley and Andrews (2013b) focused on developing a Petri net approach for modelling a 220-yard section of ballast underneath a railway line. The condition of the ballast was defined by taking the standard deviation of the vertical alignment over the section, and assigning it to five states: very good, good, satisfactory, poor, or very poor. The PN structure utilised inhibitor arcs and enabler places to allow for the maintenance policy to be adjusted, changing the threshold of when the section requires maintenance (Figure 2.6).

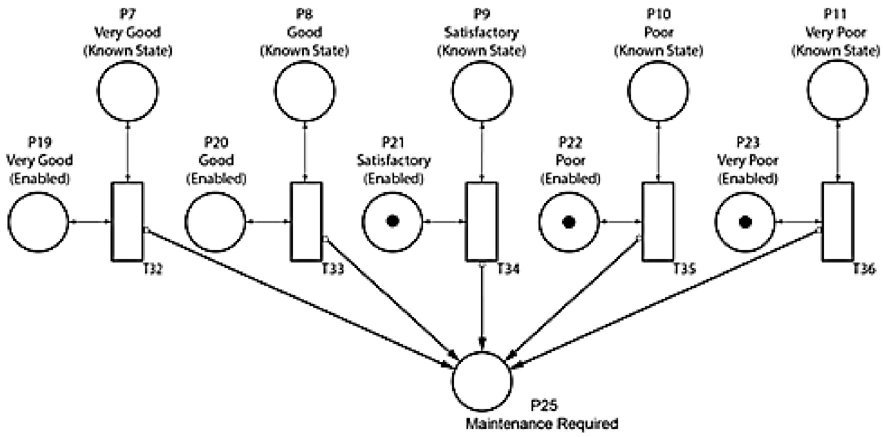

Figure 2.6: Decision making process for the threshold of maintenance (Audley & Andrews, 2013b)

The results from simulating the PN over a 50-year operational period using data from the UK's conventional speed rail network found that a ballast section within a 115-125mph zone was expected to require a renewal 3.4 times, whereas a section with speed limits between 80-110mph needed renewing only once. The modelled renewal strategy was based upon the frequency of maintenance actions – if maintenance is completed within a time (t) from the previous maintenance action, and this happens a consecutive number of times (n), a renewal is scheduled. However, as the research neglects the degradation of isolated vertical defects, which in turn can impact the frequency that maintenance is completed, the results found are unlikely to be realistic to actual engineering procedures.

In their thesis, Kumari (2019) developed a PN with four sub models: inspection, maintenance, degradation, and renewal, to model the asset management of a 200m section of ballast, using data from the Western Main Line in Sweden. A renewal is triggered when a ballast segment has reached a fixed lifetime of 30 years. The structure of the Petri net developed in this paper acknowledges that the ballast does not return to an 'as good as new' state after tamping has been completed, instead moving a token to a different state to represent that the ballast has been improved by tamping but that this has impacted its overall condition. Once within this state, the progression of degradation continues as normal. Kumari discussed the need for a line model, not just a single segment model, however the execution of this was severely limited. In the data analysis stage, consecutive sections sharing the same 'object type' – the feature that sits upon the ballast, e.g. plain line rail or S&C – were grouped and used to find the best fitting Weibull distribution for the transition times between degradation states. This clustering of identical ballast sections is the only consideration toward a line model, neglecting the opportunity to investigate asset management strategies such as opportunistic renewal. By contrast, in his thesis, Audley (2014) developed on his previous research (Audley & Andrews, 2013b) by using a PN methodology to create a line model for the ballast. The updated model considers opportunistic renewal by modelling the availability of renewal equipment and number of sections within a 3.2km region that require replacement, to predict the lifespan of a track section. This involves a far more complex PN structure than has been seen in his previous research but results in likely far more accurate predictions, as it considers the interactions between the track sections instead of the assumption that they exist in isolation. The author does note that it is computationally expensive, with the segment model running 30,000 times faster than the line model, making it inefficient for quick analysis or optimisation techniques.

The 220-yard section of ballast was also modelled via a PN structure in Fecarotti and Andrews' paper (Fecarotti & Andrews, 2020). Again, the standard deviation of the vertical alignment over the section length was used as the criterion of degradation and tamping or renewal are the only maintenance actions considered to rectify the condition of the ballast. The structure of the Petri net (shown in Figure 2.7) has allowed for three renewal strategies to be considered: the ballast section can be renewed when either a fixed lifetime has been reached, when the section has been tamped too many times (a maintenance limit), or when the sleepers present at the section also require renewal. The sleeper degradation sub model is separate (not seen in Figure 2.7) but is connected through place P37. The model was simulated for the length of a single lifetime of the asset, using estimated parameters for the probability distributions used in the transitions. Place P37 is defined in the paper as a conditional place for the three renewal transitions, T35, T36, and T37. Only one renewal strategy can be chosen at once – for example, if P37 is marked with two tokens, only T36 is enabled, and renewal is based on the condition of the sleepers. This either/or approach is unlikely to be seen in practice, as a track engineer should consider all definitions of wear before scheduling a renewal, and the predictions made from the model could be improved by allowing for multiple strategies to be considered concurrently.

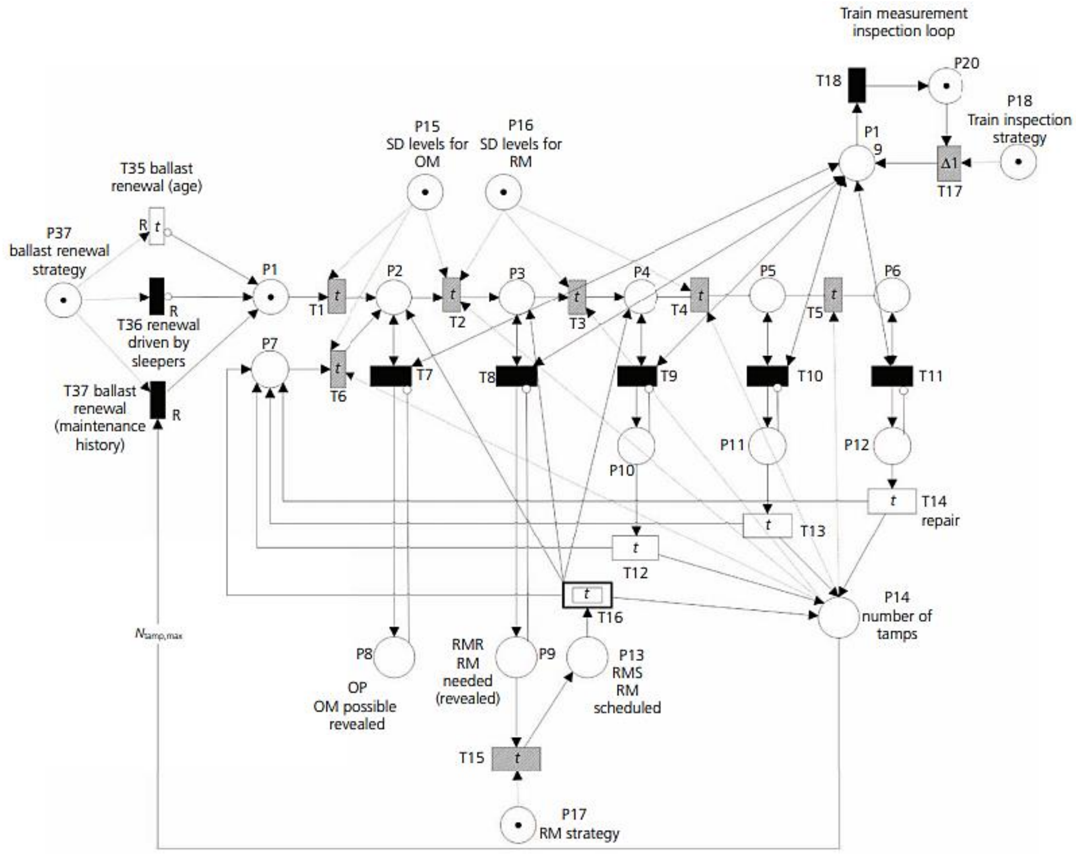

Figure 2.7: Renewal strategy for the ballast section Petri net (Fecarotti & Andrews, 2020)

In 2013, a PN approach was applied to a 1/8th mile section of ballast to model its degradation, inspection, maintenance and renewal processes (Andrews, 2013). The model used the SD of the vertical alignment of the section as an indicator of degradation, and the incorporated renewal strategy was a fixed lifetime of 30 years, although the author highlighted several other strategies that could also be appropriate. This model was built on in the later paper by Andrews, Prescott, and Rozières (2014). Monte Carlo simulation of the developed PN was used to analyse the impact of changing the lifetime from the baseline 30 years to either 20 or 40 years. 24,000 simulations were completed over a 180-year simulation period, and the results had converged to an acceptable limit after 6000 simulations. It was found that changing the fixed lifetime did not correspond with a noteworthy change in the proportion of time the section spent in a degraded state, implying that the ballast has not yet begun to experience significant degradation and that the lifetime limit should be larger than 40 years, or a different renewal strategy should be considered.

The condition of the ballast in this study is characterised into four states: good, minor degradation, major degradation requiring speed restrictions, and major degradation requiring line closure. This characterisation, as well as the ones used in previously discussed research (Audley & Andrews, 2013b; Saleh et al., 2023), all differ from one another, despite modelling the same asset and using the same parameter for degradation (SD of the vertical alignment over a 220-yard section). Different railway operators will have different thresholds and definitions for degradation; NRHS for example uses only three categories – good, satisfactory, and poor – to describe the condition of the vertical alignment's SD. Therefore, the previously discussed PN structures, and its associated results, may not be applicable to a researcher wishing to study a railway line that operates under a different definition of deterioration.

Unlike the previously discussed state-spaced modelling techniques, Petri nets have memory and choice, through strategic placing of tokens. This makes them ideal for the complicated decision-making processes that are inherent when operating a track asset. The discussed literature has demonstrated this capability, such as being able to choose the threshold for maintenance (Audley & Andrews, 2013b), acknowledging differing degradation rates between components (Litherland & Andrews, 2023b), and considering the availability of maintenance machinery (Audley, 2014). Their popularity may suggest PN's are the most appropriate modelling tool for this research, but they are not flawless. As argued by Zurawski and Zhou (1994), PN's are complex and difficult to construct, with researchers often required to develop a completely new PN every time a new feature is considered. Unless the distributions used for the transitions are estimated, PN also require previous analysis of the modelled track asset, making them labour intensive.

2.3.3 Summary of the literature surrounding asset management modelling

This section has reviewed studies focusing on applying a state-space model to the considered track assets and evaluating how they were used to predict when a renewal was required. The methodology varied between Markov models and Petri nets, with ultimately far more research conducted into PN analysis of asset management modelling. The issues highlighted in section 2.2.4, namely neglecting to consider the different failures and severities impacting degradation, were compounded in the asset management models. For example, no existing model considers both the isolated and standard deviation of geometry faults, making any renewal strategy beyond a fixed lifetime flawed as the actual management of the ballast is not accurately modelled.

Table 2.5 describes the considered renewal strategy in the reviewed literature. The decision of when an asset is eligible for a renewal will naturally impact its lifespan prediction, but most models have simply implemented a fixed lifetime, with very few examining the actual behaviour of the asset through methods such as a threshold on the number of maintenance interventions. A limited number of the studies allowed for a researcher to choose between renewal strategies, and no literature was found that could consider multiple strategies simultaneously.

Renewal strategy	**Literature**
Condition based	(Zhang et al., 2017)
Fixed lifetime	(Andrews, 2013; Fecarotti & Andrews, 2020; Kumari, 2019; Litherland & Andrews, 2023b; Prescott & Andrews, 2015; Sachan & Donchak, 2020; Saleh et al., 2023)
Maintenance frequency	(Audley & Andrews, 2013b; Fecarotti & Andrews, 2020)
Maintenance limit	(Fecarotti & Andrews, 2020; Litherland & Andrews, 2025)
Maintenance policy	(Hokstad, 2005; Litherland & Andrews, 2023b; Sachan & Donchak, 2020; Saleh et al., 2023)
Opportunistic	(Audley, 2014; Fecarotti & Andrews, 2020; Litherland & Andrews, 2025)

Table 2.5: Implemented renewal strategies within the considered literature

Numerous research has been conducted to prove that in practise, a fixed lifetime is an overly simplistic predictor for asset lifetime; in 2013, a decision support system created by Guler (2013) was built around surveys completed by railway experts, who identified multiple parameters including the tamping frequency and total number of tamps, for influencing ballast renewal. Research by Ait-Ali et al. (2024) using a lifecycle cost model found that after 10 years of service, the cost of maintaining the studied crossing panels per year is more expensive than a renewal. Also, Bai et al. (2019) proved that modelling using a renewal strategy based on a threshold for failures was a more realistic reflection of the actual behaviour of rails on the Beijing Metro compared to implementing a simple fixed lifetime based on the passing MGT. From this, it is evident that the developed asset management model needs to acknowledge the condition and economic viability of the railway when deciding if a renewal is required, and should allow for multiple parameters defining wear-out to be considered at once.

2.4 Conclusions

This research aims to predict when a track asset should be replaced, which requires models to be developed that can predict future degradation and understand how maintenance can be applied to control the asset's condition. The review has appraised various methodologies to achieve this aim, which are summarised below:

<u>Degradation modelling</u>

- Deterministic models can be based on results from laboratory experiments or estimated, requiring very little historical failure data to be functional
- If the input is unchanged, deterministic models will output the same results every time making them computationally inexpensive, but unable to model stochastic behaviour
- A reliability-based model can fit a continuous probability distribution to model the stochastic behaviour of degradation, which is easily applied alongside an asset management model
- Investigating different parameters that influence degradation requires a completely new reliability-based analysis for each subgroup
- Artificial neural networks can consider multiple parameters for degradation at once and evaluate their relative importance.
- Training an artificial neural network requires a time-consuming approach to find the ideal learning parameters.
- Both a reliability-based model and artificial neural networks will suffer from bias and random variation if sample sizes are low – considered as consisting of less than 30 uncensored times-to-failure for the reliability-based model (Indrayan & Mishra, 2021), or less than 10 times the number of weights used in the net for artificial neural networks (Alwosheel et al., 2018).

The HS1 line has been fully operational since 2007. Although this is still relatively young, evidenced by the lack of significant renewals, several failure modes have occurred frequently enough to create a significantly sized fault database. This justifies the use of reliability-based models or artificial neural networks over a deterministic model, as degradation is known to be stochastic in nature. The reliability-based methodology will be mainly considered to model degradation, as the resulting fitted continuous probability distribution can be used to indicate characteristics of a fault (such as when it is expected to appear during the operational period), and be easily implemented into an asset management model. However, given that the studied railway line is roughly 109km long with varying features throughout such as speed restrictions and traffic levels, an artificial neural network will also be employed to investigate how different factors can impact degradation rates and inform of the areas of the track that are driving renewals.

Asset management modelling

- Markov models are a relatively simple approach to modelling asset management
- Transition rates between states must be constant for Markov models.
- Markov models also increase in size exponentially as the complexity of the model also increases.
- Petri nets can incorporate more complex firing times for transitions, such as those modelled by the Weibull distribution, and do not suffer from 'state space explosion' when additional components are added.
- Markov models consider events successively whereas Petri nets can consider events simultaneously
- Petri nets are potentially computationally expensive

Ultimately, the requirements to make a Markov model functional were deemed too prohibitive for full scale research, and as such, they were not considered as a viable state-space modelling option for this research. Petri nets utilising degradation predictions found from reliability analysis, and later artificial neural networks, will therefore be used to predict when a renewal is necessary for the four considered assets. The ballast has been reasonably investigated using this technique, but never with a PN structure that reflected NRHS conduct, and therefore any results found so far are not applicable to the HS1 line. Both the plain line rail and S&C are under researched, with the current models not considered in-depth enough to adequately reflect the entire degradation and maintenance processes.

Reviewing the literature has also highlighted two areas where current models are limited, specifically:

1) Excluding or grouping failures impacting the asset
2) Implementing simplistic renewal strategies with no option for choosing between or combining multiple criteria to create a renewal strategy

Addressing these limitations will improve the ability of a model to predict when the lifespan of a track asset has been exhausted and should be replaced. The results found can be used to refine NRHS's current 40-year renewal strategy and implemented back into the original OMRC model to calculate the estimated renewal cost for the next control period.

Chapter 3

Summary of the methodology and concepts utilised

3.1 Introduction

Throughout this research, many mathematical concepts and methodologies are introduced to achieve the aim of predicting when high speed railway track requires a renewal. This chapter explains the theory and structure of the methods considered. The content is discussed in order relevant to their introduction in this thesis.

3.2 Censored data

Times-to-failure are collected and used to analyse the behaviour of different faults that affect the considered assets. Due to the nature of the historical failure and maintenance records provided, some of these times will be considered 'censored'. Censored - or incomplete - data is where the exact time to an event is unknown. There are several ways in which this can happen:

1) <u>Left censored</u>
   Data is considered left censored if the event happened before the data was collected. In context of this research, left censored data would include faults that have occurred on the track before the observation period of the provided data began.
2) <u>Right censored</u>
   Data is considered right censored if the event happens after the observation has ended, hence the true time-to-failure is longer than what is inferred. This could

include faults that occur after the observation period of the provided data has ended, or assets that have yet to experience the fault.

3) Interval censored

   Data is interval censored if the event occurs within a known interval of time, but the exact date is unknown. This is the case for many faults on a railway line as they are only detected once an inspection occurs.

Some studies exclude censored data in favour of wholly complete data, however doing so will reduce the size of the available datasets, potentially leading to biases and impacting the results. Regardless, the impact of left censored data had to be excluded as this could not be inferred from the available information, and an assumption was made that any faults recorded in the failure history occurred the day of the inspection, hence there is no interval censored data. This assumption was necessary in order to provide complete data to be used in the analysis.

Additionally, there were many occasions where the exact time of the failure was known, but the installation of the unit or any previous restorative maintenance was unknown as it happened before the observation period of the data. A similar problem is prevalent in Novak et al.'s research (2017), who argued that because the actual service time of the studied assets were longer than what was observed, this could be treated as right-censored data (Figure 3.1). Therefore, all censored data in this research is considered right censored and will be included in the data analysis.

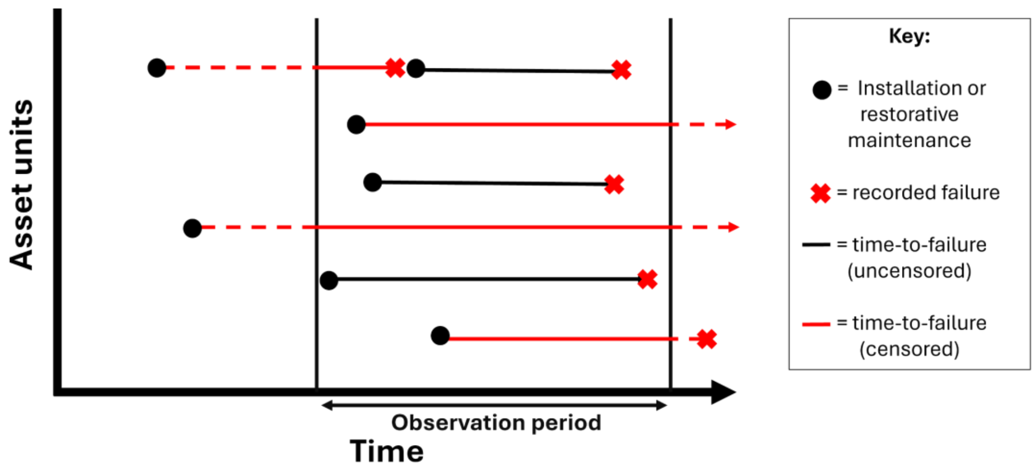

Figure 3.1: Collecting censored data

3.3 Continuous probability distributions

Continuous probability distributions are fitted to a dataset of times-to-failure to describe the behaviour of each fault. Four of the most common distributions in reliability engineering were considered: the one parameter exponential distribution, the normal distribution, the two-parameter lognormal distribution, and the two-parameter Weibull distribution. The choice of distribution and their corresponding parameter values each imply different characteristics of the fault, including when it is expected to occur during the asset's operational period.

3.3.1 One-parameter exponential

The probability density function (PDF) of the one-parameter exponential distribution is:

$$f(t) = \lambda e^{-\lambda t} \tag{3.1}$$

Where λ is the degradation rate

λ can be inverted to find the expected number of faults per unit of time. If an exponential distribution is chosen as the most suitable to fit the fault lifetimes, it is implied that the likelihood of a fault is constant throughout the assets lifespan as there is no association between failure and age (Ahsanullah & Hamedani, 2010). Degradation is expected to occur randomly if the exponential distribution is selected.

3.3.2 Normal

The normal distribution is best used to represent times-to-failure that are very symmetric around a mean. The PDF of the normal distribution is:

$$f(t) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{t-\mu}{\sigma}\right)^2} \tag{3.2}$$

Where μ and σ are the mean and standard deviation of the lifetimes respectively

The degradation rate of a fault distributed by a normal distribution is expected to be low initially, then rapidly increase as the time increases towards the mean parameter. After this is passed, the degradation rate begins to slow again (whilst still increasing), making the normal distribution unsuitable for modelling early installation or late wear-out faults. Additionally, there is a non-zero chance of the normal distribution predicting a negative time-to-failure which is impossible in dynamic systems.

3.3.3 Two-parameter lognormal

The exponential of a normally distributed random variable will result in a random variable that is best described through the lognormal distribution. The lognormal PDF is:

$$f(t) = \frac{1}{t\sigma'\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\ln(t)-\mu'}{\sigma'}\right)^2} \tag{3.3}$$

Where μ' and σ' are the mean and standard deviation of the natural logarithms of the lifetimes respectively

Unlike the normal distribution, the random variables of the lognormal distribution are bounded below by zero. Datasets with a large right skew of times-to-failure are typically best represented via the lognormal distribution, as the skewness increases as the shape parameter σ' increases for a constant scale parameter μ' (Ginos, 2009). If the shape parameter is significantly larger than 1, this implies that the degradation rate of the modelled fault is expected to increase rapidly immediately after installation/correction of the asset then decrease sharply after the peak has been reached, mimicking early-onset faults.

3.3.4 Two parameter Weibull

One of the most common distributions used in reliability modelling is the Weibull distribution, who's PDF is as follows:

$$f(t) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1} e^{-\left(\frac{t}{\eta}\right)^{\beta}} \tag{3.4}$$

Where β is the shape parameter and η is the scale parameter.

The scale parameter η represents the time at which 63.2% of the failure lifetimes have been reached, but the Weibull distribution is particularly versatile as its shape parameter β influences the degradation rate. If $\beta < 1$ the rate is expected to decrease over time, if $\beta = 1$ the rate remains constant and the distribution reduces to the exponential distribution, and if $\beta > 1$, the degradation rate increases over time, suggesting the fault is correlated with wear-out. This is demonstrated in Figure 3.2, taken from a thesis studying the Weibull distribution's use in reliability engineering.

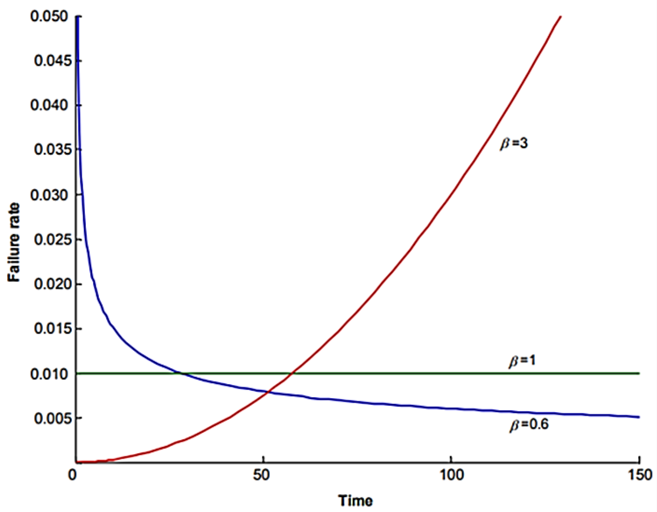

Figure 3.2: Impact of the shape parameter of the Weibull distribution on the degradation rate (Yong, 2004)

3.4 Parameter estimation and goodness of fit methods

A parameter estimation method and goodness of fit test/s are used to evaluate which of the previously described continuous probability distributions, with which parameters, best fit a dataset of times-to-failures. There are many different types of both, but only those considered are described below.

3.4.1 Parameter estimation

RRX and RRY

Rank regression (also known as least squares estimation) is a method used to estimate parameters for a chosen distribution, by transforming its cumulative density function (CDF) into a linear regression line, of the form:

$$y = mx + c \tag{3.5}$$

Where m and c are the gradient and intercept of the line respectively.

By finding the values of m and c that minimise the sum of the squared error between the values of the observed fault lifetimes (residuals) and the corresponding values on the regression line, then transforming m and c back into the distribution's original parameters, suitable parameters can be estimated. Minimising the sum of the squared error can either be performed on the x axis – rank regression on X (RRX), or on the y axis – rank regression on Y (RRY). Rank regression is easily computable however Equation 3.5 cannot incorporate data that is censored, and if this is a high percentage of the dataset, these missing values can reduce the applicability of any results (Schneider et al., 2010).

MLE

Maximum likelihood estimation (MLE) is a separate way of estimating the parameters of a distribution and unlike the RRX or RRY, can be applied when datasets are heavily censored. Using MLE during the analysis when there is not a large sample of times-to-failure, however, can bias results, as it tends to 'overfit' the data (Widyaningsih et al., 2023).

The MLE works by optimising the parameters of the chosen distribution so that they maximise the likelihood that the assumed distribution will produce values equal to the ones found in the observed data. Right censored data is included by incorporating the reliability function in the likelihood calculation, described in Equation 3.6:

$$L(\boldsymbol{\theta}\,|\boldsymbol{UC},\boldsymbol{C}) = \prod_{i=1}^{N} f(UC_i;\ \boldsymbol{\theta}) \times \prod_{j=1}^{M}[1 - F(C_j\ ;\ \boldsymbol{\theta})] \tag{3.6}$$

Where $\boldsymbol{\theta}$ is a vector of the distribution's parameters, $\boldsymbol{UC}$ is a vector of the uncensored times-to-failure, $\boldsymbol{C}$ is a vector of the right-censored times-to-failure, N is the number of uncensored fault lifetimes, M is the number of right-censored fault lifetimes, and f and F are the PDF and CDF of the considered distributions respectively

3.4.2 Goodness of fit

When the parameters for each continuous probability distribution have been estimated, the Akaike's Information Criteria (AIC) and Bayesian Information Criteria (BIC) are employed to assess how well the distribution models the behaviour of the times-to-failure dataset. Both goodness-of-fit tests incorporate the log-likelihood of the selected distribution allowing for the impact of censored data to still be considered.

The AIC is defined as in Equation 3.7:

$$AIC = -2\log(L) + 2k \tag{3.7}$$

Where L is the maximum likelihood of the distribution and k is the number of parameters used in the distribution.

The BIC's formula is slightly adapted to consider the amount of failure lifetimes in the dataset (n):

$$BIC = -2\log(L) + (\ln(n) * k) \tag{3.8}$$

The two are frequently recommended to be used in conjunction to evaluate the best fitting distribution (Kuha, 2004). A property of the AIC and BIC is that their scores are meaningless when considered in isolation, but the lowest score between a selection of distributions fitted to the same dataset implies the best fit. An additional property is that if the score for a distribution is within two units of the score for another distribution, both are considered as effective at fitting the dataset (Symonds & Moussalli, 2011). This can be used to consider the relative importance of the best-fitting distribution.

3.5 Petri nets

The literature review indicated that Petri nets (Reisig, 2012) are a suitable modelling technique for investigating the stochastic degradation behaviour of railway assets. These are weighted, bi-partite graphs, first developed in 1962 by Carl Adam Petri as a way of modelling chemical processes. Given their flexible nature, they are frequently used in the field of engineering, business, computer sciences and more to model dynamic systems.

3.5.1 Components

The basic structure of a Petri net consists of four components, demonstrated in Figure 3.3:

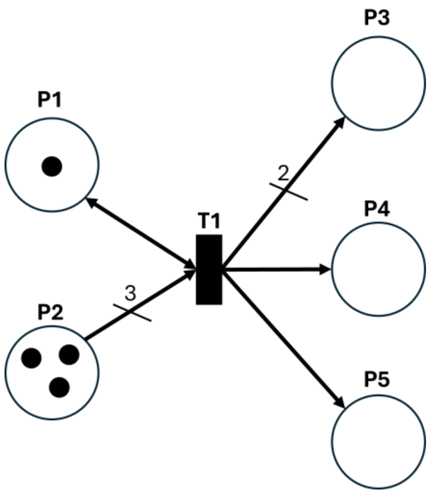

Figure 3.3: Components of a Petri net

1) Places: used to represent the condition of an asset. They are graphically represented by a circle.
2) Transitions: used to indicate actions, such as maintenance. They are represented by a black rectangle.

3) Tokens: reside in a place to indicate the current condition of an asset. They are denoted by a small black circle.
4) Arcs: represented as directed arrows. They connect places to transitions (input arcs) and transitions to places (output arcs). They are characterised by a multiplicity (Goutsias & Jenkinson, 2013), or an associated 'weight', which is shown by a positive integer located next to the arc. If no number is present, the multiplicity is one.

The movement of tokens round the Petri net represents the asset management dynamics. Statistics are collected based on the location of tokens in the net, or the duration of time a token spends in a place. Movement is caused by transitions 'firing', the rules of which must first be defined.

3.5.2 Firing rules

A transition has two firing rules:

1) A transition is only able to fire after being enabled; every input place leading into the transition must hold a minimum number of tokens consistent with the weight of the associated input arc. Once enabled, the transition will fire after a predetermined length of time t has passed, known as the firing time.
2) Once a transition fires, a given number of tokens equal to the weight of the associated input arc are removed from the input places, and a number of tokens equivalent to the weight of the output arc are placed in the output places. In this way, Petri nets are not conservative - a transition will not necessarily move the same number of tokens from input to output.

The firing process is shown in Figure 3.4.

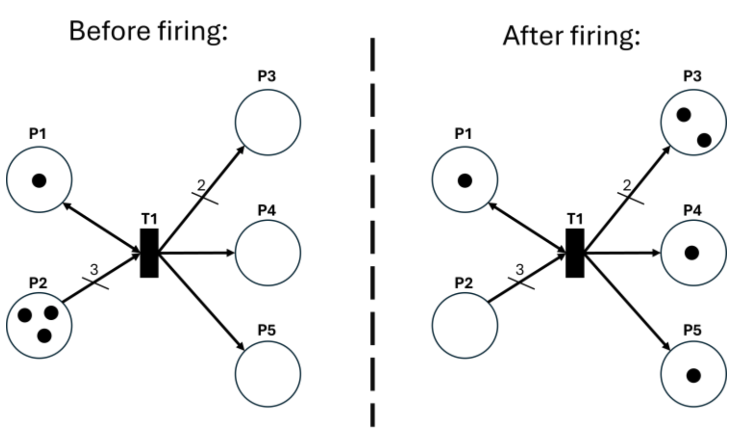

Figure 3.4: Example of a transition firing

Note the double-sided arc leading from P1 into the transition, which means that a token is placed back in P1 after firing. The firing time of a transition can be deterministic (a fixed time), or stochastic (sampled from a continuous probability distribution). When a transition fires, it's firing time is resampled.

3.5.3 Additional features

A system of places, arcs, tokens, and transitions alone is typically not sufficient to represent some of the more intricate systems in asset management – for example, the inspection interval may decrease when the asset is in a known degraded condition. This will require the transition relating to inspection intervals to adapt its firing time depending on the marking of places in the net. The additional features that can be included in the Petri net structure are described below:

1) Inhibitor arcs

   Inhibitor arcs (Christensen & Hansen, 1993) will prevent a transition from firing if the weight of the inhibitor arc is satisfied by the number of tokens present in the input place. They can lead only from a place into a transition and are graphically represented by a traditional arc with a small circle at the end. An example of inhibitor arcs is shown in Figure 3.5, where even though the firing time of t has passed and T1 is enabled, a token was not moved from place P1 into P5, as the token present in P2 inhibited T1. Note that because the multiplicity of the inhibitor arc connecting P4 to T2 was higher than the number of tokens in P4, T2 was still able to fire after time t had passed.

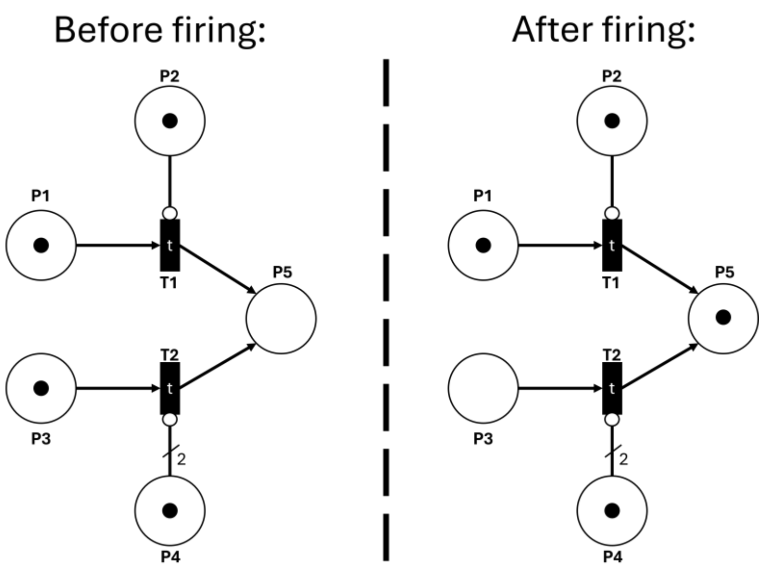

Figure 3.5: Inhibitor arcs

2) Place dependent transitions

Place dependent transitions (or conditional transitions (Andrews, 2013)) can modify their firing time by sampling from different distributions depending on the number of tokens present in associated places. They are graphically represented via a dashed line connecting the transition to the relevant places. An example of place dependent transitions is given in Figure 3.6, where the firing time t for transition T1 is sampled from a Weibull distribution with an eta parameter of 400. After time t has passed, the firing time t' is now sampled from a Weibull distribution with an eta parameter of 300, due to 1 token residing in P3. If an additional token is placed in P3, T1 will change to a constant transition with a firing time of 200 days.

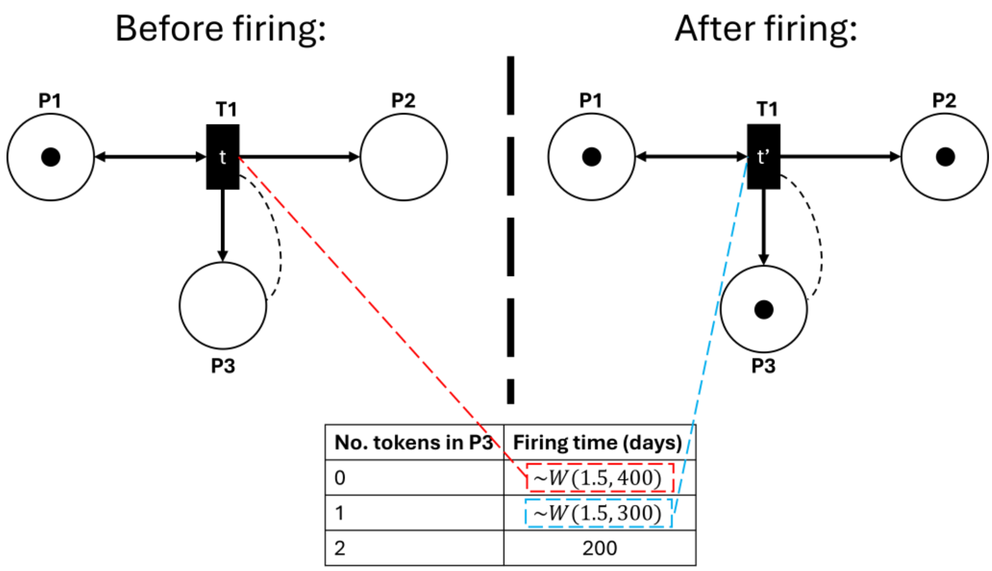

Figure 3.6: Place dependent transitions

3) Reset transitions

A combination of strategically placed arcs and transitions can restore the PN back to its original place markings when required, however this can make the net needlessly large and expensive to compute. Reset transitions on the other hand will automatically mark a listed number of places back to their initial condition upon firing, whilst still functioning as a regular transition to move tokens between places. For this research, reset transitions have been adapted to also restore other transitions, so that firing times are resampled from the relevant distribution if the transition is reset. This was to prevent transitions that were enabled when the net was reset from having an unnaturally short firing time when enabled again. Reset transitions are denoted by a striped, blue transition.

3.6 Monte Carlo simulation

Stochastic transitions within a Petri net are sampled from a given distribution. The probability (P) of the transition firing by time t is found from the distribution's cumulative distribution function (CDF):

- Exponential

$$P = 1 - e^{-\lambda t} \tag{3.9}$$

- Normal

$$P = \Phi\left(\frac{t - \mu}{\sigma}\right) \tag{3.10}$$

Where $\Phi()$ is the 'Phi' function (Marsaglia, 2004).

- Lognormal

$$P = \Phi\left(\frac{\ln(t) - \mu'}{\sigma'}\right) \tag{3.11}$$

- Weibull

$$P = 1 - e^{-\left(\frac{t}{\eta}\right)^{\beta}} \tag{3.12}$$

For the Weibull and exponential distributions, the firing time can be derived by rearranging Equations 3.9 and 3.12 to find a formula for t, which is then populated by a randomly generated probability P between 0 and 1. A similar approach cannot be as easily computed for the normal or lognormal distribution as there is no simple inversion of the Phi function, but many statistical software such as Excel are able to derive t through built-in functions.

For every randomly generated value of P, the firing time t will vary to represent the stochastic nature of asset degradation and management. However, this causes a high uncertainty in the found result from simulating the Petri net. Repeated random sampling (known as Monte Carlo simulation (Mooney, 1997)) is applied to find a range of outcomes, and statistics regarding the average performance of the model are collected. By the law of large numbers, repeated simulation of a Petri net with stochastic transitions will result in the average performance moving closer to the true value, making the results found more reliable.

3.7 Artificial neural networks

Artificial neural networks (Zou et al., 2008) are a form of machine learning, that mimic the structure of neurons in the brain to identify patterns in a dataset and predict future occurrences of events. There are multiple forms of artificial neural networks, including recurrent or convolutional, but the focus of this research is on multilayer perceptron networks, henceforth known as ANN. Additionally, all ANN considered will be 'supervised' as it aims to use known characteristics of the dataset to predict an output.

3.7.1 Structure

Artificial neural networks consist of neurons, arranged in layers, known as the input layer, hidden layer/s, and the output layer. Every neuron is connected to all the neurons in the following layer, and every connection has an associated weight. A bias is also present in every layer bar the output layer. Information is placed in the input layer, then transformed and processed through the hidden layer/s until values are inserted into the output layer, to give the prediction made by the ANN. An example of the structure of an ANN is given in Figure 3.7.

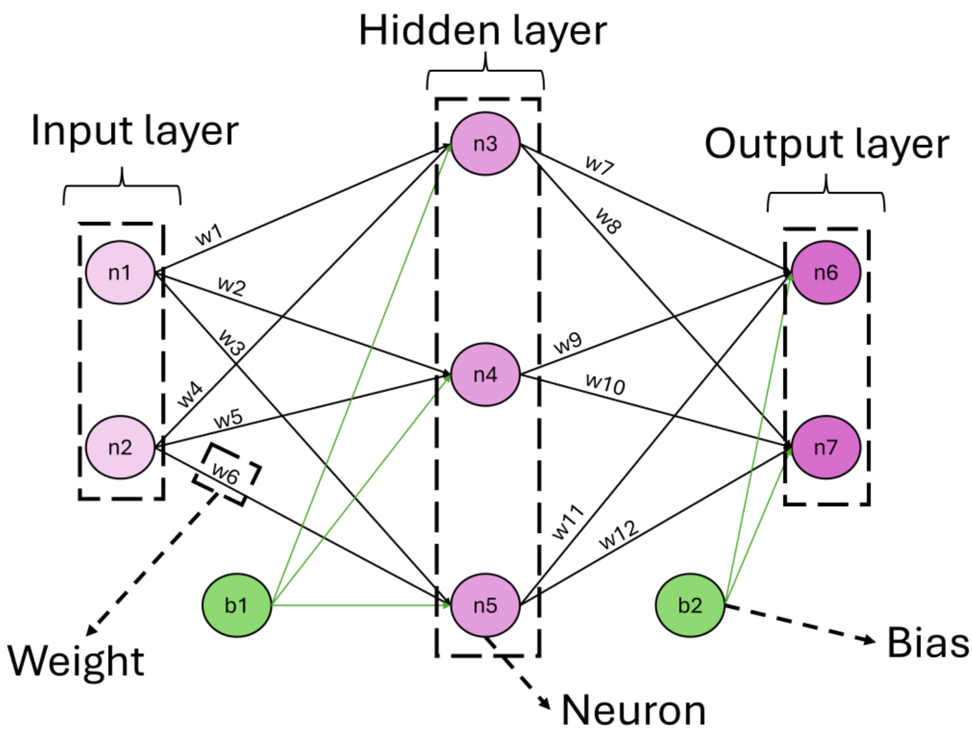

Figure 3.7: Structure of an artificial neural network

Before an ANN can be used to make predictions, it must first be trained. The available data – consisting of the values input into the ANN, and the known values that should be predicted in the output layer – is split into two datasets, training and testing. The relative size of each dataset can impact the ANN's effectiveness (Nguyen et al., 2021), but a 70:30 ratio is assumed.

3.7.2 Normalisation

Before the training or testing data is considered in the artificial neural network, all inputs and expected outputs are first scaled to equal magnitude using a normalisation technique, to reduce errors and runtime (Sola & Sevilla, 1997). Max-Min normalisation (Equation 3.13) is used throughout:

$$x_{norm} = \frac{x - \min(x)}{\max(x) - \min(x)} \tag{3.13}$$

Where x is the original variable value.

If the dataset contained extreme outliers, Equation 3.13 could result in skewed values hence 'clipping' the data beforehand was employed. For each input, a range was defined between 1.5 times the value of the interquartile range added to the upper quartile or subtracted from the lower quartile (known as fences). Any value outside this specified range was replaced by the maximum or minimum fence value. Once scaled, the training dataset is used to train the ANN until it can make accurate predictions. This is completed in three stages: forward propagation, calculating loss, and back propagation.

3.7.3 Forward propagation

The process of forward propagation is how the ANN takes the input information and uses it to make a prediction. Each neuron stores information regarding the input value (n_c) and an output value ($f(n_c)$). For neurons in the input layer, the input and output value are both set equal to the initial values used for processing. The bias is a constant, with an initial value of 1. To calculate the input value for neurons in the hidden layers and beyond, the output values from neurons in the previous layer and the relevant weights connecting the neurons in the previous layer to the neuron in the current layer are transformed into a single value using Equation 3.14.

$$n_c = \sum_{p=1}^{m} w_p f(n_p) + b_p \tag{3.14}$$

Where n_c is the input value of the neuron in the current layer, m is the number of neurons in the previous layer, w_p is the weight of the connection between the pth neuron in the previous layer and n_c, $f(n_p)$ is the output value of the pth neuron in the previous layer, and b_p is the bias of the previous layer.

For example, in Figure 3.7, the input value of n_3 is set equal to $w_1 f(n_1) + w_4 f(n_2) + b_1$. The summed value is then passed through an activation function, to transform the data into an output value that would be passed onto the next layer of the ANN. There are many activation functions that can be used such as:

- Binary activation function

$$f(x) = \begin{cases} 0 & for\ x < 0 \\ 1 & for\ x \geq 0 \end{cases} \tag{3.15}$$

- Linear activation function

$$f(x) = x \tag{3.16}$$

- ReLU activation function

$$f(x) = \max(0, x) \tag{3.17}$$

- Sigmoid activation function

$$f(x) = \frac{1}{1 + e^{-x}} \tag{3.18}$$

- Tanh activation function

$$f(x) = \frac{(e^{x} - e^{-x})}{(e^{x} + e^{-x})} \tag{3.19}$$

The aim of an activation function is to introduce non-linearity so that more complex patterns in the dataset can be detected. The choice of function is based on the features of the available dataset; for example, classification problems often perform better with the sigmoid activation function, whereas ReLU is typically chosen for most other prediction analyses (Sharma et al., 2020). Often, a different activation function is chosen for the output layer compared to the hidden layer/s, as the input tends to become very large in the hidden layers, making a lot of activation functions unsuitable. Forward propagation is repeated through the layers of the network until finally values are placed in the output layer. The output values are unscaled and used as the predicted results, hence all information is processed in a singular direction.

3.7.4 Calculating loss

During training, it is unlikely that the predictions described in the output layer will be accurate to their true value. A cost function is used to evaluate the performance of the ANN by comparing the predicted results to the actual results from the data. There are several cost functions that can be utilised, such as:

- Mean Absolute Error (MAE)

$$\frac{1}{o}\sum_{i=1}^{o}|y_i - f(n_i)| \tag{3.20}$$

Where o is the number of output neurons, y_i is the actual result from the data, and $f(n_i)$ is the predicted result from the artificial neural network.

The MAE is good for ignoring the impact of a few predicted outliers and providing an overall measurement of the model's performance.

- Mean Squared Error (MSE)

$$\frac{1}{o}\sum_{i=1}^{o}(y_i - f(n_i))^2 \tag{3.21}$$

Contrary to the MAE, the MSE is good for recognising the impact of outlier predictions and can account for them if they occur.

3.7.5 Back propagation

Once the prediction error has been obtained from the cost function, the final stage of training the ANN is to update the weights and biases, in a method known as backpropagation. Backpropagation works by calculating the gradient of the error (E) with respect to the weight (w):

$$\frac{dE}{dw} \tag{3.22}$$

This requires knowledge of the chain rule, demonstrated in Equation 3.23:

$$\frac{dy}{dx} = \frac{dy}{du} * \frac{du}{dx} \tag{3.23}$$

In the following demonstration, it is assumed that the ANN has used a sigmoid activation function throughout and the MSE cost function to measure error.

The backpropagation calculation varies depending on whether the gradient is calculated for weights leading into the output layer or into a hidden layer. Initially, the output layer is discussed. The final error, E, does not directly relate to the individual weights, such as w_7 in Figure 3.7.

Instead, the chain rule can be used to split Equation 3.22 into three derivatives, to calculate the error:

$$\frac{dE}{dw_p} = \frac{dE}{df(n_c)} * \frac{df(n_c)}{dn_c} * \frac{dn_c}{dw_p} \tag{3.24}$$

As MSE is used to calculate the error, E will be:

$$E = \frac{1}{o}\sum_{i=1}^{o}\left(y_c - f(n_c)\right)^2 \tag{3.25}$$

Therefore, the derivative of the error at $f(n_c)$ is:

$$\frac{dE}{df(n_c)} = \frac{2}{o}\,(f(n_c) - y_c) \tag{3.26}$$

As the sigmoid function is chosen as the activation function, $f(n_c)$ is:

$$f(n_c) = \frac{1}{1 + e^{-n_c}} \tag{3.27}$$

Hence:

$$\frac{df(n_c)}{dn_c} = \left(\frac{1}{1 + e^{-n_c}}\right) * \left(1 - \frac{1}{1 + e^{-n_c}}\right) \tag{3.28}$$

Finally, from forward propagation, n_c is equivalent to

$$n_c = \sum_{j=0}^{p} w_p f(n_p) + b_p \tag{3.29}$$

Hence:

$$\frac{dn_c}{dw_p} = f(n_p) \tag{3.30}$$

Substituting the found derivatives into Equation 3.24 gives:

$$\frac{dE}{dw_p} = \frac{2}{o}\,(f(n_c) - y_c) * \left(\frac{1}{1 + e^{-n_c}}\right) * \left(1 - \frac{1}{1 + e^{-n_c}}\right) * f(n_p) \tag{3.31}$$

Equation 3.31 can be used to calculate the gradient of the error for every weight connecting neurons into the output layer.

If the weight is connected to a neuron in the hidden layer, finding $\frac{dE}{df(n_c)}$ requires extra calculations, as a change in the value $f(n_c)$ alters the input into the neurons for the subsequent layer, which in turn affects the value of the error. This is completed by splitting the partial derivative of the error into the respective contributions from the next layer:

$$\frac{dE}{df(n_c)} = \sum_{s=1}^{n} \frac{dE_{n_s}}{df(n_c)} \tag{3.32}$$

Where n is the number of neurons in the subsequent layer

The chain rule is used to separate the derivatives to make them easier to calculate:

$$\begin{aligned} \frac{dE_{n_s}}{df(n_c)} &= \frac{dn_s}{df(n_c)} * \frac{dE_{n_s}}{dn_s} \\ &= \frac{dn_s}{df(n_c)} * \left(\frac{dE_{n_s}}{df(n_s)} * \frac{df(n_s)}{dn_s}\right) \end{aligned} \tag{3.33}$$

Again, due to forward propagation:

$$\frac{dn_s}{df(n_c)} = w_c \tag{3.34}$$

And $\left(\frac{dE_{n_s}}{df(n_s)} * \frac{df(n_s)}{dn_s}\right)$ can be found by multiplying the previously found values of Equations 3.26 and Equation 3.28, hence the backpropagation algorithm can update all weights and biases of the ANN.

The method described in Equations 3.22-3.34 indicates the direction the weight and biases should correct update in in order to minimise error – a positive gradient suggests the weight should decrease, and conversely if the gradient is negative, the weight should increase in size. If the gradient is zero, the weight is the optimal size to minimise the error.

Once it is known what direction the weight should adjust, the amount that the network adjusts the rate is given by an optimisation algorithm. For this research, the stochastic gradient descent optimisation algorithm was chosen (Haji & Abdulazeez, 2021), which moves the weights in accordance with Equation 3.35:

The distance that the weight moves is known as dx:

$$dx = \alpha \times \frac{dE}{dw_p} \tag{3.35}$$

Where α is a predetermined learning rate

The value of the updated weight after backpropagation is therefore:

$$w_{p_{new}} = w_{p_{old}} - dx \tag{3.36}$$

The updated weights move towards the local minima until a threshold for error has been reached, and the ANN is deemed to be predicting accurate results.

Choosing the value of α is imperative to how the artificial neural network makes predictions. Too large, and the weights will vary wildly each iteration, possibly leading to oscillating predictions; too small, and the network will take a long time to converge or potentially will not converge at all.

3.7.6 Overtraining

Using the training data to complete forward propagation, calculate the loss, then perform back propagation will adjust the weights to improve the accuracy of predictions made from the artificial neural network. The previously unseen testing data is then passed through as inputs to the ANN, implemented with the optimised weights, and forward propagation is completed until a prediction is made in the output layer, after which the model's performance can be evaluated. This process is completed multiple times using a variation of ANN structure, learning rates, activation functions etc, until the optimal methodology is found for the specific problem.

The number of epochs can also impact the ability of an ANN. An epoch is defined as being the number of iterations that the training data is passed through the structure. Increasing the number of epochs could increase the accuracy of the predictions, however too many iterations and the ANN could risk learning the statistical noise in the training dataset. This will improve the ANN's prediction error for the training error, but it may be unable to generalise to unseen data, such as that present in the testing dataset. This is corrected by 'early stopping', where the ANN stops training when the performance error of the testing set is higher than that of the previous epochs.

Once the model has been sufficiently trained, by optimising its training parameters and preventing overtraining, its performance can be evaluated using the R^2 value, described in Equation 3.37:

$$R^2 = 1 - \frac{\sum_{j=1}^{n}\left(y_j - f(n_j)\right)^2}{\sum_{j=1}^{n}\left(y_j - \bar{y}\right)^2} \tag{3.37}$$

Where n is the size of the testing dataset, y_j is the actual results from the testing data, $f(n_j)$ is the predicted result from the artificial neural network, and $\bar{y}$ is the average of the actual results of the testing data

R^2 is used to evaluate how well residuals sit along a fitted line, in this case how well the predicted results align with the actual results. The closer to 1, the better the fit indicating high performance, with an R^2 value of 0 implying that simply finding $\bar{y}$ is as effective of a prediction as the current model (Saunders et al., 2012).

3.8 Summary

This chapter has discussed the various concepts that will be used throughout this thesis. This includes using censored data in the reliability analysis, using parameter estimation and goodness of fit methods to fit suitable continuous probability distributions to datasets, creating models from Petri nets, simulating the models using Monte Carlo simulation, and finally developing an artificial neural network to predict degradation trends for a track asset. This chapter is to be used as a reference point for any methodology developed in this research.

Chapter 4

Reliability analysis of the studied track assets

4.1 Introduction

To predict the expected lifetime of high-speed railway track, it is vital to first understand how it degrades. The railway can experience multiple failure types that occur independently from one another and may be associated with the life stage of the asset e.g. during the initial 'burning-in' period, or because of late-stage wear-out. The reliability analysis will be based on fitting continuous probability distributions to the historic failure and condition data. The resulting distributions can then be incorporated into the transitions representing degradation in the PN.

For clarity, throughout the remainder of this thesis the following definitions will be used:

- A failure type refers to the category of degradation – e.g. a squat, or lipping
- A fault refers to a specific stage of a failure type – e.g. a squat of length 20-35mm
- A unit is an individual section of an asset type – e.g. a 200m length of ballast, or a single crossing panel

A unit will experience multiple failure types, all of which will have a singular fault or multiple faults.

Other researchers have used the historic failure data to investigate the maintenance process, by fitting probability distributions to describe the likelihood of maintenance being effective (Audley, 2014), or model the timescale for interventions (Bryant & Andrews, 2015). This information was not available for every failure type given that many did not occur during the observation period, hence an assumption was made that maintenance is always 100% effective and interventions are left for the maximum permitted time as detailed in the railway standards.

Additionally, it is known that characteristics of the unit– such as its material, age, traffic, speed rating, curvature, weather conditions, and localised faults such as wetbeds – can impact the rate of degradation, which will in turn affect the lifetime prediction for a specific unit. The reliability analysis should take place on groups of assets that share identical characteristics, as they are expected to degrade similarly. However, a known limitation of a reliability analysis is that small datasets can introduce bias, and splitting the available dataset by, for example, curvature, will risk reducing an already small sample. The following chapter will group all similar asset types (*crossing panel, switch panel, plain line rail*, and *ballast*) together during the analysis, but in Chapter 6, an artificial neural network is introduced to investigate any evidence of changing degradation rates within an asset subgroup.

4.2 Degradation of the studied assets

Before the historic failure and condition data can be processed, it is important to first understand how the assets degrade. To detect the presence of any faults, frequent inspections are completed to assess the condition of an asset. Table 4.1 gives a brief description of the various inspection procedures undertaken for each considered track component.

Inspection type	Assets assessed	Description	Frequency
Visual	Crossing panel, switch panel	A track engineer uses a series of gauges and rulers to assess the depth of any wear and defects on the asset.	26-weekly
Ultrasonic testing unit (UTU)	Plain line rail	The UTU sends ultrasound waves through the rail, to detect any defects and build a picture of the condition of the rail. Any found defects are followed by a visual inspection to confirm the extent of degradation.	16-weekly
Track recording vehicle (TRV)	Ballast	The TRV is run over the railway line, and has a series of cameras, lasers, gyroscopes and accelerometers to record the state of the track geometry.	6-weekly

Table 4.1: Various inspection procedures and frequency

Following an inspection, the condition of an asset is known. There are a variety of failure types that can impact the railway's condition; this research takes the novel approach of considering every fault separately, instead of grouping or using a track record index. The different types of failure detected by an inspection is described in Table 4.2, which also lists which assets can be affected. Further descriptions of each failure type, including why they develop, can be found in (Cannon et al., 2003; Kumar, 2006; Lewis, 2011; Li, 2009; Magel & Kalousek, 2017).

Failure type	**Assets affected**	**Description**
Wear	Plain line rail, crossing panel, switch panel, ballast	Wear refers to the gradual loss of material from the rail or aggregate due to friction.
Lipping	Plain line rail, crossing panel, switch panel	Lipping occurs due to plastic deformation, when the rail material deforms and separates from the rail head, creating a 'lip'
Cracking	Plain line rail, crossing panel, switch panel	A crack is a vertical or horizontal fracture in the rail.
Squatting	Plain line rail, crossing panel, switch panel	A squat is a small depression of the rail surface caused by plastic deformation
Wheelburns	Plain line rail, crossing panel, switch panel	Wheelburns occur due to high levels of friction between the wheels and rail that generates heat, causing the rail surface to melt.
Shelling	Plain line rail, crossing panel, switch panel	Shelling is when a small section of the rail chips away, caused by cracking.
Crushing	Plain line rail, crossing panel, switch panel	Crushing is long areas of plastic deformation of the rail surface
Pitting	Plain line rail, crossing panel, switch panel	A pit occurs due to corrosion, causing small holes to exist on the surface on the rail.
Rolling contact fatigue (RCF)	Plain line rail, crossing panel, switch panel	RCF refers to a series of small fatigue-based cracks that occur within or on the surface of the rails. If left unchecked, RCF can develop into previously defined fault such as shelling or cracking.
Tache Ovale	Plain line rail, crossing panel, switch panel	Tache ovales are small defects, such as cracks, that are created the rail material during manufacturing
Gauge	Ballast	The gauge is the distance between the inside of the left and right rail. The HS1 line, like most railway lines across the world, use the standard track gauge of 1435mm.

Cant	Ballast	The cant is defined as being the height difference between the left and right rail. Some cant is needed, for example when travelling over a curve in the rail, and this is known as the ‘equilibrium cant’. However, if the ballast is allowed to degrade and cause a ‘cant deficiency’, the height of the rails will impact the trains centre of gravity and potentially cause a derailment.
Vertical alignment	Ballast	The vertical alignment measures the vertical rail profile e.g. if the rail dips or humps when observing from the side. The vertical alignment data is filtered over either a 35m or 70m wavelength.
Horizontal alignment	Ballast	The horizontal alignment measures the horizontal rail profile, e.g. if it buckles to the left or right when viewing the rails from above. Its measurements are also filtered over a 35m or 70m wavelength, to detect any mid-to-long-length defects
Twist	Ballast	The twist is defined as being the rate of change in the cant of the rails, measured over a 3-meter distance.

Table 4.2: Different failure types affecting the railway

The railway standards (Network Rail (High Speed) Ltd, 2020a; 2020b; 2020c; 2020d) have been used to further specify and define each failure type; many of those discussed in Table 4.2 can be considered separately depending on the location. For example, cracking on the switch panel is separated into ‘vertical cracking on the switch rail’ and ‘transverse cracking in RCF’. The literature review has highlighted the importance of recognising these separate failures, however upon consultation with railway engineers, the four failure types that describe RCF on the full section of plain line rail (each distinguishing between the severity of RCF and geographical location of the rail) have been grouped into a singular failure to be more reflective of management procedures. The standards also define the necessary maintenance action per fault as well as the maximum allotted time to complete maintenance, via a minimum action code (MAC).

Table 4.2 details five geometry parameters to describe the degradation of the ballast. The standard deviation of the vertical alignment over a 35m wavelength for a 200m section is considered as an overall indicator of the ballast condition and is used to schedule high-output tamping over the 200m unit. However, poor isolated measurements of all five parameters will also impact the unit’s condition and can trigger other maintenance, such as localised tamping or speed restrictions. It is assumed that any degradation of the ballast's aggregate will be reflected in every geometry parameter as they are not independent from one another, hence only the vertical alignment over a 35m wavelength was considered for this analysis. This data is divided into the standard deviation measurements and isolated measurements, as both will dictate the type of maintenance required, potentially impacting the lifetime predictions for the ballast.

4.3 Data processing

Studying the times-to-failure obtained from the failure and condition history can reveal the likelihood of future occurrences for the faults described in Table 4.2. Prior to any analysis, the relevant information from provided databases must be extracted and cleaned to create a comprehensive timeline of all degradation, maintenance and renewal events across a unit's life. There are two types of databases used in this research: fault data (*rail defect history*), which is only collected after a fault has been detected during an inspection, and condition data (*track geometry history, rail profile history*), which is collected via regular inspections of the track and records the condition of the asset even if a fault was not present. This section will describe how each database was processed.

4.3.1 Rail defect history

When either a UTU vehicle or a track engineer detects a fault on a rail-based track component, it is recorded in the rail defect history in a Microsoft Excel spreadsheet. This dataset was available from 2007-2019, spanning an observation period of 4223 days. Information was extracted regarding the specific unit that the fault occurred on, the definition of the failure type which was cross referenced with those described in the standards, the date that maintenance was completed, and any accompanying engineers' comments. The comments were often needed to decipher what severity of fault had been reached within the failure type. If a fault was recorded as occurring on a unit within 10 days of that same fault having occurred on the same unit, an assumption is made that this is the original fault that has accidentally been reported twice. Repeat observations are then removed from the dataset, to create a timeline of degradation per failure type for each unit.

Finding the times-to-failure per fault is best described through an example. On the crossing panel, the squat failure type is categorised into three faults: less than 20mm long, between 20-35mm long, and over 35mm long. Figure 4.1 shows the typical degradation process for squatting on a single crossing panel unit.

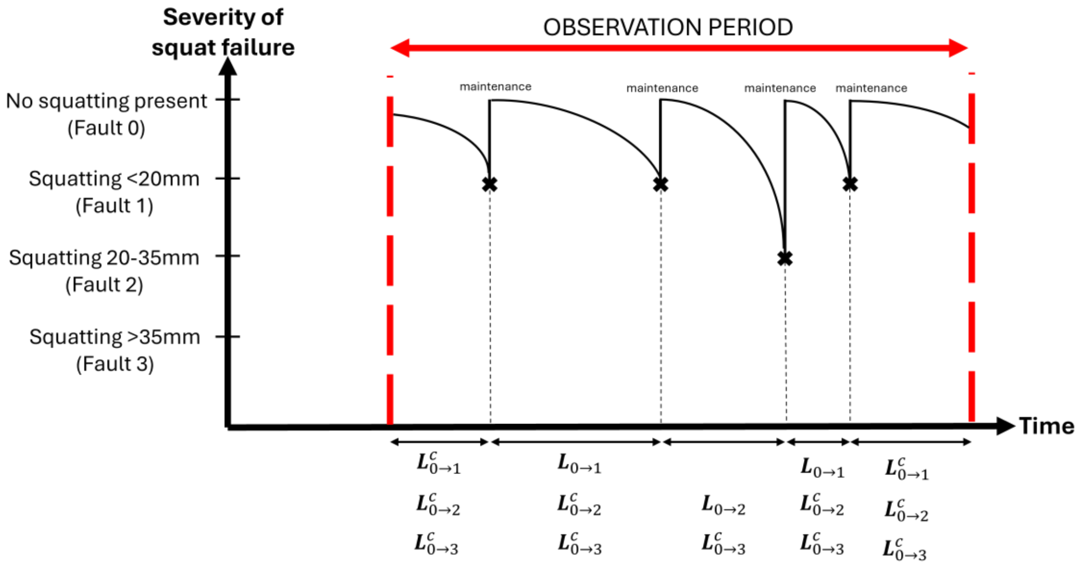

Figure 4.1: Typical behaviour of squatting on a crossing panel unit

Where $L_{0\to2}^{c}$ would be a censored time-to-failure for the crossing panel to degrade to a squat length between 20-35mm, and $L_{0\to1}$ is an uncensored time-to-failure for a squatting fault less than 20mm long to occur.

Note that maintenance is always 100% successful at removing the fault, and the initial times-to-failure found for each fault are all censored as the crossing panel was installed before the observation period had begun. The most severe squat fault – when its length is over 35mm – was not observed, which is common when inspections are frequent as the failure is noticed before it can develop further. The process described in Figure 4.1 was repeated for every failure type occurring on every unit of the studied track assets, to create large datasets of times. An example is given in Table 4.3, which contains all the $L_{0\to1}$ and $L_{0\to1}^{c}$ values for the first fault of squatting (<20mm) on a crossing panel.

Unit no.	Times-to-failure for the first fault of squatting (years)	Censoring
1	3.45	Censored
1	2.67	Uncensored
1	5.21	Uncensored
1	1.91	Censored
2	7.58	Censored
2	4.11	Uncensored
⋮	⋮	⋮

Table 4.3: Example of a dataset containing times-to-failure for a fault

The railway standards and rail defect history include defective weld repair faults, which require lifetimes to be derived from a previous weld action, rather than installation/maintenance of the unit or the beginning of the observation period. This analysis was completed for all weld faults of the switch and crossing panels, and for the MMA welds of the plain line rail, where it is assumed that the oldest recorded weld repair on the unit was defective. For faults on the plain line rail related to defective aluminothermic welds or simply classified as 'weld' without specification of what kind, the analysis was unable to distinguish whether these faults have occurred from a weld repair or those made to join the CWR, hence the lifetimes were derived from installation of the unit.

The collection of the rail defect history has several limitations, namely:

- The observation period starts in 2007, but some of the line began operating in November 2003. This means the early behaviour of the assets cannot be investigated, and any failure recorded when the initial condition was unknown must be considered right censored
- Given that the severity of the faults is typically determined by deciphering engineers' comments, there is a potential for some to be mislabelled
- Acknowledgement of a fault relies on engineers correctly filling out the rail defect history. Some faults may have been missed due to this
- The rail defect history contained faults such as ballast pitting that were detected on the switch or crossing panels of the HS1 despite not being explicitly described in the respective standards. It was found that removing said faults from the

structure of the Petri net models impacted the life expectation prediction, so they were considered during this reliability analysis

4.3.2 Track geometry history

To model the degradation of the ballast, track geometry measurements taken at 0.2m intervals, from TRV inspections completed between installation to March 2022, have been provided. In total there were 186 inspections completed during the observation period, with the results from each inspection stored in separate Microsoft Excel files. An algorithm was written to extract and store the 35m wavelength vertical alignment measurements pertaining to each 200m ballast unit, for the left and right rail, for each inspection. This was used to calculate the standard deviation (SD) of the measurements across the unit, then find the average SD across the two rails. Additional documents were provided containing the times and location for both high-output and local tamping interventions. The SD of the vertical alignment and high-output tamping history were combined to create a timeline for the behaviour of a ballast unit.

Plotting this timeline revealed events when the quality of the ballast increased despite no recorded tamping activities scheduled (as shown in Figure 4.2a, where 'RTI' indicates recorded tamping interventions). Examining the maintenance history confirmed that local tamping only coincided with 10.5% of these events, hence its impact was ignored, and it was instead assumed that these are a result of unrecorded high-output tamping interventions. Similar issues were prevalent in Khajehei et al.'s research (2019), who stated that a tamping intervention can be recognised through the track geometry if the difference in the SD of the vertical alignment after the assumed event decreases by at least 15% from the SD before the event (Equation 4.1):

$$\frac{SD_{after}}{SD_{before}} < 0.85 \tag{4.1}$$

An identical approach has been adopted here, with artificial tamping events created to explain the ballast's behaviour (Figure 4.2b, where 'ATI' indicates artificial tamping events). Once all high-output tamping events were successfully identified, the SD values were split into segments between interventions, as shown in Figure 4.3. Each segment was given a unique identification number: segment SD_375_01 would refer to the SD measurements of the vertical alignment from the ballast unit no.375 between 1-2 tamps, for example.

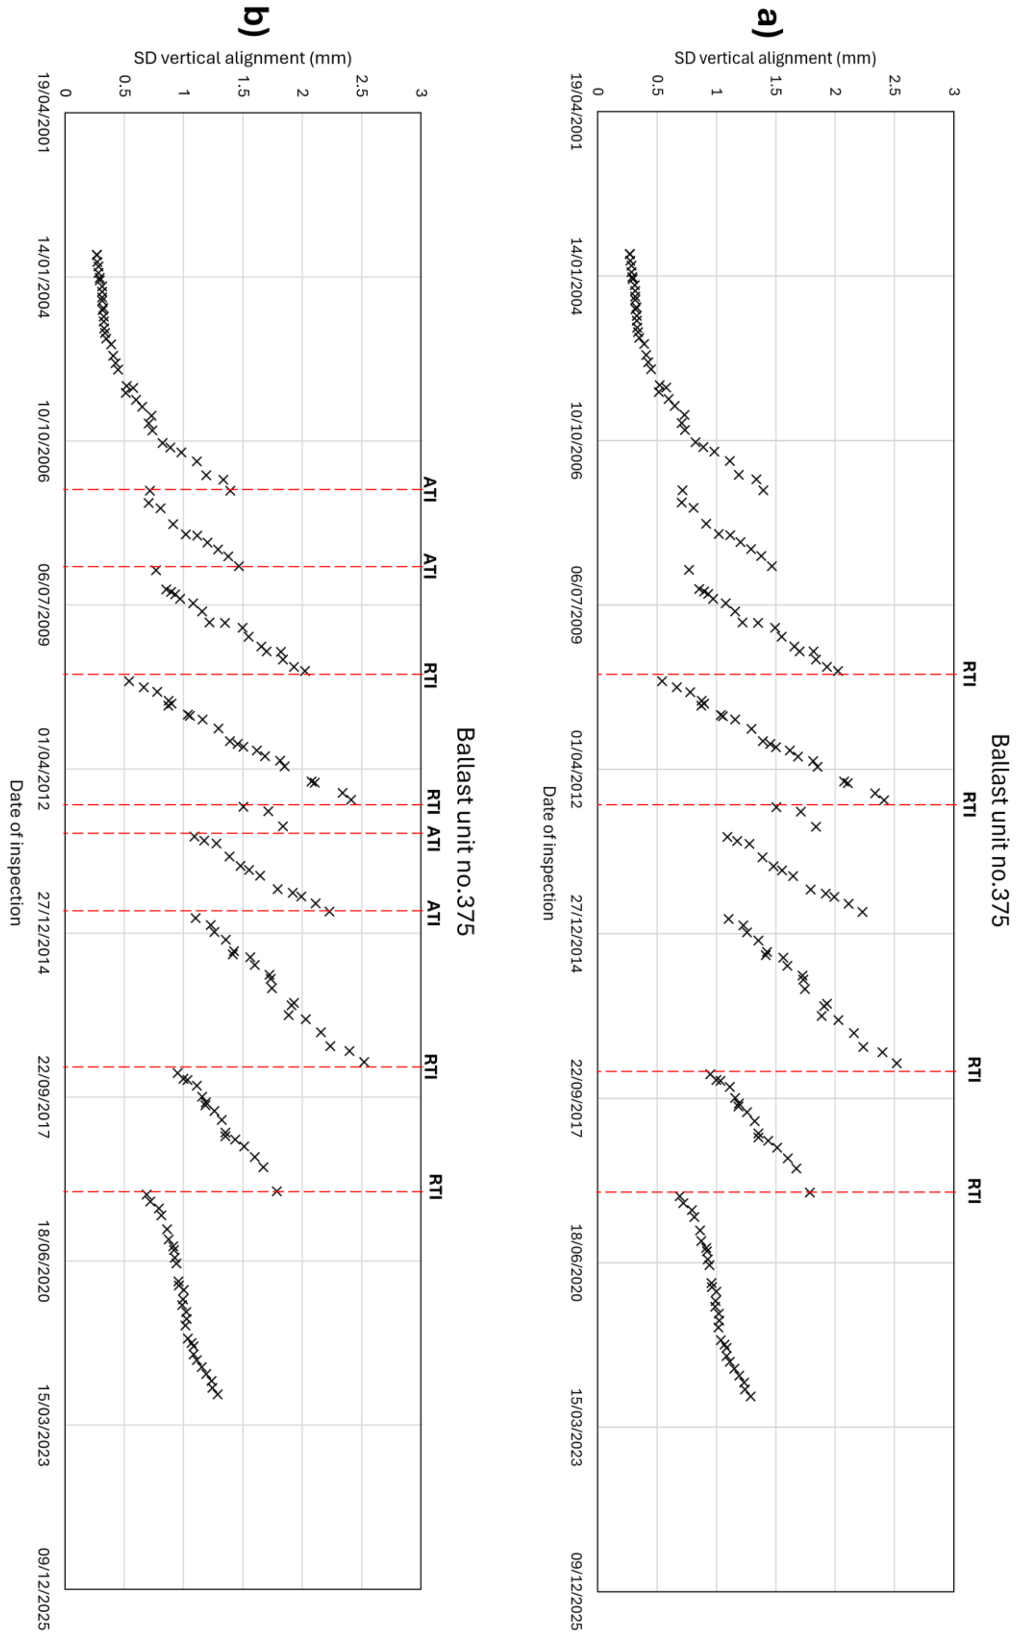

Figure 4.2: a) recorded and b) artificial and recorded high-output tamping interventions on a ballast unit

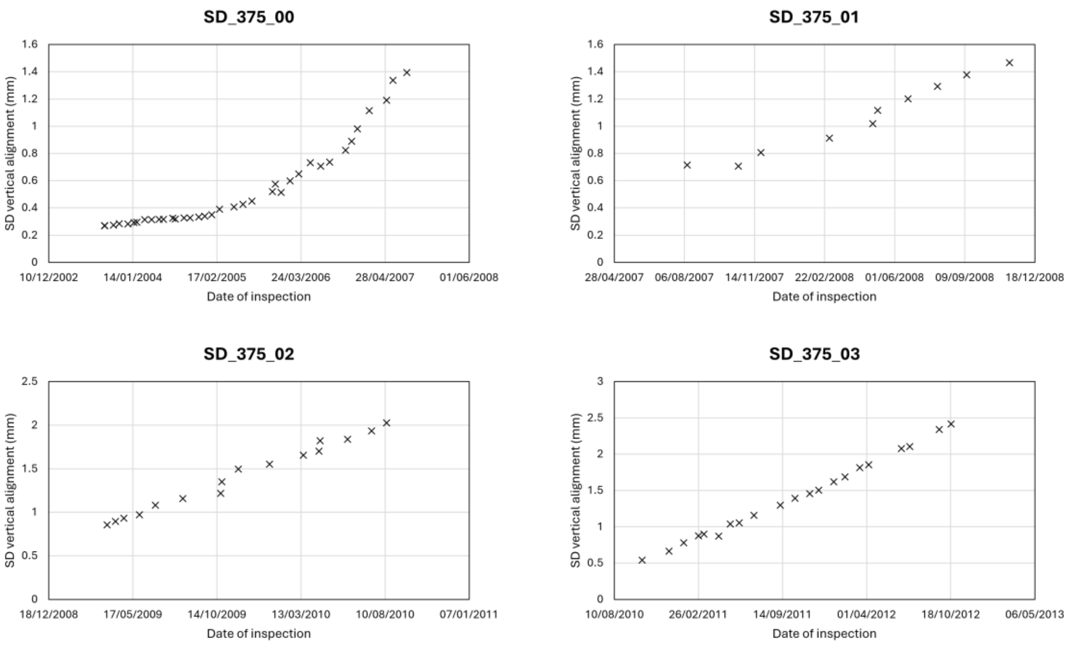

Figure 4.3: Segmentation between high-output tamping interventions

Due to the use of dynamic track stabilisation after tamping, which vibrates the ballast laterally whilst placing pressure from above to stabilise the aggregate, no evidence was found for a 'wear-in' period where the ballast quality is initially low but increases as the ballast settles. However, a few segments did experience several 'jumps', where the SD of the vertical alignment increased by over 15% between two inspections, but with no visible change to the degradation rate. These events were identified by Equation 4.2 and any segments containing a random jump event were removed from the study, as it was assumed to result from a reconfiguration of the TRV equipment.

$$\frac{SD_{after}}{SD_{before}} > 1.15 \tag{4.2}$$

The SD values within each segment can be used to find the relationship between the independent variable (time passed) and the dependent variable (condition of ballast), by fitting a regression line to the data. An assumption is made throughout that there must be 3 or more measurements within a segment to be able to fit an appropriate regression line. A visual inspection of the SD of the vertical alignment has largely indicated a linear relationship between the independent and dependent variables, which is based on the equation of a straight line:

$$y = A + Bx \tag{4.3}$$

Where y and x are the dependent and independent variables respectively, A is an intercept (the initial condition of the asset) and B is the gradient of the line.

However, previous research has found that as an asset ages, the degradation rate begins to increase (He et al., 2015), and the exponential function (Equation 4.4) is best used to describe the behaviour of the ballast:

$$y = De^{Fx} \tag{4.4}$$

Equation 4.4. can be modified using series expansion to be in a more approachable format:

$$y = A + Bx + Cx^2 + \cdots \tag{4.5}$$

The least squares formula using Cramer's rule (Robinson, 1970) was used to find the parameters in Equations 4.3 and 4.5, fitting the regression lines to the segments described in Figure 4.3. Comparing the average fitted parameters revealed there was no evidence to justify the use of an exponential fit to describe the deterioration of the SD of the vertical alignment (Table 4.4). The difference between the mean value for the A and B parameters was negligible, and the mean value of the C parameter was very close to zero, as there was low indication of curvature in the segments. Additionally, this value was negative which implies that given enough time, the ballast would stop deteriorating and potentially even improve without maintenance required.

Regression line parameters	**Mean [SD] parameter value**		
	$\boldsymbol{A}$ **(*mm*)**	$\boldsymbol{B}$ **(*mm/day*)**	$\boldsymbol{C}$ **($\boldsymbol{mm/day^2}$)**
Linear	0.668 [0.331]	0.000351 [0.000444]	N/A
Exponential	0.663 [0.383]	0.000421 [0.000969]	-0.000000317 [0.00000336]

Table 4.4: Parameters for the linear and exponential regression lines fitted to the track geometry history

The linear regression line was therefore used throughout the remaining analysis. Equation 4.3. was inverted to derive times-to-failure at a given SD value for each segment, which are placed into a dataset. The thresholds for the condition of the ballast depends on the maximum permittable speed of the section, defined further in the relevant railway standard (Network Rail (High Speed) Ltd, 2020c), but the condition itself is categorised into three states – *good, satisfactory*, and *poor*. As regression lines will extrapolate until the failure state is reached, there is no censored data in this dataset.

For a small number of segments, typically those containing less than six SD measurements, the gradient (B parameter) of the fitted line was negative implying that the condition of the ballast is expected to improve over time. Such lines are unable to derive a time-to-failure, and as such they were discarded from the research. Furthermore, the SD values in some segments were too scattered to obtain an adequate regression line to explain its behaviour- indicated by an R^2 value of less than 0.3 – and were removed before any analysis could occur. Studying the fitted A parameters, which represent the initial condition of the ballast after a restorative action, revealed that 86% of the segments were rectified to a good condition after tamping (Table 4.5), supporting the assumption that maintenance is always effective. However, inverting Equation 4.3. for a segment that was not initially in a good condition will result in negative times-to-failure, hence such segments were also removed from the study. After all data cleaning, 78% of the initial segments (3195 in total) were available for analysis.

Condition after tamping	**Good**	**Satisfactory**	**Poor**
%	85.95	11.71	2.34

Table 4.5: Condition that the ballast is restored to after high-output tamping interventions

The condition of a 200m section of ballast is indicated by its vertical alignment geometry, using the standard deviation and isolated measurements. A similar approach of plotting measurements and fitting a regression line was also attempted to model the development of defects due to isolated measurements, akin to research by Soleimanmeigouni et al. (2020). The geometry data demonstrated good positional alignment between inspections, however when attempting to identify the largest isolated vertical alignment measurement in a ballast section, the same location was rarely identified across subsequent inspections. It is thought that due to the new quality of the ballast, and because the vertical alignment is kept in a very good condition that is far stricter than the safe working limits, this very localised trend could not be seen. Instead, the algorithm used to extract the vertical alignment per 200m unit was adapted so that when a measurement was above a given threshold, a failure event was recorded. Again, the thresholds are dependent on the maximum permitted speed over the unit, but the condition of isolated measurements is defined by seven states – *good, satisfactory, poor, critical 1, critical 2, critical 3,* and *critical 4*. This process is shown in more detail in

Figure 4.4, using the vertical alignment from an inspection of a ballast unit in 2015 to identify 5 defects. It is assumed that multiple recorded faults within a 3m proximity describe the same fault, and only the most severe condition reached for a fault is recorded – for example, the first fault shown in Figure 4.4 would be recorded as *poor,* with the rest all considered *satisfactory*. The datasets of lifetimes to reach each state was then found using the methods described in 4.3.1, with the times-to-failure derived from the most recently recorded high-output tamping event.

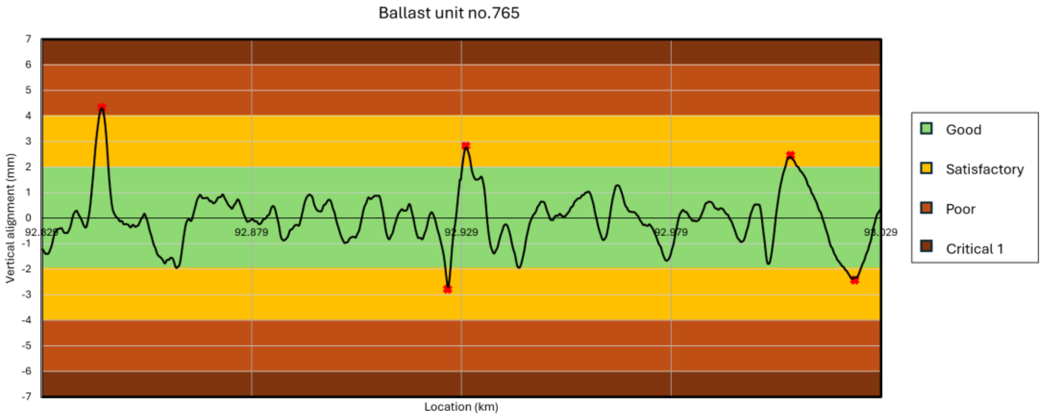

Figure 4.4: Detecting isolated faults for the vertical alignment measurements within a unit of ballast

4.3.3 Rail profile history

When the plain line rail is inspected using the UTU, the measurement equipment onboard uses a series of lasers pointed at the rail, crossing at the centre, to identify where any damage is located. This can subsequently reveal the depth of the rail from either side, vertically, and in the foot, to indicate how much it has worn since installation. A report regarding the rail profile taken at approximately 2m increments is created in Excel after every inspection, but correspondence with track engineers revealed that these measurements are largely unreliable and recommended only the vertical rail depth was used. Reports from 21 UTU inspections, spanning from 2012-2020, were provided. An algorithm was developed to extract the vertical rail depth data for the left and right rail over a 109m unit of plain line rail per inspection. As no renewals have taken place on the plain line rails, these measurements did not require segmentation due to previous maintenance actions.

The UTU history was found to be in a poor condition, with the positional alignment between inspections often being over 10m out, the left and right rail showing contrasting levels of wear that would alternate between inspections, and differences in measurements depending on which of the three UTU vehicles in the fleet was used. Considerable effort has been made to clean this database, by averaging the wear value over a 109m unit to mitigate the positional alignment issues, then averaging this value between the left and right rail. For each unit, only the measurements from a single UTU vehicle (UTU1, UTU2, or UTU3) was used, dependent upon which vehicle had been involved in the most inspections of the unit.

Similarly to the ballast, fitting Equation 4.3. and 4.5. to the data revealed that the wear of the rail was largely linear, as the mean value of the C parameter showed very little curvature (Table 4.6). As no information was available regarding the initial dimensions of the rail profile, it is assumed that the rail was in a perfect condition when installed hence the A parameter is zero.

Regression line parameters	**Mean [SD] parameter value**		
	$\boldsymbol{A}$ **($\boldsymbol{mm}$)**	$\boldsymbol{B}$ **($\boldsymbol{mm/day}$)**	$\boldsymbol{C}$ **($\boldsymbol{mm/day^2}$)**
Linear	0*	0.000669 [0.000381]	N/A
Exponential	0*	0.000326 [0.000421]	0.0000000720 [0.000000969]

Table 4.6: Parameters for the linear and exponential regression lines fitted to the rail profile history
**assumed parameter*

With the linear regression line chosen as the most appropriate, the available dataset was cleaned by removing all lines with a negative B parameter and those with an R^2 value of less than 0.3, leaving 1365 regression lines available for analysis (79% of the original dataset). A closer look at the discarded lines revealed that when the B parameter was negative, it was also incredibly small, implying that the unit was most likely not degrading over the observed period- it is impossible for the rail head to gain material over time without a renewal. Equation 4.3, with the fitted parameters, was inverted to find the time until the vertical rail depth had reduced by 14mm, consistent with the lowest safe working limits described in the standards (Network Rail (High Speed) Ltd, 2020d). The times-to-failure were then placed into a dataset for further analysis.

4.3.4 No-occurrence faults

There are many faults listed in the standards that have not been observed yet for an asset on the HS1 line. This could be because the railway line is not yet old enough to experience wear-based failure, the fault itself is incredibly rare, or because the current management strategy is very effective and doesn't allow the asset to reach critical states. Data from other high-speed railways (such as France's SCNF) was not available, and although some data was initially obtained from conventional-speed lines at the Kings Cross St Pancras station in London, it cannot be applied as the dataset does not specify when the asset was last replaced, nor the specific type and severity of fault that occurred.

Current no-occurrence faults are unlikely to contribute largely to the need to renew an asset. However, the developed asset management model will consider the impact of every potential fault, so for no-occurrence faults an assumption is made that one failure occurred on one unit of an asset, one day after the observation period ended. This creates a dataset consisting of a single uncensored, and multiple censored, lifetimes, each exactly the length of the observation period.

4.4 Fitting distributions

The datasets of times-to-failure for each fault are used to fit a continuous probability distribution that best describes the fault's occurrence over a unit's operational period. Four of the most common distributions in reliability engineering were considered:

- One-parameter exponential $\sim E(\lambda)$
- Normal $\sim N(\mu, \sigma^2)$
- Two-parameter lognormal $\sim L(\mu, \sigma^2)$
- Two-parameter Weibull $\sim W(\beta, \eta)$

A description of each distribution is given in Chapter 3.

4.4.1 Mean-time-to-failure

A simplistic way of analysing the fault databases is to find the mean-time-to-failure (MTTF), derived from Equation 4.6.

$$MTTF = \frac{\sum_{i=0}^{n_{UC}} TTF_i + \sum_{j=0}^{n_C} CTTF_j}{n_F} \tag{4.6}$$

Where n_F is the total amount of failures during the observation period, n_{UC} and n_C are the total number of uncensored and censored lifetimes respectively, and TTF and $CTTF$ are the times-to-failure and censored-times-to-failure respectively.

Equation 4.6. can also be used to find the expected rate of degradation, λ:

$$\lambda = \frac{1}{MTTF} \tag{4.7}$$

The value of λ can be implemented into a 1-parameter exponential distribution or alternatively, a property of the Weibull distribution means that it can mimic the behaviour of the exponential distribution if the beta parameter is set to 1 (Nielsen, 2011). The MTTF can then be used as the scale parameter, to allow for the Weibull distribution to model the databases. Either fitted distribution will imply that the fault occurs randomly at a constant rate.

4.4.2 Parameter estimation method

A parameter estimation method and goodness-of-fit test can be used in tandem to evaluate the best fitting distribution for the fault datasets, as an alternative to using the MTTF. The Reliasoft Weibull++ software (ReliaSoft, 2020) has been used for this capacity throughout; for the analysis to work, there must be a minimum of two uncensored failure times present in the database. The parameters estimation method is applied to the fault databases created in section 4.3, which contain mostly censored, and very few uncensored, lifetimes. There is no set rule for deciding the best method, but typically small datasets are best explained using rank regression (RRX), and datasets containing high censoring is best represented by maximum likelihood estimation (MLE). To investigate which parameter estimation method would be the most appropriate, Weibull++'s SimuMatic tool was used. This tool works by generating a sample of specified size and censoring threshold from a chosen distribution using Monte Carlo simulation. The parameters of the sample are then approximated using the chosen estimation method. This process is repeated hundreds of times, and the probability plots from each sample are combined, creating a visual aid which can be used to analyse and evaluate the effectiveness of different parameter estimation methods.

For this investigation the sample size was set to 143 and any lifetime generated larger than 12 years was considered right censored. For the considered distributions, parameters were chosen that are consistent with characteristics expected from the fault history – a range of degradation rates (increasing, decreasing, or constant) with an average lifespan large enough so that between 60%-95% of the lifetimes are censored. 500 samples were created per distribution type and parameter estimation method, and the resulting probability plots are shown in Figures 4.5-4.7:

I. $\sim Exponential(0.017)$. Constant degradation rate with roughly 82% of the database censored:

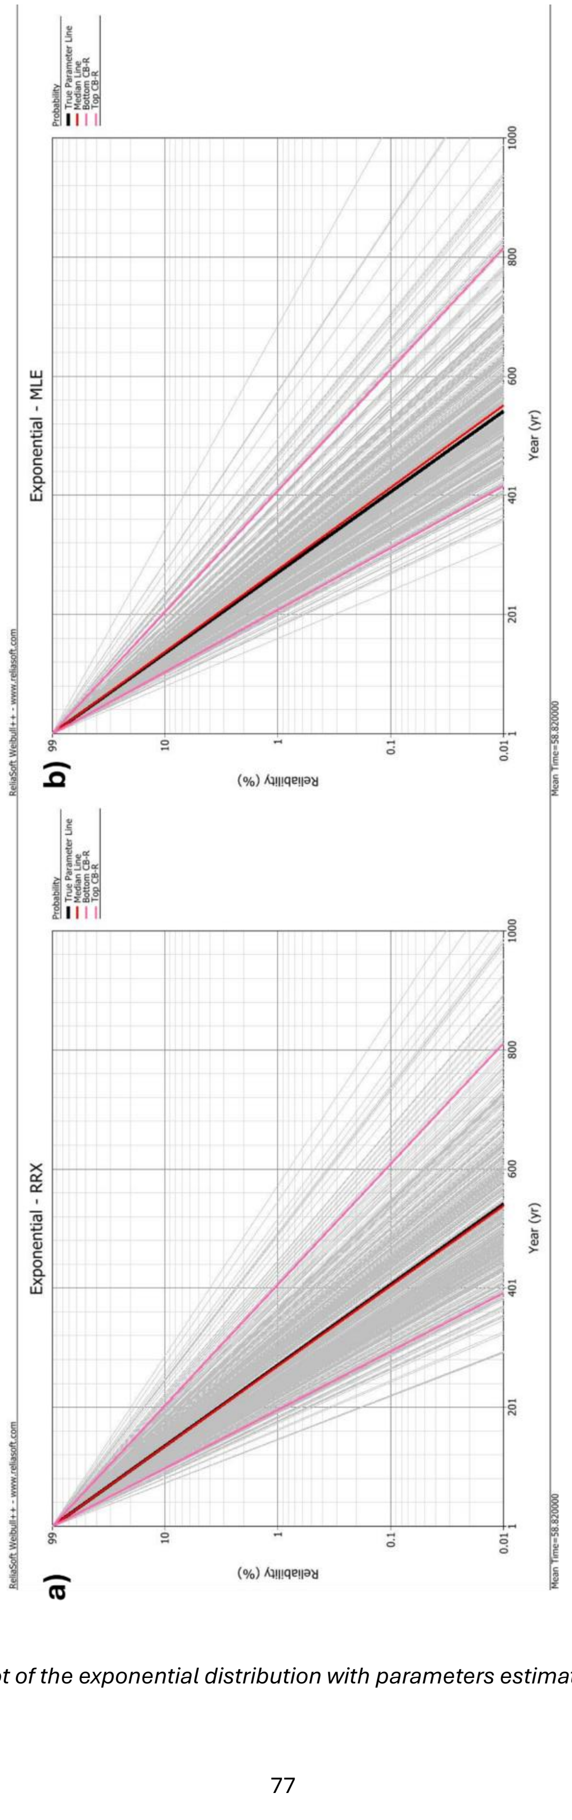

Figure 4.5: SimuMatic plot of the exponential distribution with parameters estimated by a) RRX and b) MLE

II. $\sim Lognormal(3,2)$. Decreasing degradation rate with roughly 60% of the database censored:

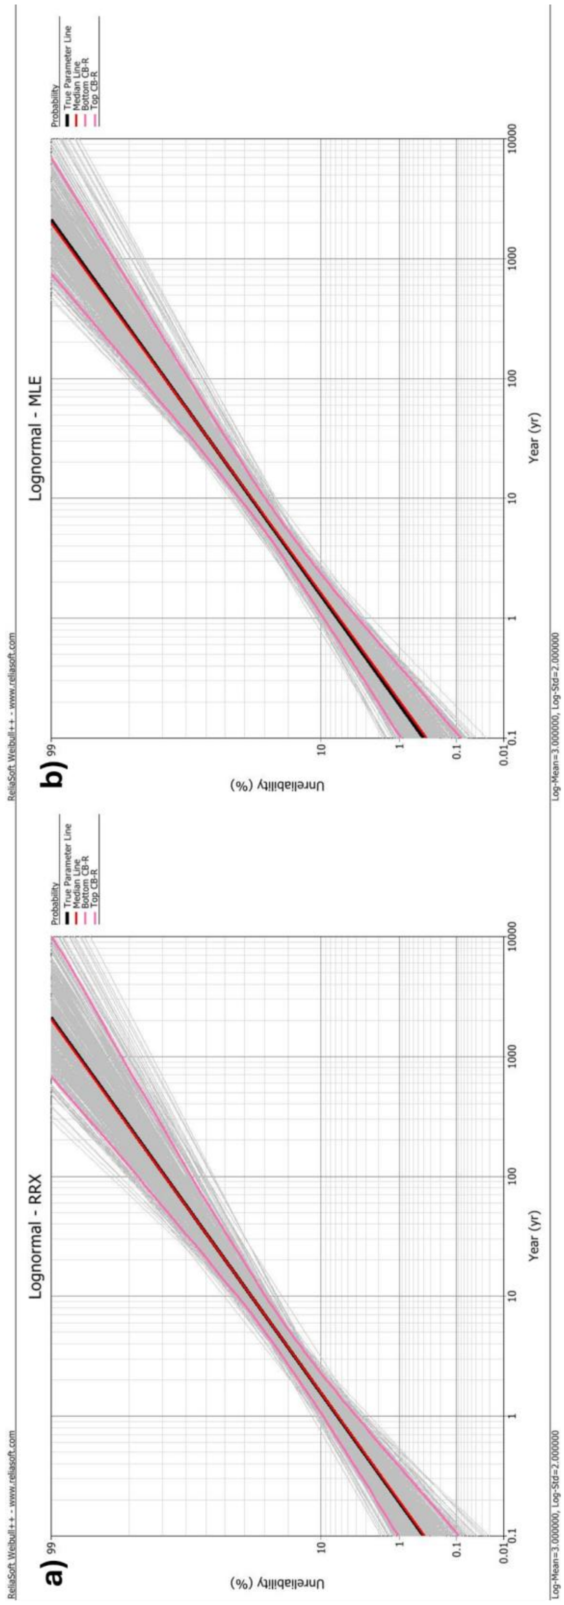

Figure 4.6: SimuMatic plot of the lognormal distribution with parameters estimated by a) RRX and b) MLE

III. $\sim Weibull(1.5, 60)$. Increasing degradation rate with roughly 91% of the database censored:

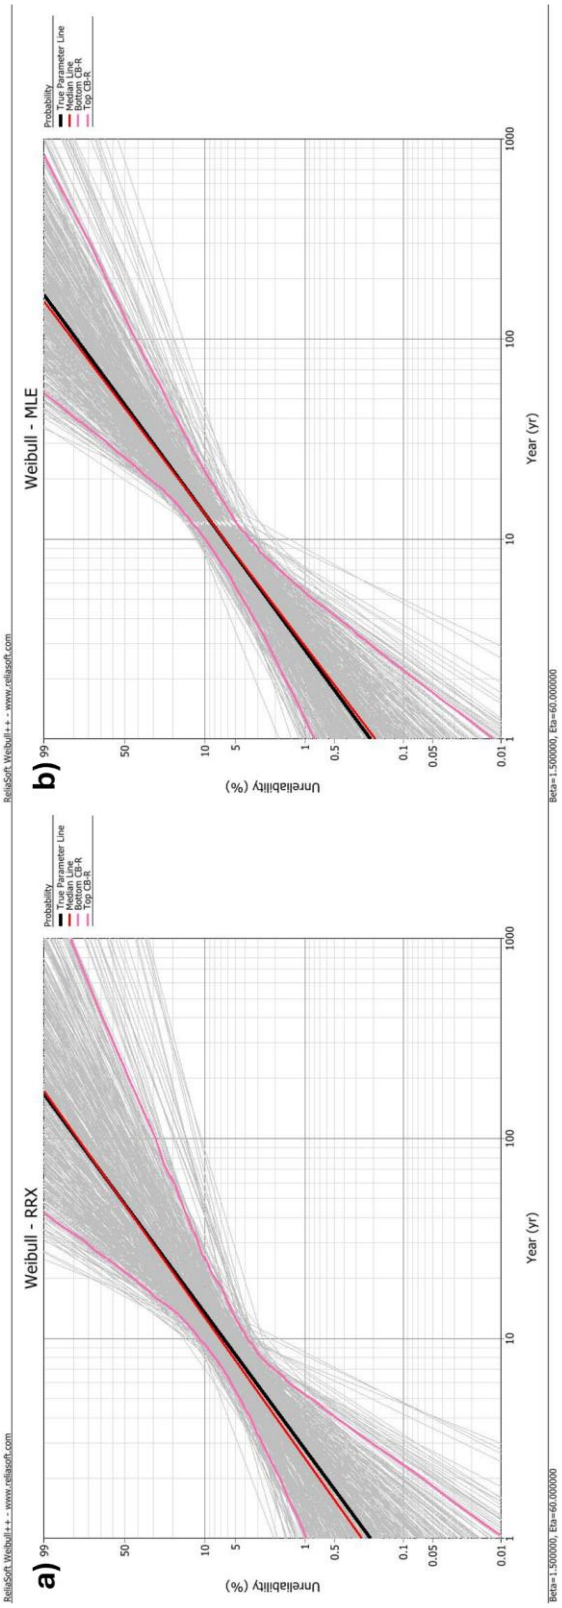

Figure 4.7: SimuMatic plot of the Weibull distribution with parameters estimated by a) RRX and b) MLE

The grey lines represent the probability plot of each individual sample. Figures 4.5 – 4.7 all show that the median parameter line (red) matches closely with the true parameter line (black), suggesting that either estimation method has little deviation from the true parameter values. For datasets with heavier censoring, the median parameter line for the MLE has a slightly better fit to the true parameter line at the extreme values, exemplified in Figure 4.7. There is a noticeable difference in the placement of the upper and lower confidence bounds (pink lines) of the reliability, with the MLE having a smaller 95% confidence interval and therefore less variation across the estimated results. The results of the investigation prove that the MLE remains the best parameter estimation method for datasets with low amounts of uncensored lifetimes, which are typically obtained when observing failures on the HS1 track, and will be used for the remainder of this research.

4.4.3 Goodness-of-fit tests

Weibull++ considers how well the distribution and estimated parameters fits the dataset using 5 goodness-of-fit tests:

- Kolmogorov-Smirnov test (GOF)
- Correlation coefficient test (PLOT)
- Likelihood value test (LKV)
- Bayesian Information Criterion (BIC)
- Akaike Information Criterion (AIC)

The result of each test for every considered distribution is ranked between 1 and 4, with 1 indicating the distribution is the best fit and 4 the worst. The overall ranking of a distribution (OR) is found according to Equation 4.8.:

$$OR = (GOF_r * w_{GOF}) + (PLOT_r * w_{PLOT}) + (LKV_r * w_{LKV}) + (BIC_r * w_{BIC}) + (AIC_r * w_{AIC}) \tag{4.8}$$

Where GOF_r, $PLOT_r$, LKV_r, BIC_r, AIC_r are the results from the ranking, and w_{GOF}, w_{PLOT}, w_{LKV}, w_{BIC}, w_{AIC} are the weights, of the associated goodness-of-fit test.

The lowest OR value represents the most suitable distribution to fit the dataset, but this is dependent on the perceived importance of each goodness-of-fit test. Due to the strengths of the AIC and BIC highlighted in section 3.4.2, these were weighted at 50% each and the other tests are discarded, consistent with the recommended weightings given by the software. The property of the AIC and BIC that states that if the score for a distribution is within two units of the score for another distribution both are considered as effective at fitting the available dataset, can be used to evaluate the importance of the four considered continuous probability distributions. For example, studying two failures of the crossing panel, ‘flaking or shelling of the wheel contact surface’ and ‘transverse defect from RCF in a casting’, revealed that not only were the best-fitting distributions significantly a better fit than the others considered, but also that they varied between the different faults (Table 4.7). Typically in asset degradation research, a single distribution type is applied to every dataset within the study, but as Table 4.7 illustrates, a similar approach for this project would result in many faults being fitted with distributions that are an inadequate representation of their behaviour.

Distribution type	**Flaking or shelling of the wheel contact surface**		**OR**	**Transverse defect from RCF in a casting**		**OR**
	AIC	**BIC**		**AIC**	**BIC**	
2-parameter Weibull	1186.6	1190.6	200	1502.2	1506.9	200
2-parameter normal	1220.3	1224.5	400	1608.9	1613.5	400
1-parameter exponential	1184.6	1186.7	**100**	1519	1521.3	300
2-parameter lognormal	1187.8	1191.9	300	1495.1	1499.8	**100**

Table 4.7: Best fitting distribution per fault

4.4.4 Best fitting distribution per failure type

For fault databases with two or more uncensored lifetimes, the Weibull++ software utilising the MLE and AIC and BIC were used to fit an appropriate distribution, with the MTTF used to analyse all remaining faults. Following the results of Table 4.7, it was concluded that there was no significant reason for faults from separate failure types to share a single distribution, however, this was still required for faults within a failure type.

The best fitting distribution per failure type with more than one fault was found by averaging the OR value across the associated faults for all four of the considered distributions, then selecting the distribution with the lowest score. If a fault distribution was derived from utilising the MTTF, an OR value cannot be calculated. In these cases, as the mean-time-to-failure can be implemented as the lambda parameter in the exponential distribution, or the eta parameter in the Weibull distribution, the average OR score is found per applicable faults, then the lowest score is chosen between the exponential and Weibull distribution. If the failure type included only faults whose distribution was derived from the MTTF, the exponential distribution was considered as the most appropriate. An example of this process is shown in Table 4.8 for the squatting failure on the crossing panel which has three associated faults: <20mm long, between 20-35mm long, and over 35mm long. Despite the lifetime for the crossing panel to go from a perfect condition to a squat over 35mm long being best described via an exponential distribution, overall, the two-parameter lognormal was the best fit of the squat fault databases:

		OR score per distribution type				
Fault	**Distribution before**	**Exponential**	**Lognormal**	**Normal**	**Weibull**	**Distribution after**
<20mm	$\sim L(8.257, 1.835)$	300	**100**	400	200	$\sim L(8.257, 1.835)$
20-35mm	$\sim L(9.401, 2.157)$	250	**100**	400	250	$\sim L(9.401, 2.157)$
>35mm	$\sim E(0.00000497)$	**100**	200	400	300	$\sim L(13.520, 2.543)$
	Average OR	217	**133**	400	250	

Table 4.8: Overall best fitting distribution per failure type

It was found that variation in the best fitting distributions within a failure type tended to only happen when the fault databases had a low number of uncensored lifetimes, with the AIC and BIC often indicating that any of the considered distributions would be an equally appropriate fit. The analysis described in Table 4.8 was completed for each failure type with multiple faults.

4.5 Analysis of the fitted distributions

After fitting the most appropriate distribution to the databases of each fault occurring on a high-speed asset, the associated parameters can be utilised to infer characteristics of the faults behaviour. For example, the parameters can be used to calculate the mean lifespan of the asset before the fault is experienced through the following equations:

I. Exponential

$$\frac{1}{\lambda} \tag{4.9}$$

II. Lognormal

$$e^{\mu+\frac{\sigma^2}{2}} \tag{4.10}$$

III. Weibull

$$\eta\Gamma\left(1+\frac{1}{\beta}\right) \tag{4.11}$$

IV. Normal

$$\mu \tag{4.12}$$

However, the lognormal and Weibull distributions are not typically symmetrical and if there is a large skew to the data, the mean lifespan will be influenced and unrepresentative of the typical times-to-failure that were found. Therefore, for these distributions, an alternative median measurement correlating to the time t at which 50% of times sampled from the distribution will be lower than t ($P(T < t)$ = 0.5) was used.

The times-to-failure for each fault – and by association, the estimated parameters of the found distributions – were measured in days, consistent with the scaling for actions like inspections or maintenance. The mean or median lifespan was then converted to years as it was easier to conceptualise the likelihood of the fault occurring over the operational period of the unit. It is noted that initially, the average lifespan of an asset before the fault is experienced can seem exceptionally large- many faults have lifespans of over 100 years. The structure of the developed life-cycle model is such that a unit does not require every fault to occur before a renewal can be scheduled, and the expected lifespan should be interpreted as an indicator of how likely that fault is to occur during a simulation of the model – the smaller the time, the more likely the unit is to experience the fault at least once during its operational period. The large expected lifespans may also be compounded by the low number of uncensored times-to-failure in many of the fault datasets, as this is likely to distort results and introduce bias. Without a larger dataset available, the fitted distributions were recognised as the best current predictions for the fault behaviour, however the resulting expected lifespan are highlighted to recognise that this is likely an untrue representation of the actual behaviour. Such faults are currently considered irrelevant for driving renewals of the considered asset.

Fitting a continuous probability distribution to every fault impacting a high-speed railway asset is novel, and the results are discussed in Tables 4.9-4.12. Each fitted distribution has associated parameters (Equations 3.1– 3.4) – the ‘Parameter 1’ column refers to the degradation rate, mean, shape, and shape parameters of the exponential, normal, lognormal, and Weibull distribution respectively, and the ‘Parameter 2’ column refers to the standard deviation, scale, and scale parameter of the normal, lognormal, and Weibull distribution respectively. For each parameter, the 95% confidence interval (CI) is also included to evaluate how uncertain the estimation is. Additionally, for convenience the ‘No occurrence’ fault is used to refer to all faults without a recorded failure during the observation period and hence whose distribution was estimated using the assumption in section 4.3.4 and the MTTF.

4.5.1 Switch panel

The best fitting distribution for faults occurring on the switch panel is described in Table 4.9:

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
No occurrence	-	0 (1)	142	Exponential*	0.0000017	-	1654.49
Ballast pitting	-	3	149	Lognormal	26.32 [6.45, 46.18]	8.79 [3.09, 25.02]	N/A
Defective weld repair	-	7	189	Lognormal	12.99 [9.31, 16.66]	3.03 [1.61, 5.70]	1200
Wheelburn	-	6	146	Weibull	1.92 [0.87, 4.22]	21876 [5553, 86172]	49.52
Transverse cracking in RCF	-	29	173	Lognormal	11.05 [9.89, 12.20]	2.69 [1.98, 3.66]	172.45
Shallow flange contact angle	Below 60	45	143	Weibull	0.53 [0.40, 0.70]	42146 [17424, 101946]	57.83
	Below 55	1	143	Weibull*	1	603889	1654.49
Switch rail height relative to the stock rail	Projecting above the stock rail	8	147	Exponential	0.000014 [0.0000071, 0.000028]	-	192.35
	Forming a ramp	0 (1)	142	Exponential*	0.0000017	-	1654.49

(cont.)

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
Lipping in the machined length of the switch or stock rail	-	510	179	Lognormal	6.48 [6.37, 6.58]	1.29 [1.21, 1.37]	1.78
Gauge 2 contacts below any damage or crack	Less than 200mm	89	147	Exponential	0.000160 [0.00013, 0.00020]	-	17.12
	More than 200mm	17	216	Exponential	0.000028 [0.000018, 0.000046]	-	94.47
	Multiple cracks (from a crack < 200mm long)	24	208	Exponential	0.000041 [0.000027, 0.000060]	-	67.65
	Multiple cracks (from a crack > 200mm long)	1	143	Exponential*	0.0000017	-	1654.49
Gauge 2 contacts above any damage or crack	Less than 200mm	2	143	Exponential	0.00000331 [0.00000083, 0.000013]	-	827.71
	More than 200mm	0 (1)	142	Exponential*	0.0000017	-	1654.49
	Multiple cracks (from a crack < 200mm long)	0 (1)	142	Exponential*	0.0000017	-	1654.49
	Multiple cracks (from a crack > 200mm long)	0 (1)	142	Exponential*	0.0000017	-	1654.49

*Table 4.9: Fitted distributions for faults on the switch panel *estimated distribution*

Of the 13 failure types specific to a switch panel described in the railway standards, only six had occurred during the observation of the HS1 line, with the additional faults 'Wheelburn', 'Ballast pitting', and 'Defective weld repair' also included in the analysis as they appeared in the rail defect history (despite not explicitly mentioned in the switch panel standards). By using the MTTF of these no-occurrence datasets to fit an exponential distribution, the associated degradation rate reveals that there is a 0.06% chance per year of these faults occurring, and are therefore not expected to have a significant impact on the switch panel requiring a renewal due to the extremely low chance of appearing during a simulation of the Petri net. Of the 196 repair welds completed over the observation period, only 7 were defective; the likelihood of a defective weld is best described through a lognormal distribution with a large scale parameter, suggesting the probability of failure is highest almost immediately after the repair but the risk decreases the more the weld ages.

Most faults for the switch panel are fitted with an exponential distribution, as such they are expected to occur randomly across the asset's lifespan at a constant rate. The normal distribution was not chosen as an appropriate representation for any of the faults, with the remaining being distributed either lognormally or via the Weibull distribution. Studying the fitted parameters for these distributions revealed there is very little evidence of faults on the switch panel having an increasing degradation rate – the scale parameters of the lognormal distributions, and the beta value of 0.53 for a shallow flange contact angle below 60 degrees implies that the asset is most at risk of failure shortly after installation. This correlates with similar research by Litherland and Andrews (2023a), who found that the occurrence of faults on a switch rail had a decreasing degradation rate. The exception to this finding is the wheelburn fault, which is fitted with a Weibull distribution with a shape parameter of 1.92.

The 95% confidence intervals frequently display a large range relative to the parameter value. This is particularly evident for the ballast pitting fault which has an incredibly large CI for both fitted parameters, suggesting a high degree of uncertainty in the fitted distribution. Occasionally, the CI implied contrasting characteristics of the fault – for example, the shape parameter CI for the previously mentioned wheelburn fault was between 0.87 and 4.22, casting doubt over whether the fault is truly associated with an increasing degradation rate. The large CIs are thought to transpire due to the low amount of uncensored times-to-failure throughout the switch panel and any findings should be treated with caution. This uncertainty is rectified only through larger datasets of uncensored lifetimes, which can be seen in common faults such as lipping or a shallow flange contact angle below 60, where the beta parameter CI further validates the assumption of a decreasing degradation rate over time.

The expected lifespan of the asset before the fault occurs is often large enough (>100 years) to indicate that they are not consistently expected during a unit's operational period. The same cannot be said for lipping, which dominated switch panel degradation in the historical fault data and was responsible for almost 70% of all failures. Lipping is best described by a lognormal distribution with a moderately large scale parameter, resulting from a large right skew in the dataset, and is expected to occur roughly every 1.78 years on a switch panel. The railway standards divide cracking on the switch rail into two categories depending on whether the crack occurs above or below a gauge used during the inspections. For either category, the degradation process is as shown in Figure 4.8. Once a crack less than 200mm long has occurred, if left untreated it can degrade into either multiple cracks, or a longer crack which in turn can also develop into multiple cracks.

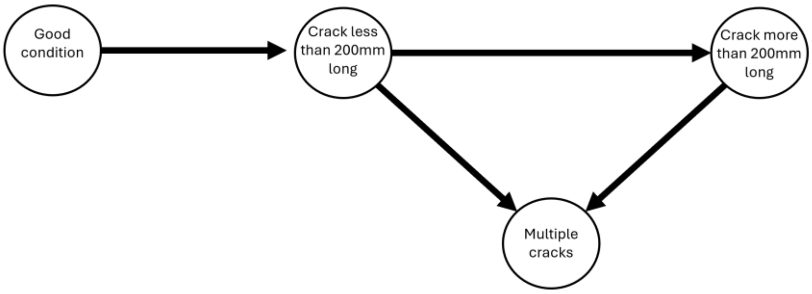

Figure 4.8: Degradation process of cracking on the switch rail

For switch blade damage that occurs above the gauge, it can be inferred from the fault history and expected lifespans that a crack <200mm long on the switch panel is more likely to progress into multiple cracks than it is to a longer, single crack. Although all three faults are rectified by grinding, the maximum time allotted to complete maintenance varies. This stresses the importance of considering the severity of faults during a reliability analysis, as results like this will impact how frequently maintenance is scheduled on the asset which may be paramount for various renewal strategies.

4.5.2 Crossing panel

The best fitting distribution for every fault that could occur on the crossing panel is described in Table 4.10.

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
No occurrence	-	0(1)	142	Exponential*	0.0000017	-	1654.49
Ballast pitting	-	4	143	Exponential	0.0000067 [0.0000025, 0.000018]	-	408.91
Defective weld repair	-	7	387	Exponential	0.000012 [0.0000056, 0.000025]	-	228.31
Wheelburn on a casting	-	5	145	Exponential	0.0000083 [0.0000034, 0.000020]	-	330.09
Wear/damage in wheel transfer area	-	27	157	Exponential	0.000042 [0.000029, 0.000062]	-	65.23
Lipping on a casting	-	115	174	Exponential	0.00019 [0.00016, 0.00022]	-	14.42
Transverse defect from RCF on a casting	-	76	157	Lognormal	9.08 [8.58, 9.57]	2.41 [2.01, 2.88]	24.05
Squat on a casting surface length	<20mm	103	165	Lognormal	8.26 [7.94, 8.57]	1.83 [1.58, 2.13]	10.59
	20-35mm	59	205	Lognormal	9.40 [8.86, 9.94]	2.16 [1.77, 2.63]	33.12
	>35mm	3	149	Lognormal	13.52 [7.82, 19.22]	2.54 [0.90, 7.20]	2038.77

(cont.)

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
Flaking or shelling of the wheel contact surface	-	53	149	Exponential	0.000088 [0.000067, 0.00011]	--	31.13
Transverse cracking of casting in the wing rail to a depth below the running surface	Up to 15mm	9	149	Exponential	0.000015 [0.0000078, 0.000029]	-	182.65
	15-25mm	1	143	Exponential*	0.0000017	-	1654.49
	25-75mm	1	143	Exponential*	0.0000017	-	1654.49
	>75mm	0(1)	142	Exponential*	0.0000017	-	1654.49
Transverse cracking of casting in the crossing vee to a depth below the running surface	Up to 15mm	18	152	Weibull	1.46 [0.94, 2.26]	17042 [8350, 34779]	36.32
	15-25mm	0(1)	142	Weibull*	1	603889	1654.49
	>25mm	0(1)	142	Weibull*	1	603889	1654.49

*Table 4.10: Fitted distributions for faults on the crossing panel *estimated distribution*

The railway standards described 22 failure types, of which only 8 were experienced between 2007-2019. Again, 'ballast pitting' and 'defective weld repair' were not directly referenced in the crossing panel standards but were included in the analysis as they are present in the rail defect history. Compared to the switch panel, the crossing panel experiences a lower frequency of faults but a wider variety, with lipping remaining the single most common fault as 115 uncensored times-to-failure could be obtained from the rail defect history. Interestingly, this fault was best described by an exponential distribution, implying no association with the life stage of the crossing panel and with a 7% chance that lipping will occur in any given year. Another very common failure type is squatting, accounting for 35% of all recorded failures on the crossing panel. The recommended maintenance action for lipping is to grind the fault away, whereas any squats should be rectified by adding material via welding. This highlights the importance of considering faults separately, as they have different impacts on the condition, and therefore the lifespan, of a crossing panel unit.

The fitted distributions indicate that most of the faults either occur randomly or have a brief settling in period where the degradation rate rapidly rises, but after this has been reached, the expected occurrence of the fault decreases as the asset ages. This may be a characteristic of crossing panels on the HS1 line, but these findings do not correlate with other research which has found the crossing nose to have an increasing degradation rate over time (Litherland & Andrews, 2023a). Given that the reliability analysis was completed before any crossing panels had been replaced due to wear, it may be that late-stage degradation has simply not been observed yet and therefore is not reflected in the fitted distributions. Additionally, the results in Table 4.10 have revealed that the lognormal distribution is more commonly used to describe the behaviour of a fault compared to the Weibull or normal, which was never chosen as the most appropriate distribution. A lognormal distribution typically predicts times-to-failure that occur early into the assets operational period with a large right skew, but research by Kittaneh et al. (2022) found that for datasets with high amounts of censoring, the lognormal distribution was better at minimising the error between the true and predicted MTTF compared to the Weibull distribution. Given the large quantities of censored data present, this may explain why the lognormal distribution – and therefore early failures – were so frequently associated with the crossing panel.

The expected lifetime of the asset before each fault is expected to occur has indicated which failure types are most at risk of developing into more severe faults. There is a difference of 1468 years between the mean lifespans for a transverse crack in the vee growing from less than 15mm long to between 15-25mm in length, suggesting that this will be an extremely rare occurrence. Conversely, there is only a 22-year difference between a crossing panel going from a perfect condition to a squat fault <20mm and between 20-35mm. Whilst still long, this is a significantly shorter length of time and suggests an effective management strategy should aim to rectify the crossing panel's condition as soon as a squat fault is detected.

4.5.3 Plain line rail

The best fitting distribution for every fault that could occur on the plain line rail is described below in Table 4.11.

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
No-occurrence	-	0(1)	1738	Exponential	0.00000014	-	20119.99
Crushing	-	9	1767	Exponential	0.0000012 [0.00000055, 0.0000022]	-	2515.00
Defective aluminothermic weld with MPT indication	-	82	1892	Exponential	0.000011 [0.0000091, 0.000014]	-	243.32
Damage on the running surface	<15mm	11	1773	Exponential	0.0000015 [0.00000083, 0.0000027]	-	1829.09
	15-25mm	0(1)	1738	Exponential*	0.00000014	-	20119.99
	>25mm	0(1)	1738	Exponential*	0.00000014	-	20119.99
Transverse defect from a tache ovale, squat, wheelburn, or weld repair	1mm-10mm depth	12	1762	Exponential	0.0000016 [0.00000093, 0.0000029]	-	1676.67
	10-15mm depth	2	1741	Exponential	0.00000027 [0.000000068, 0.0000011]	-	10060.00
	15-25mm depth	0(1)	1738	Exponential*	0.00000014	-	20119.99
	>25mm depth	0(1)	1738	Exponential*	0.00000014	-	20119.99

(cont.)

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
Transverse crack in rail foot	-	2	1739	Exponential	0.000000028 [0.000000068, 0.0000011]	-	10059.00
Untestable rail	-	11	1742	Lognormal	19.37 [13.11, 25.62]	4.42 [2.54, 7.67]	N/A
Arc weld repair with an ultrasonically detected defect	<15mm long	13	92	Exponential	0.000066 [0.000038, 0.00011]	-	41.42
	>15mm long	0(1)	104	Exponential	0.0000023	-	20119.99
Visible cracking not covered in other descriptions	-	6	1745	Lognormal	24.59 [12.17, 37.01]	6.01 [2.84, 12.71]	N/A
Vertical wear	14mm depth	1365	-	Lognormal	10.02 [9.99, 10.04]	0.54 [0.52, 0.56]	61.32
Flaking	-	2	1743	Exponential	0.00000027 [0.000000068, 0.0000011]	-	10060.00
Wheelburns	-	35	1794	Exponential	0.0000048 [0.0000035, 0.0000067]	-	568.62

*Table 4.11: Fitted distributions for faults on the plain line rail *estimated distribution*

Figure 4.9 shows the total amount of faults attributed to the switch panel, crossing panel, and plain line rail between 2008-2019. There are 12.16 times the number of plain line rail units on the HS1 line compared to the amount of S&C units, yet the total number of faults recorded is typically equal to or lower in comparison. It is thought that NRHS's cyclic grinding campaign, where the top layer of the plain line rail is routinely removed to prevent faults and increase ride quality, is incredibly effective as most units had simply not experienced a failure during the 12-year observation period.

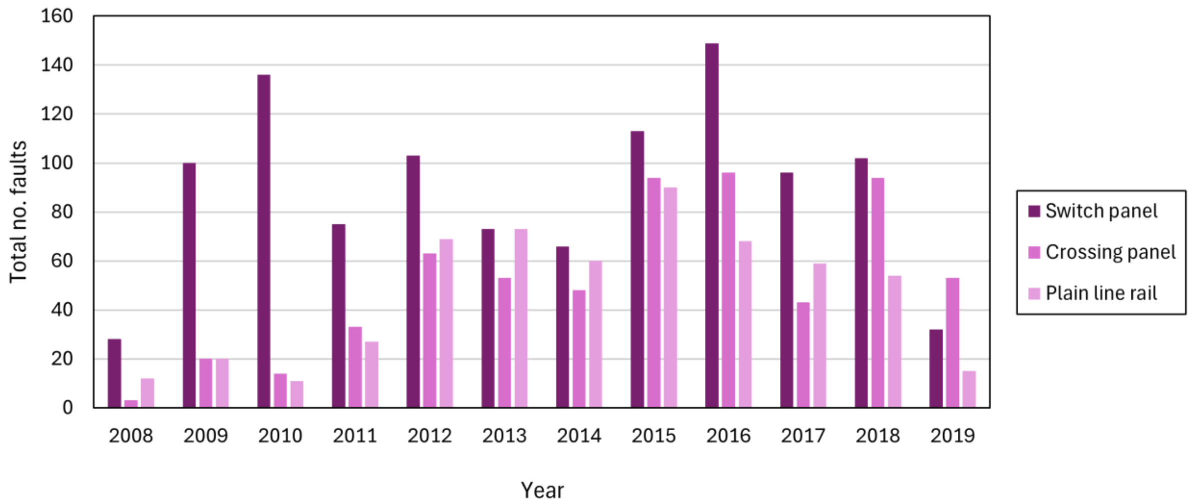

Figure 4.9: Total number of faults experienced per year by each rail-based asset on the HS1 line

It is not unexpected for the plain line rail to experience comparatively small amounts of failures – research by Kassa et al. (2006) and Cornish et al. (2016) (using databases from Great Britain and Sweden) both state that S&C assets take up the larger proportion of the fault history compared to the plain line rail. The low number of failures is further exacerbated by the observation period beginning after many of the units were installed, meaning that 65% of all found times-to-failure were censored.

Issues stemming from the low numbers of uncensored times-to-failure in the datasets were reflected in every metric described in Table 4.11:

- Most distributions were fitted with an exponential distribution as no association could be found between the lifetime of the asset and the fault occurrence
- Of the 29 failure types described in the plain line rail standards only 11 had been experienced during the observation period.
- The expected lifespan of the plain line rail before the faults were expected to occur were exceptionally large to the point where most are not expected to occur

at all during a simulation (faults that did not occur during the observation period for example have a 0.005% chance of developing at any given year).

- There was almost no evidence of a failure type ever degrading beyond the first fault.
- There is heavy uncertainty around the fitted parameters, as the 95% confidence intervals are also incredibly large.

These flaws are not observed for the vertical wear on the head of the rail to reach 14mm in depth. This was the only failure type able to be derived from the UTU history and was best described with a lognormal distribution. The low scale parameter of the fitted distribution suggests that there is less variation between the found times-to-failure, and less of a right skew seen in the dataset compared to other faults such as an untestable rail.

The expected lifespan of the asset before it has reached the maximum vertical wear depth is 61 years and if reached, a renewal is automatically scheduled as it is unsafe to continue to allow rolling stock travel over the rails. There are five other failure types in the standards that relate to wear of the rail, all of which also require renewal when detected. However, because the observed plain line rail was not yet old enough to have worn to a threshold where action is required, and the UTU measurements for these failures were deemed insufficient to obtain regression lines from, no times-to-failure could be obtained. Due to the assumptions for no-occurrence faults, they can still be fitted with an exponential distribution to model their behaviour, however it is acknowledged that this will impact the renewal predictions for the plain line rail.

4.5.4 Ballast

The best fitting distribution for failures of the vertical alignment of the ballast is given in Table 4.12.

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
SD vertical alignment – 0 tamping interventions	Satisfactory	650	-	Weibull	1.09 [1.03, 1.15]	4759 [4417, 5128]	9.31
	Poor	650	-	Weibull	1.19 [1.12, 1.26]	8968 [8374, 9604]	23.16
SD vertical alignment -1 tamping intervention	Satisfactory	493	-	Weibull	0.87 [0.81, 0.92]	3833 [3440, 4270]	6.91
	Poor	493	-	Weibull	0.94 [0.89, 1.01]	8360 [7570, 9233]	15.51
SD vertical alignment - 2 tamping interventions	Satisfactory	410	-	Weibull	0.82 [0.77, 0.88]	3950 [3489, 4472]	6.92
	Poor	410	-	Weibull	0.92 [0.86, 0.99]	8662 [7745, 9686]	15.93
SD vertical alignment - 3 tamping interventions	Satisfactory	323	-	Weibull	0.77 [0.71, 0.84]	3284 [2828, 3815]	5.59
	Poor	323	-	Weibull	0.88 [0.81, 0.95]	7507 [6581, 8563]	13.56
SD vertical alignment - 4 tamping interventions	Satisfactory	222	-	Weibull	0.85 [0.77, 0.94]	2283 [1940, 2688]	4.06
	Poor	222	-	Weibull	1.01 [0.92, 1.12]	5328 [4639, 6119]	10.15
SD vertical alignment - 5 tamping interventions	Satisfactory	176	-	Weibull	0.80 [0.72, 0.90]	2293 [1890, 2783]	3.97
	Poor	176	-	Weibull	0.98 [0.88, 1.09]	5576 [4753, 6541]	10.51

(cont.)

Failure type	Fault	Uncensored lifetimes	Censored lifetimes	Fitted distribution	Parameter 1 [95% CI]	Parameter 2 [95% CI]	Expected lifespan (years)
SD vertical alignment- 6 tamping interventions	Satisfactory	122	-	Weibull	0.92 [0.81, 1.06]	2759 [2252, 3379]	5.07
	Poor	122	-	Weibull	1.05 [0.92, 1.20]	6265 [5241, 7489]	12.11
SD vertical alignment- 7 tamping interventions	Satisfactory	75	-	Weibull	0.76 [0.64, 0.90]	1589 [1165, 2170]	2.69
	Poor	75	-	Weibull	0.98 [0.83, 1.15]	4213 [3301, 5377]	7.94
Isolated vertical alignment	Satisfactory	445	643	Weibull	0.42 [0.39, 0.46]	26597 [20308, 34833]	30.45
	Poor	36	813	Weibull	0.46 [0.33, 0.64]	5789300 [573377, 58453720]	7149.93
	Critical 1	0(1)	848	Weibull*	1	5682860	15569.48
	Critical 2	0(1)	848	Weibull*	1	5682860	15569.48
	Critical 3	1	847	Weibull*	1	5682860	15569.48
	Critical 4	0(1)	848	Weibull*	1	5682860	15569.48

*Table 4.12: Fitted distributions for faults on the ballast *estimated distribution*

Due to the large amounts of available data, the times-to-failure for the SD of the vertical alignment were separated into groups with an equal amount of high-output tamping interventions before the analysis, to investigate whether tamping affects the ballast condition. After 7 tamps, the subgroups were too small for any significant analysis, hence it is assumed that the behaviour of the ballast beyond 7 interventions is identical to the behaviour of the ballast at 7 interventions. This separation was not applied to isolated vertical alignment faults however, because dividing the already small dataset further would introduce high uncertainty. An assumption was made that isolated defects would occur at the same rate regardless of the previous maintenance actions across the unit.

The datasets for the standard deviation of the vertical alignment across a unit of ballast to degrade from a good condition to a satisfactory or poor condition were best described by a Weibull distribution. The scale parameter mostly decreased as the number of tamps increased, finding that tamping reduces the threshold for when 63.2% of the predicted times-to-failure have been reached. The negative impact of tamping was also reflected in the expected lifetimes of the ballast before the fault was experienced, as this generally decreased the more maintenance was performed. The shape parameter always increased between the time taken to go from a good to satisfactory condition and from a good to poor condition but was often less than 1, implying failures are less common over time. There was minimal evidence of an increasing degradation rate in Table 4.12 as even when the fitted shape parameter was above 1, it was still very close by, hence degradation can be considered as randomly occurring rather than associated with the ballast's age. A general trend could initially be seen of the shape parameter decreasing with the number of tamping interventions performed, however this appears to have stabilised once 4 tamps had been completed.

A brief investigation was held to determine whether the presence of outliers in the datasets was impacting the fitted distribution parameters. 8.6% of the times-to-failure obtained from inverting Equation 4.3. were relatively large (>35 years) and could be considered outliers, as they existed beyond the fences of the interquartile range of the dataset (Walfish, 2006). Removing the outliers from the dataset did change the parameters of the Weibull distributions as the scale parameter decreased (which is to be expected) and the shape parameter increased slightly, however it mostly remained below or very close to 1. With no significant reason to exclude the outliers – they were not near the boundaries of any recorded sections of slab track which is expected to have a much longer lifetime, and often had data from over 15 geometry inspections to fit an appropriate regression line – they were included in the analysis, and it can be concluded that ballast units on the HS1 line are most likely to experience failure early after installation or rectification by high-output tamping, but this risk decreases as time goes on.

The 95% confidence interval for the parameters show a higher level of certainty compared to what has been seen in the rail-based assets, which is indicative of the much larger uncensored datasets that are available for analysis. The expected lifespans of the ballast to reach a poor condition (where high-output tamping is required to restore the ballast back to a good condition) obtained from the distributions were considered large compared to what was observed in the data. For example, the expected lifespan for a section of ballast with 0 previous tamping interventions to require tamping is 23.16 years using the parameters of the fitted Weibull distribution; finding the same statistic from the high-output tamping history using Equation 4.6 revealed an average time of just 6.90 years. Although the railway standards do state that a ballast section in a satisfactory condition may be considered for high-output tamping, the high disparity between the two times suggest that the ballast is rarely permitted to degrade to a poor condition and is maintained prematurely, potentially minimising its useful lifetime. This is investigated further in Chapter 5.

Focusing on the isolated vertical alignment faults indicate that these are a far less common occurrence for a unit of ballast, as the associated expected lifespans for the ballast to degrade due to isolated faults are significantly higher than the lifespans for faults arising from the standard deviation of the vertical alignment. The Weibull distribution was again chosen as the most appropriate to represent these faults, but this did require the assumption that no-occurrence faults could be modelled with a Weibull distribution with shape parameter of 1 and scale parameter equal to the MTTF. The shape parameters found from using the Weibull++ software, and its associated 95% CI, strongly confirm that isolated vertical alignment faults have a decreasing degradation rate over time.

4.6 Summary

This chapter has taken the novel approach of modelling the degradation of a track asset by individual faults within each failure type, as described by the railway standards. The times-to-failure for each fault were obtained by constructing a timeline of failure across every asset unit for the rail defect history and fitting regression lines to describe the development of faults described through the rail profile or track geometry history. For faults that hadn't occurred over the observation period, an assumption was made that a single unit experienced the fault one day after the observation period ended. The times-to-failure – whether censored or uncensored- were placed into a dataset and maximum likelihood estimation and the AIC and BIC were used to rank how well the one-parameter exponential, normal, two-parameter lognormal, and two-parameter Weibull fit the datasets. For datasets with a single uncensored time-to-failure, the MTTF was used to fit either an exponential or Weibull distribution with the beta parameter set to 1, and both were ranked equivalently. The overall best distribution for faults within each failure type was chosen by comparing the ranked values of the associated faults and choosing the distribution with the lowest average ranking.

The results discussed can be used by a track engineer to comprehend the likelihood of a fault occurring over a unit's lifetime, which can aid in maintenance management decisions – e.g. having very frequent inspections immediately after installation if many faults show a decreasing degradation rate. Additionally, the fitted distributions can be implemented into an asset management model, to study how maintenance can prolong the life of a railway component and predict when a renewal is necessary. Throughout this chapter, the small amount of available data has frequently been discussed. This research was completed when no wear-based renewals had occurred, and in many cases, the units were less than halfway through their expected lifespan. Condition data, such as the rail profile history or track geometry history, can help to understand the trend of degradation even if the unit itself not in a degraded condition, but this is not available across every fault for every asset.

It is commonly understood in degradation modelling that degradation follows a 'bathtub curve' (Sun et al., 2011; Thurlby, 2013). This is where the degradation rate is initially high when first installed but decreases as the asset enters its 'useful lifetime', where faults rarely occur. As the asset begins to age, the degradation rate starts to increase and faults are expected to occur more frequently- this is known as the 'wear-out phase'. The observation period for the data was only for the initial 12-19 years of service, hence it is unlikely that evidence of a wear-out phase has been observed and would explain why very few of the distributions described in section 4.5 indicate an increasing rate of degradation. This Is the biggest limitation to this research – if the modelling of degradation is not truly representative of real-world processes, the asset management model will not predict accurate results. Going forward, it is understood that the distributions found are the best representation so far for understanding the behaviour of an asset on the HS1 line, but it is important to stress that the developed Petri nets are largely an example of a methodology for modelling degradation and renewal on a high-speed line, and any results found may be unrealistic. Revisiting this research when all units have experienced at least one renewal will improve the accuracy of any predictions made.

Chapter 5

Modelling high speed railway track renewal

5.1 Introduction

Petri nets (PN) were previously chosen as an appropriate modelling technique to represent the management of track and predict when renewal is required, due to their flexible stochastic nature. This chapter discusses how the structure of a PN, including its various components and key features (defined fully in Chapter 3), has been applied to represent the complexities of the inspection, degradation, maintenance, and renewal processes. The developed model is populated using the distributions established in Chapter 4 and results are found through Monte Carlo simulation of the nets. These results, and several key examples where the PN can improve upon known limitations of current asset management models and renewal predictions, are discussed.

5.2 Track model

The various features of Petri nets are combined to create models representing the asset management strategy of a track component. The models can be considered as a series of interconnected sub-models – inspection, degradation, maintenance, and renewal. This section describes the structure of each sub-model.

5.2.1 Inspection

Every asset has an identical inspection process sub-model, shown in Figure 5.1.

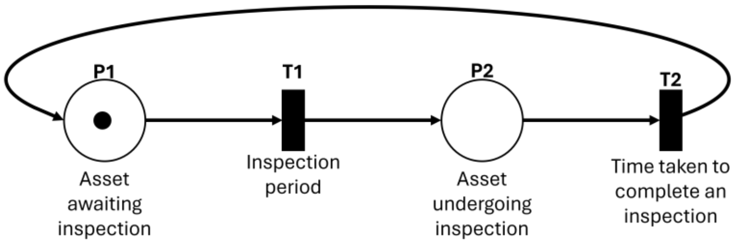

Figure 5.1: Inspection process sub-model

Although some areas of the track have multiple types of inspection – the plain line rail has the UTU inspection and then potentially a further manual inspection if a defect was detected – an assumption was made that a single inspection would detect all the faults present on the unit. This assumption was necessary as the extent to which different inspection methods could detect which faults was unknown. Initially, a token is present in P1 to indicate that the unit is awaiting an inspection. Then, after a length of time determined by the inspection frequency, T1 fires moving the token to P2, and the unit undergoes an inspection. After the inspection has been completed (which is assumed to take a total of one day), T2 sends the token back to P1, and the unit is again awaiting an inspection.

5.2.2 Degradation

Railway standards have been used to identify the different failure types occurring on a track asset, as well as the stages of degradation within this failure which are considered as separate faults. This was used to build the structure of the degradation process sub-model. As every fault has independent maintenance requirements, including type of corrective maintenance, timescales for completion, and additional preventative actions such as enacting speed restrictions or increasing the inspection frequency, it was required for every failure type to have its own degradation process modelled via a separate sub-model. This increases the overall size of the PN but makes it better able to represent real-world engineering processes, which improves the reliability of any lifetime predictions.

The PN structure for the degradation of one failure type is shown in Figure 5.2, where the failure has two defined faults (P4 & P5). P5 is the most severe fault and can only be reached once the less severe fault (P4) has already occurred. Initially a token is present in P3, to indicate that the failure hasn't appeared yet. The token moves through P3-T3-P4-T4-P5 to represent the development of the failure, e.g. the increase of a squat size on the crossing panel. The structure of the degradation process sub-model varies depending on the failure being modelled; 'Transverse cracking in the crossing vee' has 5 associated faults each increasing in severity, whereas 'Lipping on casting' is a singular fault – lipping is either present on the casting or not.

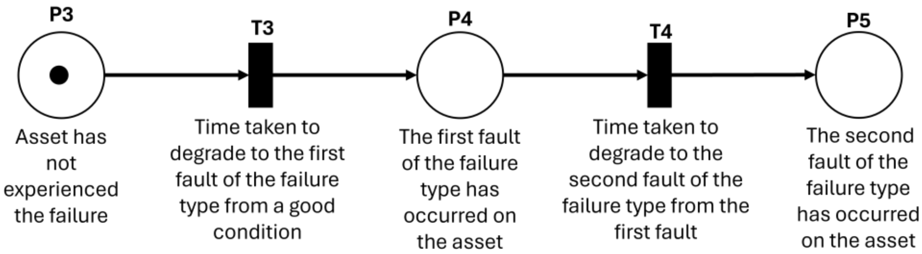

Figure 5.2: Degradation of a failure type

An asset can be in a degraded state, but a track engineer is only aware of this during an inspection. Acknowledgement of a fault is represented in the PN by additional places P6 & P7 (Figure 5.3), where a token can only move into these places if a token is present in P2, <u>and</u> a token is present in either P4 or P5. Transitions T5 and T6 have a firing time of zero, as a degraded state is immediately noticed during an inspection.

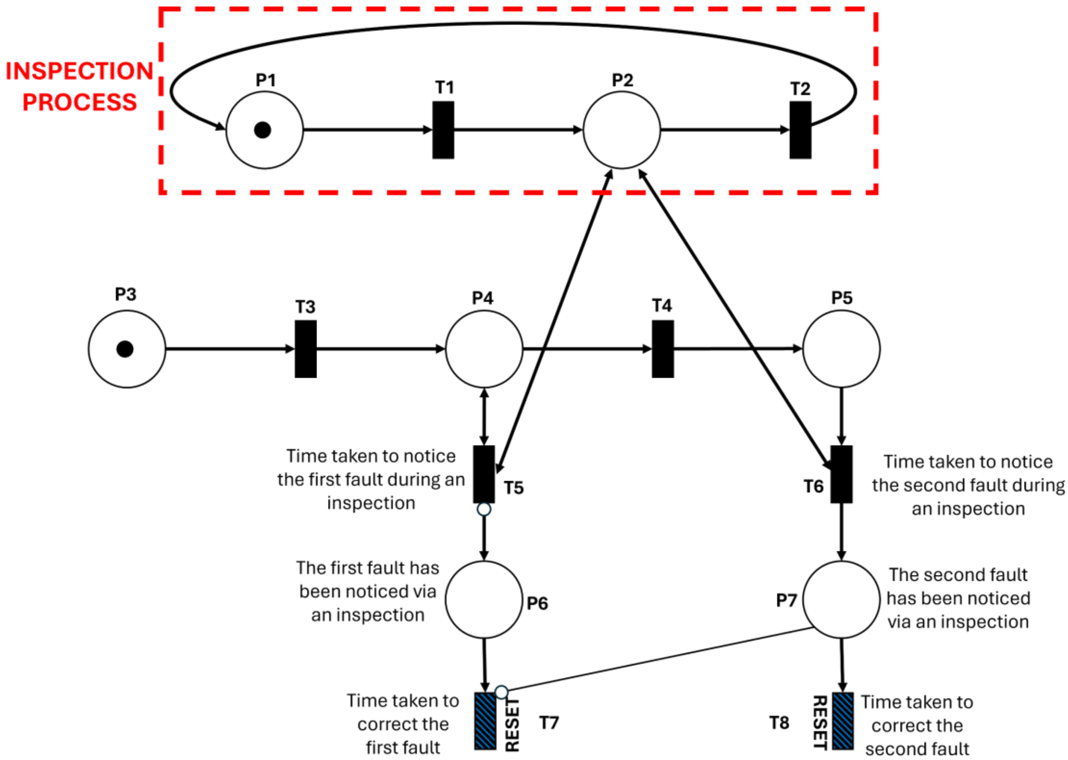

Figure 5.3: Degradation process sub-model

If an inspection has revealed that the asset is in a degraded state but that it could still degrade further before maintenance can occur (P6), a token is returned to the degraded place (P4) so that the degradation process may continue to develop, and an inhibitor arc is used to ensure that future inspections do not re-record the same fault. Similarly, when an inspection has noticed a fault, a token is returned to the 'undergoing inspection' place (P2) so that further faults can be detected within a single inspection. Inhibitor arcs are also used to ensure that if the failure type has multiple faults that have been noticed via an inspection (P6 & P7), only the most severe fault can trigger a maintenance action, so that maintenance is not counted multiple times for the same failure.

T7 and T8 represent the times for maintenance to be completed. They are also reset transitions where once fired, will reset all transitions and places associated with the specific failure (P3-P5 and T3–T6) back to their original conditions, to indicate that maintenance was completed. In total there are 2 failure types on the ballast, 16 on the switch panel, 24 on the crossing panel, and 29 on the plain line rail, all of which have their own degradation sub-model. An assumption is made that if multiple faults are awaiting repairs on the same unit, they are completed at the same time, corresponding with the fault that has the shortest timeframe for requiring maintenance. This consideration of opportunistic maintenance in the models is included to better mimic real practises, as it is more cost effective to correct multiple faults at once when an engineer is available to complete the repairs. It is realised through place dependent transitions interconnecting every 'acknowledged fault' place (P6 & P7 in Figure 5.3) to every 'correction' transition (T7 & T8 in Figure 5.3) across all separate degradation sub-models, to modify the firing time to the lowest enabled transition. An example of this process is given in Figure 5.4, where P3-P7 represents the degradation sub-model of one failure type and P8-P10 is the degradation sub-model of a separate failure type. Both are in an acknowledged degraded condition, shown by tokens in P6 and P10. T7, T8, and T11 are all place dependent transitions, with T7 and T8 dependent upon P10, and T11 dependent on P6 or P7. As a token is present in P10, the T7 firing time is adjusted from 90 days to 10 days, hence T11 and T7 will fire at the same time to indicate opportunistic maintenance has occurred. Note that the firing time for T11 is not adjusted as it is currently the lowest enabled transition.

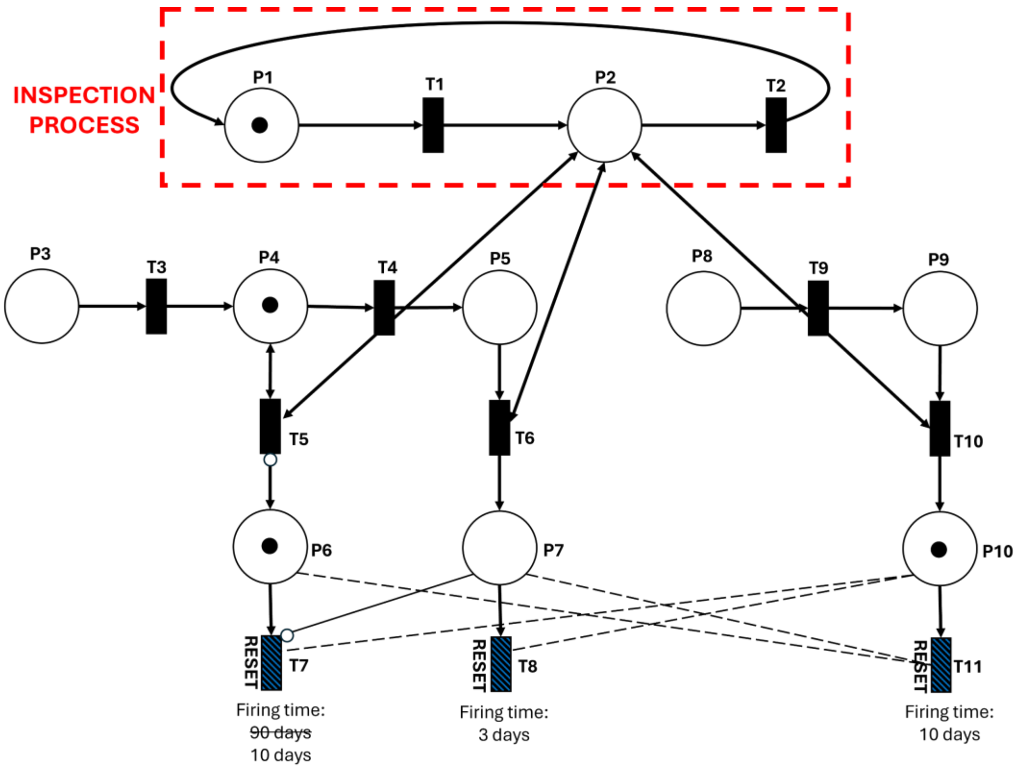

Figure 5.4: Implementation of opportunistic maintenance between separate failure types

Finally, there are several failure types on the rail-based assets that refer to defective weld repairs. For the plain line rail it is assumed that any defective aluminothermic or unstated weld types can include those made during installation, hence the associated degradation sub-model is structured as in Figure 5.3. For the switch and crossing panels which do not have any initial welds as part of their design, and for defective MMA welds on the plain line rail, the degradation sub-model is modified slightly so that a token is only placed into the 'Asset has not experienced the failure' place (P3) after a weld repair has been completed at least once on the unit. This is to ensure that these failures cannot occur before an actual weld repair has happened.

5.2.3 Maintenance

Once a fault on the track component has been noticed via an inspection, it must be corrected. Maintenance actions can be considered in two ways – corrective maintenance will restore the asset back to its optimal condition whereas preventative maintenance will not change the condition but does aim to prevent further damage developing or to protect passengers travelling on the track. Table 5.1 summarises each action, including what asset and area the maintenance is performed on.

Maintenance action	Assets applicable	Area of asset	Type	Description
Manual grinding	Crossing panel, switch panel, plain line rail	Localised	Corrective	Removes a small layer from the surface of the rail using a hand-held grinder
Cyclic grinding	Plain line rail	Entire	Corrective	Regularly removes a small layer from the surface of the rail using a grinding machine running along the track
Welding	Crossing panel, switch panel, plain line rail	Localised	Corrective	Adds material to cover defects by fusing the metals of the rail and a filler metal together
High-output tamping	Ballast	Entire	Corrective	Shifts large sections of the aggregate using vibrating tines to rectify the geometry
Local tamping	Ballast	Localised	Corrective	Shifts small sections of the aggregate using vibrating tines or manually using shovels
Lubricate	Switch panel	Entire	Preventative	Adds a lubrication material to the rail
Impose speed restrictions	Crossing panel, plain line rail, ballast	Entire	Preventative	Reduces the speed of rolling stock travelling over the asset
Block the line	Ballast	Entire	Preventative	Stops all rolling stock from travelling over the asset
Ban to-facing moves	Switch panel	Entire	Preventative	Forces the S&C to act as a section of plain line rail by stopping any movement of the switch rails
Appoint watchman	Crossing panel	Entire	Preventative	Hires a watchman to be present whilst maintenance is performed
Fit emergency fishplates	Plain line rail	Localised	Preventative	Bolts two sections of rail together
Increase inspection frequency	Switch panel, Crossing panel, plain line rail, ballast	Entire	Preventative	Reduces the interval between inspections whilst the fault remains
Renew	Crossing panel, switch panel, plain line rail	Entire	Corrective	Replaces the entire unit

Table 5.1: Viable maintenance actions per asset

The standards are used to identify which of the actions described in Table 5.1 should apply per fault. Occasionally, a fault on a rail-based asset is labelled as 'repair', as it can be rectified through either manual grinding or welding. The decision of how to repair is up to the track engineer and will depend on factors such as the availability of maintenance equipment, severity of fault, overall condition of the unit, remaining maintenance budget, and the time available to complete maintenance. To model this arbitrary choice, a decision-making strategy was implemented into the Petri nets (Figure 5.5).

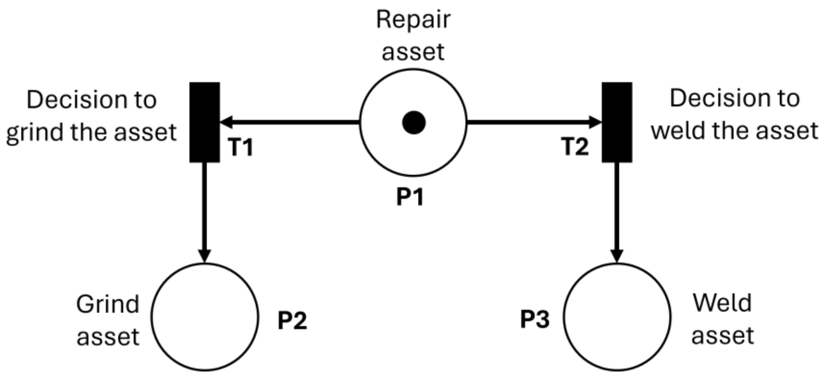

Figure 5.5: Decision-making strategy for maintenance

The ratio between grinding to welding is estimated by observing the historical maintenance data of the relevant faults of each asset. An example of a 70:30 ratio is given by equating T1 to a uniform distribution with parameters 0 and 1, and T2 a constant transition with a firing time of 0.7 days. 70% of the time, sampling from the uniform distribution will output a time lower than 0.7 days and T1 will fire, moving the token from P1 to P2 (the 'grind asset' place) before T2 can move a token to P3.

The degradation of the ballast is shown through the vertical alignment, which has two separate failure types – the standard deviation of the vertical alignment over a 200m unit, and isolated measurements of the vertical alignment. The relevant standards assert that when either failure type has degraded to a state deemed 'satisfactory', a review of the ballast must take place and corrective maintenance can be considered, but this is not a requirement until the failure degrades further to a 'poor' condition or beyond (Figure 5.6). This decision has been enacted in the ballast PN using the process described in Figure 5.5 to either correct the failure once it reaches a satisfactory condition, or to leave it until it degrades further. Place dependent transitions are then used to change the likelihood of either action to replicate the review process – if the standard deviation <u>and</u> isolated values of the vertical alignment have both degraded to 'satisfactory' or beyond, the probability of immediate maintenance changes to 1 and high-output tamping is completed to restore the condition.

Table 2 – Track Quality Status Bands	
Construction	These tolerances shall be achieved for new or renewed assets
Good	Good track quality with geometry tolerance within Target Value (TV), the normal operational band.
Satisfactory	Warning Value (WV) refers to the value which, if exceeded, requires that the track geometry condition is analysed and considered in the regularly planned maintenance operations.
Poor	Action Value (AV) refers to the value which, if exceeded, requires corrective maintenance in order that the immediate action limit shall not be reached before the next inspection.
Critical	Speed Restriction Value (SRV) refers to the value which, if exceeded, requires taking measures to reduce the risk of derailment to an acceptable level. This can be done either by blocking the line, reducing speed or by correction of track geometry. 'Block the line' faults shall be verified (on the train or by manual measurement) within 5 minutes of recording and fault control advised immediately once the defect is confirmed. All other SRV geometry faults shall be reported to the AFC within 60 minutes of discovery and local staff notified (via AFC) within 90 minutes of discovery.

Figure 5.6: Track quality status bands (Network Rail (High Speed) Ltd, 2020c)

Studying the maintenance history for the HS1 line found that 44% of 'satisfactory' isolated ballast faults were not corrected within 6 months of detection and were instead permitted to degrade further. For the standard deviation, observing the last recorded measurement before a high-level tamping intervention revealed that the ballast was frequently maintained prematurely, with only 18% of those that degrade to 'satisfactory' allowed to degrade further to a 'poor' condition. This investigation also found that 66% of the sections that were tamped were still considered in a 'good' condition. The developed PN's are structured based on the information given in the standards, and as such there is no option to correct the ballast whilst it remains in a good condition, but this finding does explain the high disparity between the lifetimes to reach a poor condition and the lifetimes between tamping interventions, discussed in section 4.5.4.

Condition when tamped	Good	Satisfactory	Poor
%	66.54	27.53	5.93

Table 5.2: Historical condition of the SD of the vertical alignment when tamped

The decision-making strategy will move a token to the correct place on the maintenance process sub-model, or in the case of the ballast, allow the degradation process to continue if 'don't maintain' is chosen. The ratios described in Table 5.3 were applied to every relevant fault through the asset, but this can be adapted to be more specific if the exact grind/weld ratio per fault is known.

Asset (failure type)	Grind/maintain	Weld/don't maintain
Crossing panel	58	42
Switch panel	56	44
Plain line rail	77	23
Ballast (SD)	82	18
Ballast (isolated)	56	44

Table 5.3: Ratios implemented in the decision-making strategy

The maintenance process sub-model varies depending on what actions are permitted on the asset (Table 5.1), but the switch panel sub-model is given in Figure 5.7 as an example:

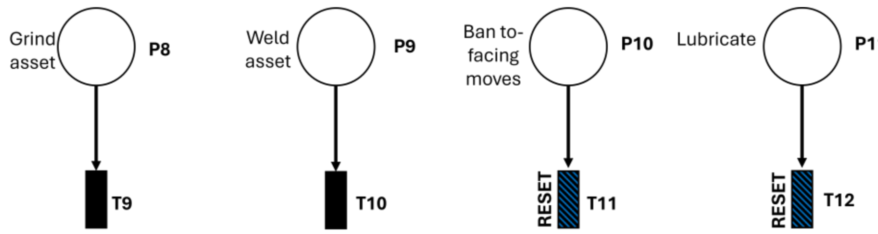

Figure 5.7: Maintenance process sub-model

Places P8-P11 represent the different maintenance actions that can be completed to rectify a degraded switch panel; often, multiple actions are required for the same fault. T9-T12 are immediate transitions with a firing time of zero- they exist to move the tokens from a maintenance place into the renewal process, so that statistics can be collected regarding what actions were completed during a simulation of the PN. Any maintenance actions that occur across the entire unit (described in the 'area of asset' column of Table 5.1) have a reset transition to return the associated input places back to their original marking of zero tokens after firing once (T11 and T12). This was necessary as an inspection can often detect multiple faults at the same time, and it must be ensured that these actions are not charged multiple times when calculating the operation costs.

Most maintenance can be modelled as in Figure 5.7, but there are several exceptions. An increased inspection frequency is realised by connecting the 'known fault' place of the relevant fault's degradation sub-model (Figure 5.3) to the inspection period transition (Figure 5.1) via a place dependent transition. For faults with maintenance actions that are required daily until the fault is rectified – imposing speed limits, banning to-facing moves, and blocking the line - an additional transition is inserted that sends a token from the 'known fault' place to the associated maintenance place with a firing time of 1 day, and returns a token to the 'known fault' place. For faults with a maintenance action of renewal, a token is sent directly to the place representing the number of renewals, and all the places and transitions of the PN (except for the 'number of renewals' place) are reset to their original condition. Finally, for faults maintained by cyclic grinding, the token is retained in its 'known fault' condition, until the transition associated with cyclic grinding fires, which is assumed to happen at a frequency of every 2 years.

For clarification, this process is shown further in Figure 5.8. In step 1, the degradation sub-model described by P1-P5 is in a known degraded state that is rectified by cyclic grinding. Then in step 2, (as the firing time t of T7 is shorter than the firing time k for T2) a token is moved from P7 to P8 to indicate that the cyclic grinding campaign of the plain line rail is underway. This enables the transition T5 to fire, and as it is an instant reset transition, the degradation sub-model is reset back to its original markings to imply that cyclic grinding has rectified the fault. In step 3, T8 (with a firing time of 1 day) has fired to move a token from P8 to P7, to indicate that the cyclic grinding campaign is over. However, had T2 fired before T5 was enabled (meaning that the failure had degraded further before the cyclic grinding campaign had taken place), T5 would have been inhibited from firing due to the token present in P3 or P5, and the rail would still be in a degraded condition after the campaign hence deeper, more localised manual grinding is required to remove the fault.

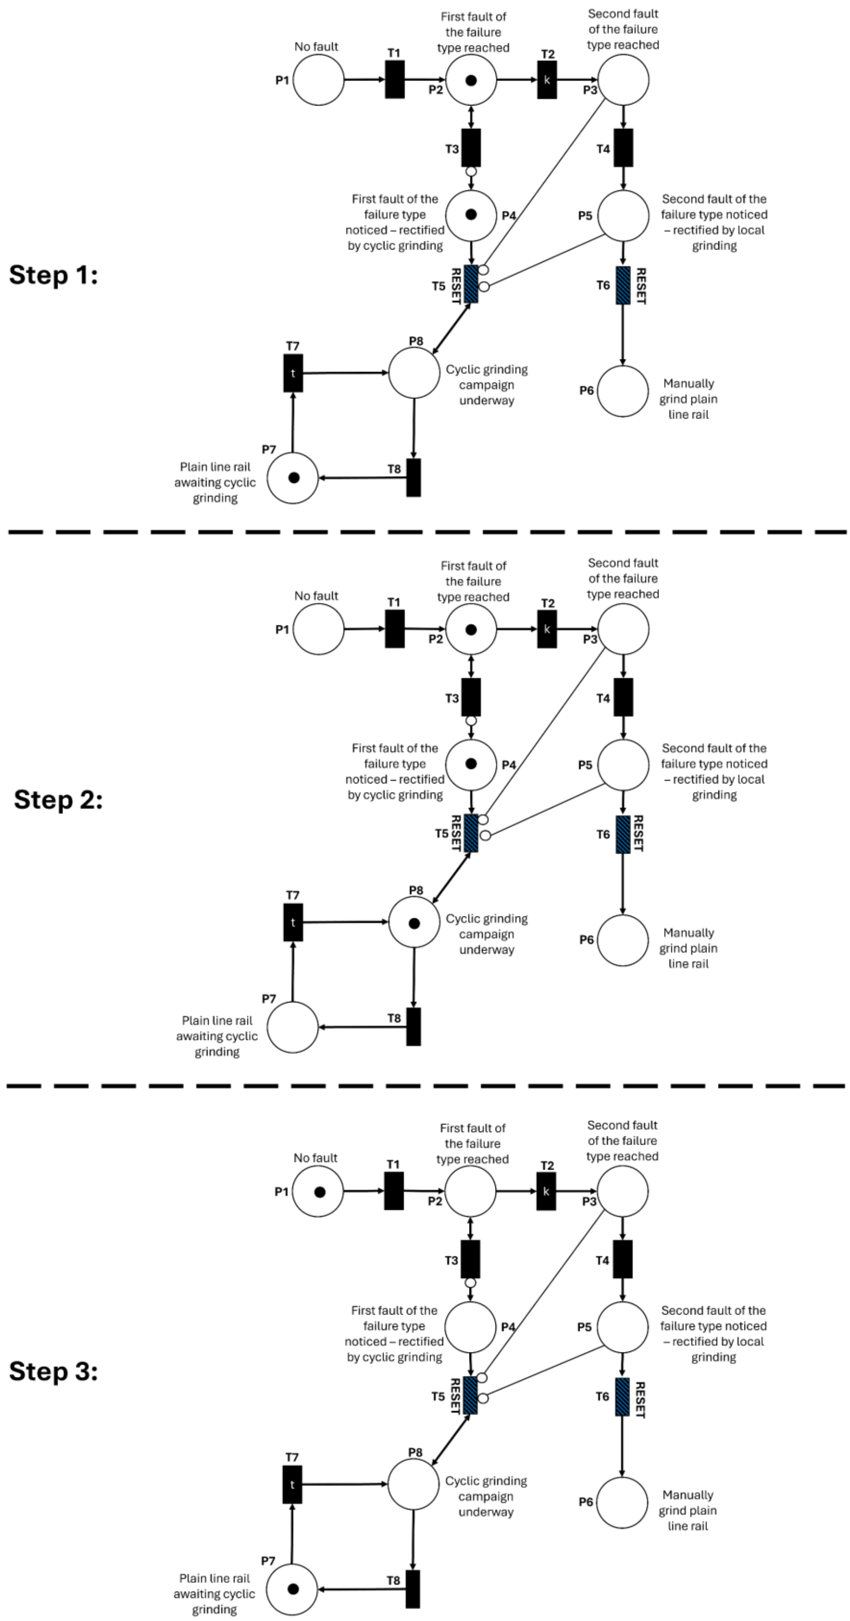

Figure 5.8: Modelling of cyclic grinding to rectify faults

5.2.4 Renewal

The final sub-model of the asset management strategy is the renewal process. Typically, the renewal threshold implemented in other models has been based around a fixed lifespan, however this is not considered as a viable strategy as it does not consider the behaviour of the asset prior to renewal, and any lifespan predictions will be equivalent to the predetermined fixed limit. This research has highlighted the need for multiple definitions of wear to be considered when triggering a renewal, with the ability to choose between them and investigate the impact on an asset's lifetime. Within the developed models, five different renewal criteria are considered – the number of maintenance interventions performed (ML), the number of interventions performed in a set period of time (SP), the frequency of successive failures (TD), the amount spent on operating the asset (AC), and finally the condition of the asset after maintenance has been completed (CB). Their implementation in the PN is shown in Figure 5.9, with the separate criterion highlighted to make distinction clearer. Additional places (also seen in Figure 5.9) are defined to represent a renewal being scheduled (P30), and to count the total number of renewals that have occurred (P31). T27 represents the time between scheduling a renewal and the renewal taking place – upon firing, it also resets the entire Petri net back to its original marking, except for P31.

Places P25-P29 represent the renewal criterion enabling places. If a token is present in these places, it blocks the corresponding transition – T17, T18, T21, T24, T25, or T26 – from firing via use of an inhibitor arc. This prevents a token from moving through the blocked transition to place P30, triggering a renewal. The complete list of renewal strategies that the PN can implement are shown by the matrix in Equation 5.1. Each row represents a different strategy, with the value in each column representing the number of tokens present in the associated place. For example, the first row of the matrix shows the renewal strategy depicted in Figure 5.9, as a single token is present in P26, P27, P28, and P29, and no token present in P25. This enables only T17 to fire and move a token to P30, meaning the chosen renewal strategy is based solely on the number of maintenance interventions performed.

$$\begin{array}{ccccc} \boldsymbol{P25} & \boldsymbol{P26} & \boldsymbol{P27} & \boldsymbol{P28} & \boldsymbol{P29} \end{array}$$
$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \tag{5.1}$$

The matrix contains 31 rows, meaning the impact of 31 different renewal strategies on the asset lifespan can be considered. The final row of the matrix corresponds to a renewal strategy considering all renewal criteria.

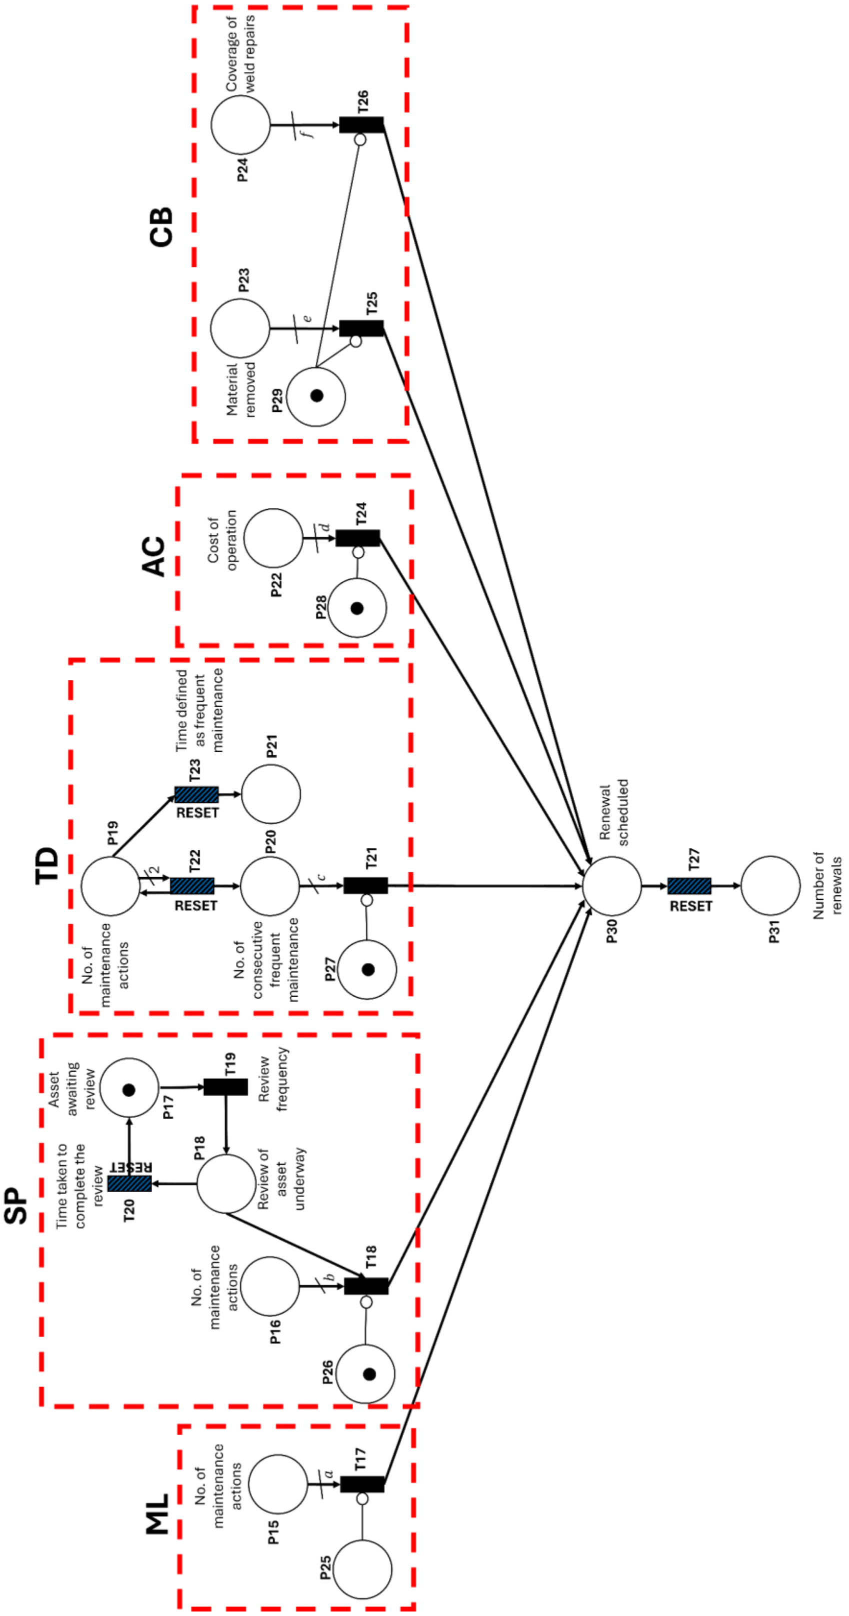

Figure 5.9: Renewal process sub-model

For additional clarification, Table 5.4 is provided to describe the purpose of each place seen in Figure 5.9 in more detail:

Place	Description
P15	Counts the number of corrective maintenance interventions performed on the asset
P16	Counts the number of corrective maintenance interventions performed on the asset since the previous review was held.
P17	Represents that the asset is awaiting a review
P18	Represents that the asset is currently undergoing a review
P19	Counts the number of corrective maintenance interventions performed on the asset during a specified duration of time
P20	Counts the consecutive number of times two interventions have been performed within a specified duration of time
P21	Represents a spare place for tokens to be moved to if two maintenance interventions were not performed within the specified duration of time, so that the TD renewal criterion threshold is not met prematurely
P22	Counts the amount spent from operating the asset
P23	Counts the amount of material removed from the asset by grinding interventions
P24	Counts the percentage of the asset covered by welding interventions
P25	Enables or prevents the ML criterion to be included in the implemented renewal strategy dependent on its marking
P26	Enables or prevents the SP criterion to be included in the implemented renewal strategy dependent on its marking
P27	Enables or prevents the TD criterion to be included in the implemented renewal strategy dependent on its marking
P28	Enables or prevents the AC criterion to be included in the implemented renewal strategy dependent on its marking
P29	Enables or prevents the CB criterion to be included in the implemented renewal strategy dependent on its marking
P30	Represents that the asset is scheduled for a renewal
P31	Counts the number of times that the asset has been replaced

Table 5.4: Description of the places seen in Figure 5.9

Finally, the implementation of each renewal criterion in the model is described below.

1) Maintenance Limit (ML)

   This criterion defines wear-out by the amount of corrective maintenance actions that have occurred on the unit. There is no distinction between what type of maintenance is completed, nor when in an asset's lifetime the faults have occurred. If the number of tokens in P15 has exceeded a given maintenance threshold (a in Figure 5.9), T17 is fired to schedule a renewal.

2) Set Period (SP)

   This renewal criteria aims to mimic a review or control period of the asset by considering wear based on the number of faults that have occurred within a given time-period. P17-T19-P18-T20 is the review process, where a token in P18 indicates a review occurring, and T19 is the review-period interval. If the number of faults that have occurred within the review-period exceeds a threshold (b in Figure 5.9), T18 is fired, and a renewal is scheduled. If this limit is not exceeded, the firing of T20 sends a token to P17 to restart the review process and removes all tokens in P16 via the use of a reset transition.

3) Time Duration (TD)

   The Time Duration renewal policy considers an asset at the end of its useful life if corrective maintenance is too frequent, implying that it is no longer effective at restoring the condition of the asset. Every time corrective maintenance is completed, a single token is placed in P19, enabling the constant transition T23. The TD criterion is then enacted in the PN by first considering the time elapsed between two successive maintenance interventions – if this is shorter than a specified duration (modelled by the firing time of T23), T22 is enabled as two tokens are present in P19, satisfying its input arc's multiplicity, and it immediately fires to place a single token in P20 and back into P19. As T22 is a reset transition, this action will also resample the firing time of T23, which is instantly enabled again due to the placing of a token back in P19. However, if the firing time of T23 has elapsed since being enabled and before two tokens are present in P19, T23 fires to move the token from P19 into P21. As a reset transition, it will also restore P20 to its original marking by removing any tokens present within. P21 acts as a spare place for tokens to be moved into, to prevent T22 from simply being enabled whenever multiple repairs have been completed on the asset. Given that two corrective maintenance interventions completed within a short period of time may simply be coincidence rather than an indication of a worn asset, this renewal criterion is based around the consecutive number of times two interventions has been completed within the specified duration, by storing tokens in P20. If this has reached a threshold (c in Figure 5.9), T21 is enabled and immediately fires to schedule a renewal. Unlike the Set Period criterion, the Time Duration strategy focuses on the time between faults occurring, not the total amount of faults in a time-period.

4) Associated Cost (AC)

   A frequent factor in many asset management decisions is the available budget, which is limited and set before the start of every control period. Renewal strategies based solely on corrective maintenance can keep the asset functional indefinitely, but this may come at a great cost. The AC renewal strategy is based on accrued expenditure from operating the unit, to ensure that the asset management decisions and scheduling of renewals will always remain within the confines of the projected budget. Wear-out in an associated cost renewal strategy is considered reached when the unit's operational cost has reached a threshold (d in Figure 5.9), causing T24 to fire and schedule a renewal. The current total cost of operating the asset is stored in P22. By considering the renewal based on a set limit of actions (in this case, costs), this criterion is similar to the Maintenance Limit criteria. However associated cost also considers the impact of all management operations (not just corrective maintenance), defined here as any hiring of maintenance equipment, staff labour, sourcing of materials, or financial penalties resulting from a reduced service.

This renewal criterion requires an assumption on the approximate cost of various actions, which was obtained through a mixture of discussions with track engineers and estimations. The cost relates to the amount spent on a single unit of the asset, and is defined by financial units (Table 5.5):

Operation	Cost (units)
Manual grinding	130
Cyclic grinding	17
Welding	200
High-output tamping	50
Local tamping	20
Lubricate	3
Impose speed restrictions	20 (per day)
Block the line	100 (per day)
Ban to-facing moves	40 (per day)
Appoint watchman	15
Visual inspection	20
TRV inspection	2
UTU inspection	2

Table 5.5: Assumed cost (in units) per maintenance action

5) Condition Based (CB)

   The Maintenance Limit, Set Period, and Time Duration renewal criteria are all based around the occurrence of corrective maintenance, but do not consider the type of maintenance used nor the severity of the faults being repaired. Both factors will affect the impact of the maintenance on the condition of the asset, which in turn will affect how quickly the asset is considered to be unrepairable – e.g. the aggregate is too small, or weld repairs cover too large a proportion of the rail. This definition of wear is instead modelled through the condition-based renewal policy. P23 represents the amount of material removed by grinding, and P24 represents the coverage of any weld repairs. In Figure 5.9, if more than $e$mm of material has been removed via grinding, or weld repairs cover f% of the asset, T25 or T26 fires to schedule a renewal. For the ballast PN, this is modified slightly so that P23 represents the percentage of aggregate that has been reduced to an unacceptable size by tamping, and if more than e% has been damaged, T25 fires to schedule a renewal. P24 and T26 are then removed from the structure of the ballast PN.

The impact of local grinding, welding, localised tamping, or high-output tamping on the condition of the asset is specialised to every fault, allowing for more flexibility within the model. For example, a weld repair of a crack >75mm long will cover a greater proportion of the asset compared to a crack between 10-25mm long. Information regarding the depth of rail manually grinded to remove a fault or the length of a weld repair, was scarcely included in the historical maintenance data; hence the assumed impact is as described in Table 5.6. For the ballast, research has found that high-output tamping was expected to damage the aggregate by between roughly 2-6% for every intervention (Aursudkij, 2007), and an assumption was made that localised tamping would be 4x less intrusive given it spans a much smaller area. The assumed impact of tamping is shown in Table 5.7.

	Severity 1	**Severity 2**	**Severity 3**
Severity ranking explained	Recommended maintenance timescale for the fault is longer than 28 days	Recommended maintenance timescale for the fault is between 8-28 days	Recommended maintenance timescale for the fault is shorter than 8 days
Impact of local grinding	1mm of material is removed from the rail	2mm of material is removed from the rail	3mm of material is removed from the rail
Impact of welding	2% of the rail is covered by a weld repair	4% of the rail is covered by a weld repair	6% of the rail is covered by a weld repair

Table 5.6: Impact of corrective maintenance on the condition of rail-based assets

Maintenance	**Impact**
High-output tamping	The aggregate is damaged by 4%
Local tamping	The aggregate is damaged by 1%

Table 5.7: Impact of corrective maintenance on the condition of the ballast

What remains unknown is the impact of cyclic grinding on the condition of the plain line rail. The railway standards state this can remove between 0.2-3mm of the rail head on every pass. A weighted arc can be used to deposit a set number of tokens into the 'Material removed' place for the Condition Based renewal criterion every time cyclic grinding is completed, however this assumes that the same amount of material is always removed for every unit of plain line rail. This approximation can be improved by considering the transition for the 'vertical wear' failure, which is modelled on the length of time for the rail to reach the absolute minimum depth of 14mm. Assuming that this rate of degradation is a combination of friction of the rail and material removed from cyclic grinding, and because it was found that the vertical wear was best fitted by a linear regression line implying degradation is constant, it can be used to approximate how much material would be removed from the rail between cyclic grinding campaigns. The implementation of this is shown in Figure 5.10.

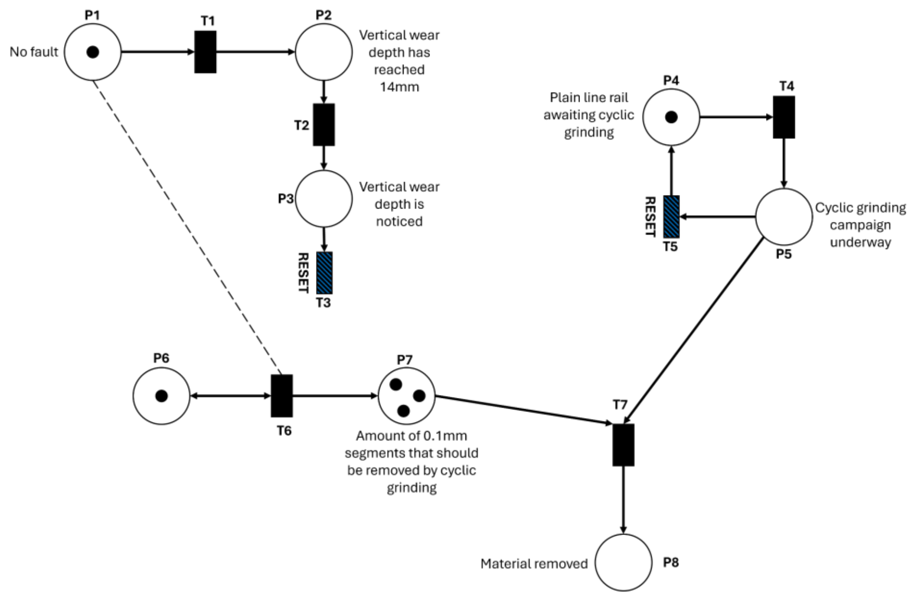

Figure 5.10: Impact of cyclic grinding on the material removed from the rail

P1-T1-P2-T2-P3-T3 represents the degradation sub-model of the vertical wear fault, and P4-T4-P5-T5 is the cyclic grinding intervention. P6-T6-P7 counts the amount of 0.1mm segments that should be removed from the rail so far by changing the firing time of T6 to the firing time of T1/140, via a place dependent transition. When the cyclic grinding campaign is underway, a token is moved to P5, enabling the transition T7, which has a firing time of zero – it fires instantly. This will move all the tokens in P7 (the place representing the amount of material that should be removed from cyclic tamping) into P8, the place representing the amount of material removed from the rail in the Condition Based renewal criterion. This implementation allows for a renewal policy based on the amount of material removed from the rail to be separate from renewal due to the absolute minimum depth of the rail reached, the advantages of which is discussed further in section 5.5.3.4.

5.2.5 Entire asset model

The complete Petri net structure for a track asset is obtained by connecting the sub-models of the inspection, degradation, and maintenance processes, and the renewal options. An example for a track asset with two failure types is shown in Figure 5.11. Due to the size of the developed PNs, it is not possible to reproduce them in graphical form here, however they all share identical sub-model structures as those described throughout this section. Appendix A.1-A.4 gives a description of these models in a textual form.

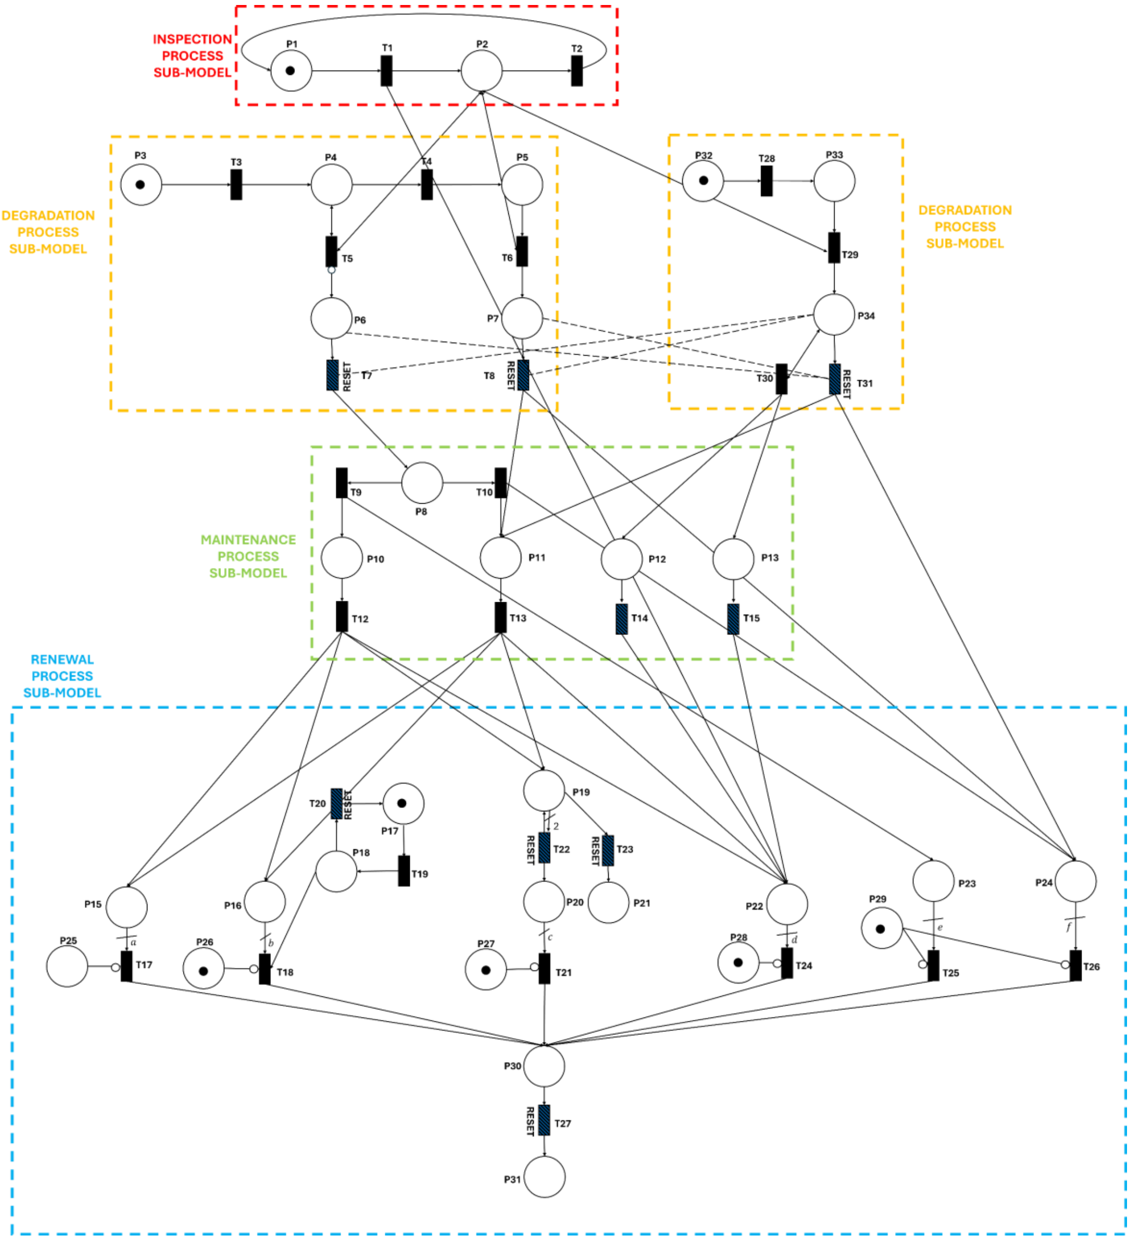

Figure 5.11: Entire asset model

5.2.6 Assumptions

The developed models were created to represent the management strategy of a track asset as closely as possible. However, some assumptions were still required to make the models functional, namely:

- A single inspection will detect every fault present on the asset
- Maintenance is completely effective at restoring the condition of the asset to as good as new
- Localised corrective maintenance will only restore the condition of the affected fault, but entire corrective maintenance (such as high-output tamping or cyclic

grinding) will restore the condition of every associated fault, even if it is not currently degraded

- Maintenance machinery is always available when needed

5.3 Analysis of models

5.3.1 Construction

The high complexity of modelling asset management decisions requires the inclusion of added features such as place dependent transitions, inhibitor arcs, and reset transitions. Most commercially available PN simulation programs are unable to accommodate these features, and a bespoke code was written in C++ to create and simulate the developed PN structure. This was used to investigate the degradation behaviour of the track over its lifetime.

Due to the separate modelling of every failure type, the PN's were comparatively large - in total, the crossing panel PN had 198 places and 275 transitions, the switch panel had 101 places and 121 transitions, the plain line rail had 224 places and 315 transitions, and the ballast PN had 43 places and 47 transitions. A single simulation of the PN lasted an average of 0.038 seconds, allowing for an extensive analysis of the track behaviour over repeated simulations.

5.3.2 Transition firing times

The firing times for transitions relating to engineering procedures – inspection intervals and timeframes for maintenance – are derived from the railway standards. The timeframe for maintenance to be completed varies depending on the type and severity of the fault, but a 'worst case scenario' where maintenance takes the maximum permitted time is assumed. An assumption has also been made that the time between scheduling and completing a renewal is 14 days, based upon similar times detailed in the standards for other corrective maintenance actions such as grinding or welding. Although a potentially unrealistic assumption, as the timetabling of rolling stock over the studied line as well as availability of staff and equipment is likely to be more impactful when completing renewals compared to less intrusive actions, the amendable structure of a PN means this time can be adapted if more information surrounding renewal procedures is available, or even derived from a continuous probability distribution such as with the stochastic transitions located between degraded states of a fault.

Figure 5.12 describes the span of the predicted times-to-failure obtained via the distributions fitted during the reliability analysis in Chapter 4. For a typical degradation of a failure type with two associated faults, the distributions T0 and T1 will predict the times to degrade from a good condition to the first fault of the failure type, and from a good condition to the second fault of the failure type respectively. However, the distribution representing the time to degrade from fault 1 to fault 2 (T2) is unknown.

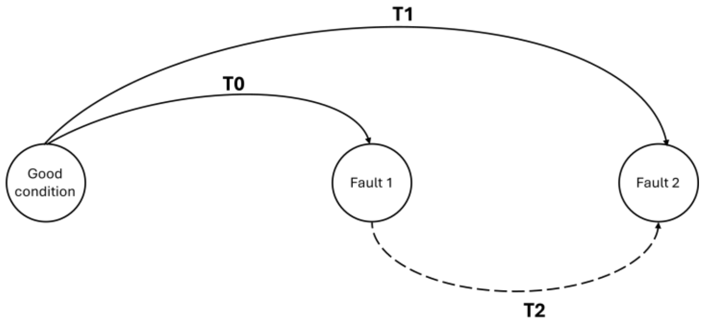

Figure 5.12: Span of distributions fitted to each fault

The degradation sub-model in the developed PN will use only T0 and T2, hence the parameters for T2 must be estimated. For failure types modelled via the exponential distribution, which utilises MTTF, this is simple:

$$MTTF_{T2} = MTTF_{T1} - MTTF_{T0} \tag{5.2}$$

T2 can then be represented by the distribution $\sim E(\frac{1}{MTTF_{T2}})$. If $MTTF_{T0}$ and $MTTF_{T1}$ have identical mean times to failure, as is often the case for distributions derived from faults that did not occur during the observation period of the HS1 line, $MTTF_{T2}$ was set equal to 1 day.

For failure types not fitted with the exponential distribution, the parameters for T2 can be estimated iteratively using convolution integrals (described further in Andrews's paper (2013)), but this requires prior analysis to be completed and was stated to be unsuitable for Weibull distributions with a beta parameter less than 1, which is common for faults in this research. An alternative approach is to try every combination of parameters and evaluate which combination best matches the expected behaviour of T2. This can be applied to any distribution type but is very time-consuming, hence it was adapted using an optimisation technique that could consider a smaller amount of parameter combinations. This was achieved by creating a Petri net with five places and three transitions, described below in Figure 5.13:

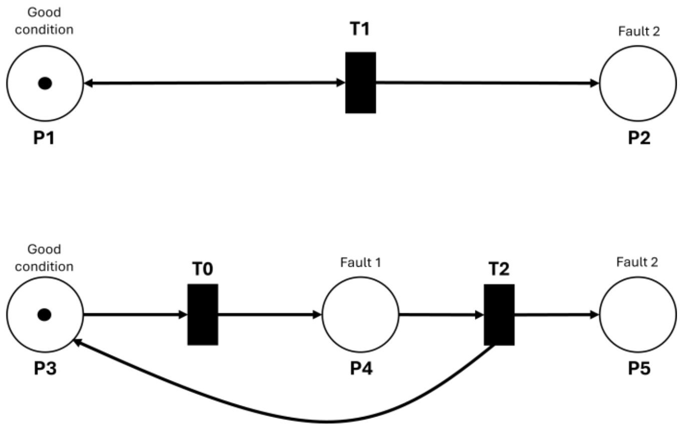

Figure 5.13: Petri net structure for the optimisation technique

T0 and T1 are known distributions. When fired, T1 places a token in P1 and P2, and T2 sends a token to P3 and P5, hence the net will run continuously. The optimisation technique is based on the genetic algorithm (Kramer, 2017), and is completed as follows:

1) A range of parameters is decided for the appropriate distribution used in transition T2
2) 20 parameter combinations are randomly chosen from the ranges
3) The Petri net is simulated for a lifetime of specified length, after which the total number of tokens present in P2 and P5 is noted. The net is reset to its original condition
4) Step 3 is repeated 100 times to find the mean absolute error (MAE) for the given parameters for the distribution of T2:

$$MAE = \frac{\sum_{i=1}^{i=100} |P2_i - P5_i|}{100} \quad (5.3)$$

Where $P2_i$ and $P5_i$ are the number of tokens present in P2 and P5 at the end of the simulation i

5) The MAE for the 20 parameter combinations is compared and ranked from lowest to highest. The top ten parameter combinations with the lowest MAE scores are retained and used in the next generation.
6) The ten lowest combinations are removed and replaced with a parameter combination made from randomly selecting between the parameters of the top ten combinations.
7) For parameter combinations ranked between 5-20, there is a 5% chance that a parameter will be removed and replaced by a random selection from the original

ranges determined in Step 1. This is to ensure that convergence is not reached too soon so that an optimal combination of parameters can be chosen.

8) Steps 3 to 7 are repeated using the new generation until the results have converged. Convergence is defined as when the highest ranked combination of parameters has not changed for 5 generations.

The optimisation technique is a practical and relatively quick solution to obtaining the parameters needed in transition T2 and will allow for the rate of degradation between faults to be analysed and utilised in the developed Petri nets.

5.3.3 Convergence

The model is simulated for the length of a single lifetime of a track asset, until a renewal is required. Statistics are found from repeated simulations of the model with random sampling via Monte Carlo simulation. The more simulations that are completed, the more accurate the results become but the time taken to run the models, and the amount of memory required, increases. Instead, a convergence threshold was set, which was defined as the point at which there is less than a 1% difference between the results found per simulation. This was always achieved when there were 1500 simulations of the model.

5.3.4 Renewal parameters

The developed PN can predict the expected lifetime of a track asset with 31 different renewal strategies (Equation 5.1). First however, the thresholds for renewal must be defined. The parameters for each renewal criterion were derived from a combination of available information in the railway standards, extrapolating based on observations of the historical failure data, and trial and error of various parameters to ensure there were no wildly unrealistic predicted lifetimes. The applied parameters are explained in Table 5.8.

Asset	Maintenance Limit	Set Period	Time Duration	Associated Cost	Condition Based
Crossing panel	7 faults	2 faults within a 2-year review period	3 consecutive faults within 2.5 years of the previous fault	2500 units spent on operations	8mm material removed by grinding or 10% weld coverage
Switch panel	11 faults	2 faults within a 1-year review period	3 consecutive faults within 1.5 years of the previous fault	3000 units spent on operations	8mm material removed by grinding or 10% weld coverage
Plain line rail	2 faults	2 faults within a 5-year review period	3 consecutive faults within 7 years of the previous fault	1500 units spent on operations	14mm of the rail removed by grinding or 10% weld coverage
Ballast	9 faults	3 faults within a 6-year review period	3 consecutive faults within 3 years of the previous fault	2000 units spent on operations	30% of the aggregate has been damaged

Table 5.8: Implemented renewal parameters

5.4 Application of models

Once built, populated with distributions found from an analysis of historical data, and asset management parameters implemented, the models are repeatedly simulated. Results are gained by observing the movement of tokens around the net and recording the number of places visited or how long a token resided in various states. Given that the renewal parameters were estimated (Table 5.8), the predicted lifespans are not directly discussed as these will vary depending on the thresholds set for renewal. Instead, the impact of the renewal strategy on the range of expected lifetimes, how the predicted lifespan on one asset compares to another given identical renewal parameters, what maintenance actions are expected across a lifespan and how the track's condition will vary, as well as any other examples of where the developed PN methodology can improve on known limitations of current track asset models and increase comprehension of the renewal process, has been included in the discussion.

There is no limit on how large a lifetime can be during a simulation, as a renewal is only scheduled once a relevant threshold has been met or as part of the maintenance requirements for a fault. Because of this, several lifetime predictions were exceptionally large (>500 years). This is unrealistic, and including these in the analysis was found to distort the results. Anomaly results were identified by any lifetime lying outside the interquartile range fences, defined as 1.5 times the interquartile range of the lifetimes added to the upper quartile value, or subtracted from the lower quartile value, and then removed from the study before lifetime predictions were presented. All lifetimes discussed have been pre-processed in this way.

5.4.1 Crossing panel

5.4.1.1 Validation

Before any results are discussed, it is important to first establish if the developed model is behaving similarly to what was observed in the fault history used to populate it, as this can be the first indication of an issue with the implementation of the model or suggest that the PN structure is unable to accurately represent the asset management strategy. Given that the observation period of the railway data did not extend across a crossing panel's entire lifecycle and no significant renewals had taken place, it is difficult to validate any lifespan predictions made from the model. Instead, the PN was simulated for a total of 4223 days (equivalent to the span of the rail defect history) with no renewal strategy implemented, and the behaviour from the simulation is compared to what was observed in the HS1 data.

Average no. actions	Observed	Simulation
Faults	3.34	2.93
Grinds	0.75	0.78
Welds	2.58	2.16

Table 5.9: Average number of faults, grinding actions, and welding actions occuring on a crossing panel unit for the observed data and simulation results

Table 5.9 shows that the modelled crossing panel is degrading at a similar rate to those present on the HS1 line, with an accurate division between performed welding and grinding actions. By comparing the contribution of each failure type to the total amount of faults experienced on a crossing panel (Table 5.10), it is noted that the simulated model also degrades due to similar failures compared to how the observed track was degrading.

Failure type	Observed	Simulation
Lipping on a casting	24%	22.09%
Squat on a casting	34.45%	31.27%
Transverse defect from RCF	15.87%	16.35%
Flaking or shelling at wheel contact surface	11.06%	11.18%
Wear in the wheel transfer area	4.59%	7.78%

Table 5.10: Contribution of each failure type to the faults on the crossing panel for the observed data and simulated results

The final validation is to ensure that the results of the simulations have converged. If no convergence is reached, the predicted results from the model are not reliable as a single solution could not be found. Plotting the average lifespan of the crossing panel over the number of simulations, for a renewal strategy considering each criterion in isolation, confirm that convergence was always reached after 1500 simulations (Figure 5.14), and as such the PN was concluded to model the asset management behaviour of a crossing panel well. Any results found are reflective of the data used to populate the model.

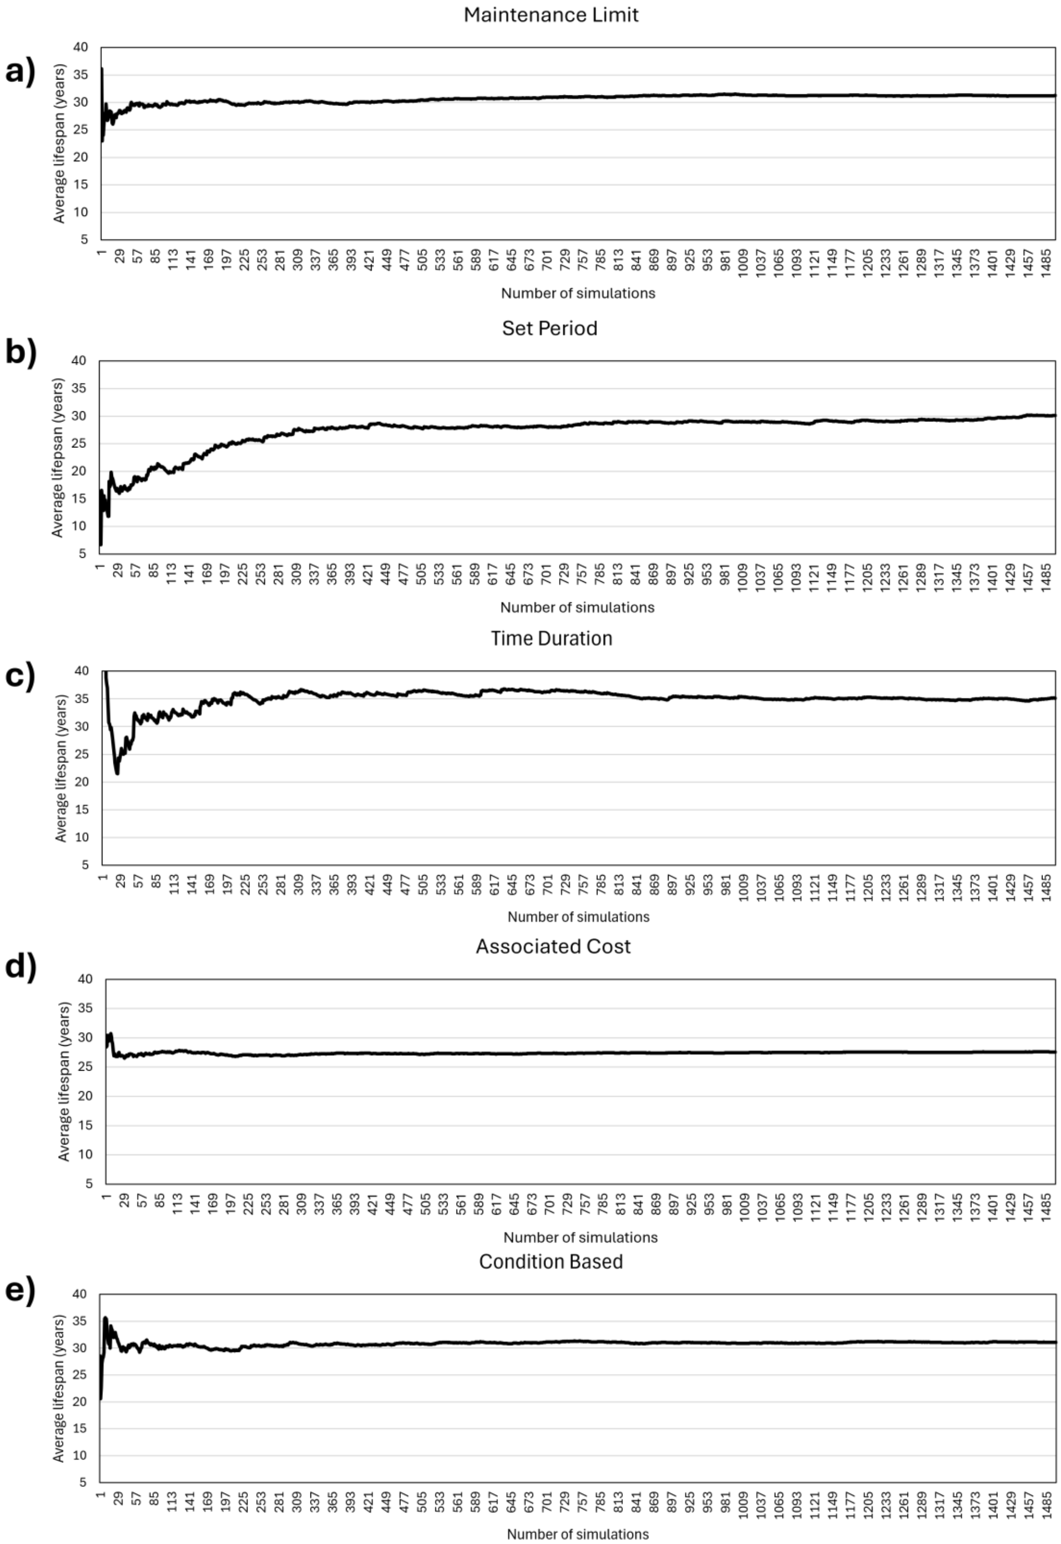

Figure 5.14: Convergence of the mean predicted lifespan for a crossing panel unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation

5.4.1.2 Expected lifespan of the crossing panel

The average expected lifespan from the models described in Figure 5.14 revealed that limit-based strategies such as Maintenance Limit or Associated Cost had a much smaller range of predicted lifespans compared to those that evaluate when maintenance is completed (Table 5.11).

Renewal strategy	ML	SP	TD	AC	CB
Mean [SD] lifespan (years)	30.41 [12.56]	23.93 [21.98]	25.49 [22.24]	27.48 [7.06]	29.17 [14.49]

Table 5.11: Average expected lifespan of a crossing panel unit with each renewal criterion considered in isolation for the renewal strategy

The comparatively high SD for the Set Period renewal strategy implies that 95% of crossing panel units are expected to require a renewal between -20.03 and 67.89 years of operation. As negative times are impossible, this prompted a further investigation. Plotting the distribution of the predicted lifespans from 1500 simulations of the crossing panel PN with the SP criterion considered in the renewal strategy (Figure 5.15) revealed a very large right-skew to the predictions, with the majority of renewals scheduled before 10 years of operation. This was most likely due to the renewal threshold of 2 faults within a 2-year control period being met early because of the 'settling-in' behaviour of crossing panels that was found during the data analysis. The size of this skew influenced the large standard deviation seen in Table 5.11.

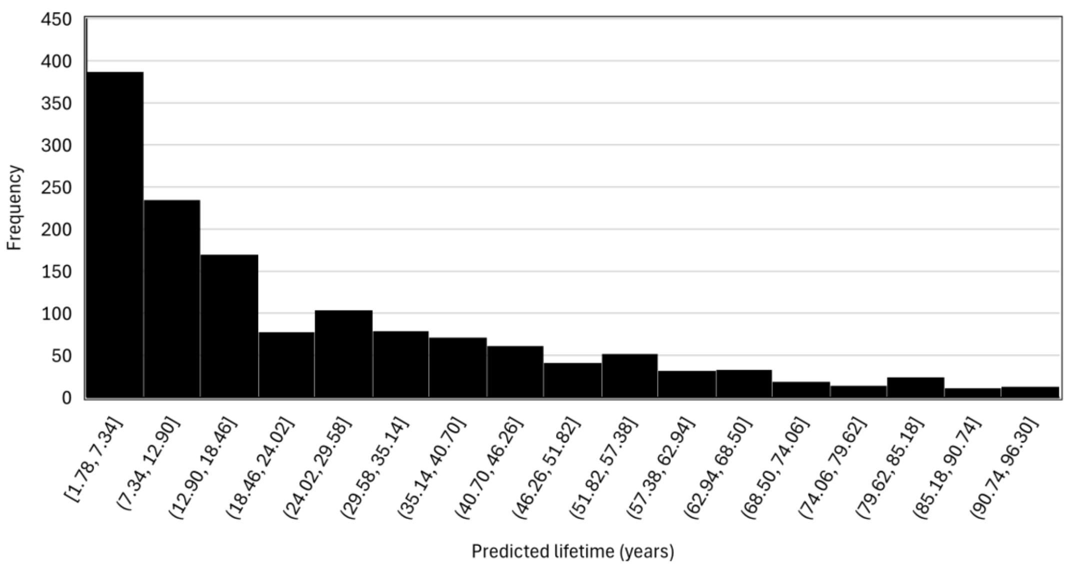

Figure 5.15: Distribution of the predicted lifetimes from simulating the crossing panel PN with the SP criterion considered in the renewal strategy

If the SP criteria with the given parameters is the only factor considered within the renewal strategy, many crossing panel units will be replaced very early into operation, but as Figure 5.15 shows, some can be expected to last for almost 100 years.

5.4.1.3 Expected condition of the crossing panel

The developed PN can be adapted to terminate the simulation when a fixed time has passed, instead of when a renewal is required. From this, statistics can be collected regarding the condition and behaviour of a crossing panel over time. This is achieved by adding two loops in the PN structure to 1) reset the entire model back to its initial condition when needed and 2) regularly evaluate the location of tokens of the net. As the degradation of the track is not expected to change between different renewal strategies, the following analysis will be implemented with only the Associated Cost criterion considered in the renewal strategy as it had the lowest variation between predicted lifespans.

Simulating the model over a 40-year observation period and monitoring the movement of tokens into the maintenance places can find the expected type and frequency of maintenance actions completed on the crossing panel. From Table 5.12 it can clearly be seen that welding far outweighs grinding to rectify faults, which is consistent with what was observed from the historical maintenance data. Additionally, the crossing panels of the HS1 line are expected to spend approximately 1 day under speed restrictions over 40 years of operation, making their impact negligible in budget planning. As restrictions are only implemented when the crossing panel has degraded to a severe condition, this suggests that the current maintenance strategy is incredibly effective at maintaining the condition of the track.

Action	**Renewal**	**Grinding**	**Welding**	**Speed restrictions (daily)**	**Appoint watchman**	**Inspection**
Expected no. over 40-year observation	1.02	2.83	6.49	0.99	0.15	102.30

Table 5.12: Expected maintenance actions completed on a crossing panel unit with the AC criterion considered in the renewal strategy over a 40-year observation period

The 40-year observation period can also be used as a budget planning tool for track engineers, by evaluating the average cost spent per control period (CP), defined as a 5-year period starting from installation of the line. The expenditure was calculated by multiplying the cost per action by the expected number of actions completed during each control period, which required the assumption that a replacement of a crossing panel is equivalent to 2000 units. Figure 5.16 shows that renewal of the crossing panel typically occurs between the 5th and 8th control period, which correlates with a higher overall expenditure during those times. The evaluation also found that the cost of inspecting the asset was consistently higher for each control period than the cost of maintaining via grinding or welding.

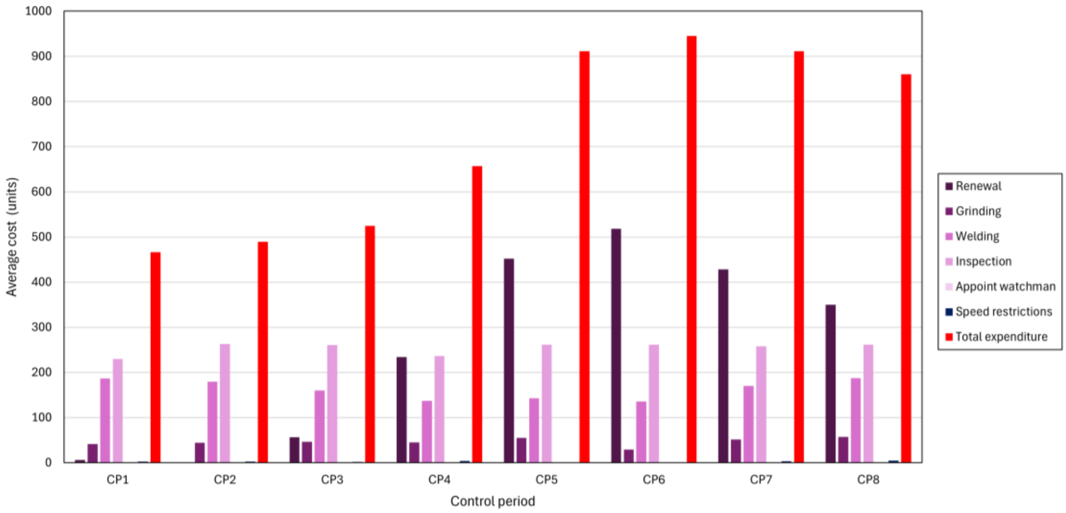

Figure 5.16: Average expenditure on the crossing panel per control period with the AC criterion considered in the renewal strategy

Observing the condition of the crossing panel over time can be used to visually represent the occurrence of faults, expanding on the findings from the reliability analysis by demonstrating how the faults occur concurrently on an asset instead of in isolation. In Figure 5.17, faults have been grouped by their failure type, with failures that occurred in less than 5% of the simulations grouped into 'other'.

It was observed that a crossing panel is expected to be in a good condition no less than 65% of the time at any point during its operational period. This again is reflective of the very effective maintenance strategy currently implemented. A total of 21 failure types were best described by an exponential distribution for the crossing panel, including lipping, which can explain why the likelihood of different failures remains mostly consistent throughout the 40 years. However, the impact of early defects due to squatting and transverse defects from RCF - both modelled by a lognormal distribution with a large scale parameter- is shown, as the crossing panel is in its worst condition shortly after installation. After this, the condition gradually increases until around 22 years, when the units are expected to start requiring renewal, and the probability of being in a good condition slightly decreases as 'settling in' defects are again more common.

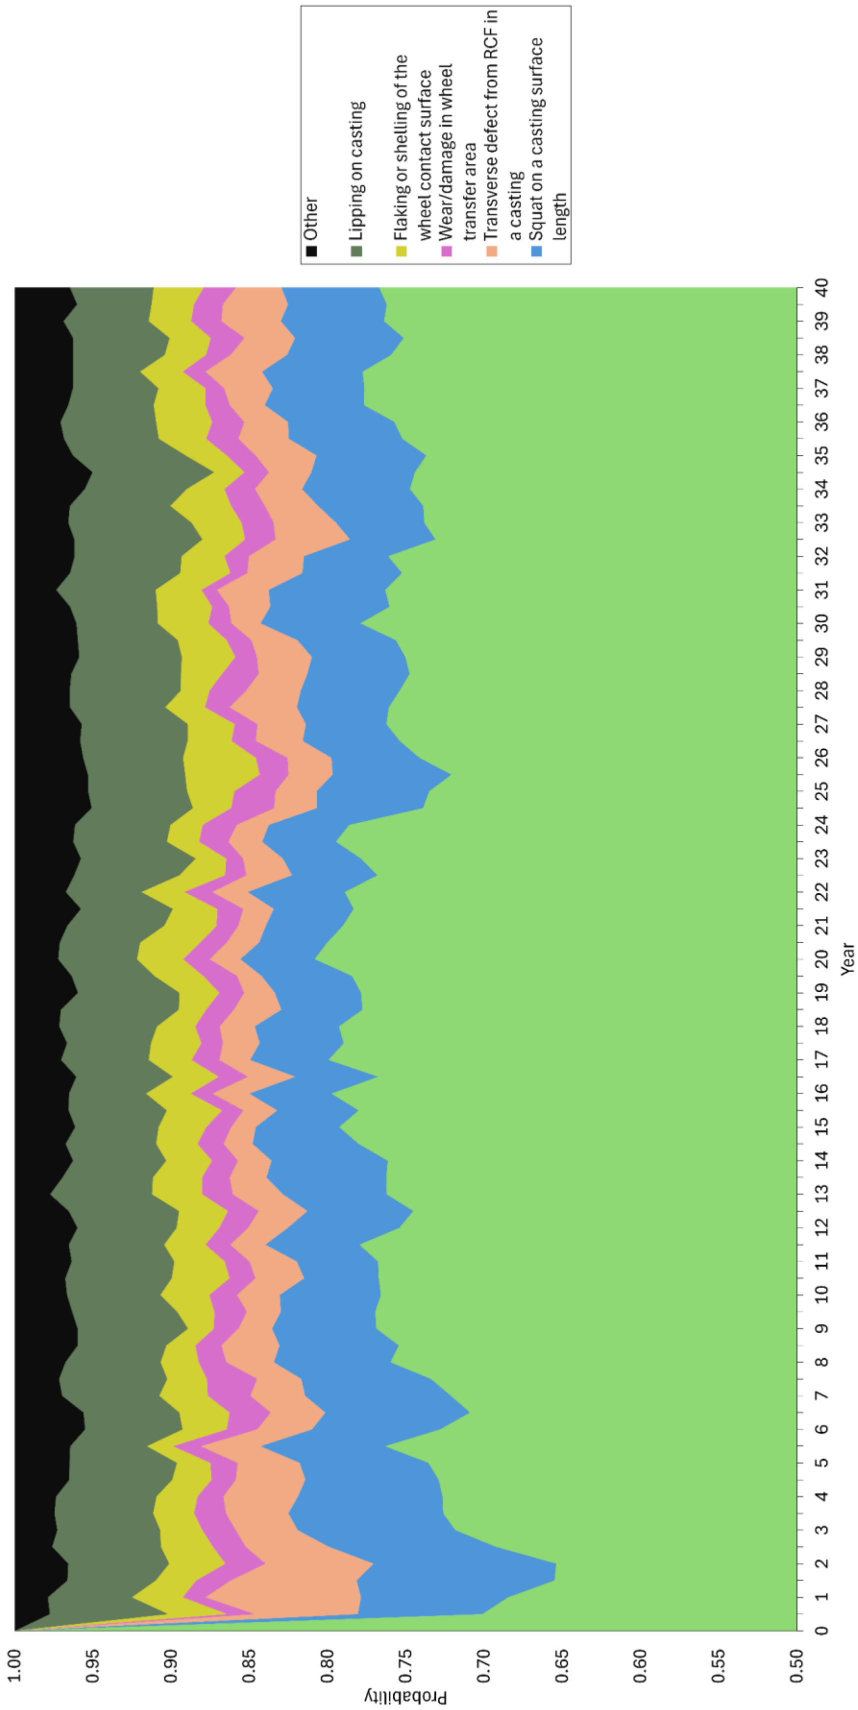

Figure 5.17: Average condition of a crossing panel unit over a 40-year observation period with the AC criterion considered in the renewal strategy

5.4.1.4 Including multiple criteria in the renewal strategy

The largest limitation highlighted in the literature review for existing asset management models for railway assets was that the majority considered only a fixed lifetime as a renewal strategy, and none allowed for choice between, or a combination of, different measurements of wear. The considered renewal criteria for this research, and how they are involved in the renewal strategy by strategic placing of tokens, is shown in Figure 5.9 - in total there are 31 different strategies that could be implemented. The results discussed so far have been based on the consideration of one criterion at a time; when multiple are selected, the criteria that reaches its renewal threshold first fires its associated transition to move a token to the 'renewal scheduled' place and resets the net to its initial markings, hence the lowest possible lifetime is always modelled.

For the crossing panel, including all 5 of the considered criteria in a single strategy will result in an unrealistically low mean lifespan of 16.66 years, but it can be used to demonstrate how the unit is wearing by considering which criteria is the most influential for triggering renewals. Figure 5.18 shows that under the current parameters, the Set Period renewal threshold triggered 34% of all replacements, whereas the crossing panel unit was rarely permitted to degrade long enough for 7 faults to occur, given that Maintenance Limit only contributed to 4% of renewals during the simulations. 'Failures' refers to all renewals that were due to a severe fault that couldn't be rectified through maintenance, as determined in the standards.

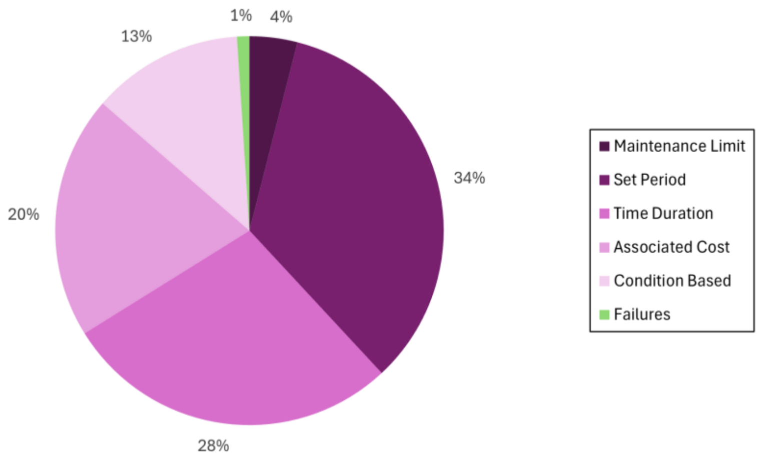

Figure 5.18: Distribution of criteria triggering a renewal of the crossing panel, for a renewal strategy considering all five criteria

In Figure 5.15 it was shown that when only the Set Period criterion was considered in the renewal strategy, there was a large right skew to the predicted lifetimes and 15% of crossing panel units are expected to still be operational after 50 years of service.

Considering the average rate of degradation was 1 fault every 4 years, it is likely that the asset would be worn due to excessive grinding or high coverage of weld repairs resulting from maintenance across the 50 years of operation. This is undesirable as it may leave the unit more susceptible to serious faults and decreases ride quality for the passengers. However, by also including the Condition Based criterion in the renewal strategy, the modelled crossing panel would be removed if its condition was undesirable, totalling 34% of all replacements (Figure 5.19). Plotting the updated distribution of predicted lifespans from 1500 simulations of the crossing panel in Figure 5.20 demonstrated that the inclusion of CB caused the skew of the lifetime predictions to be less pronounced, and found that only 4.3% of units were now expected to still be in operation after 50 years.

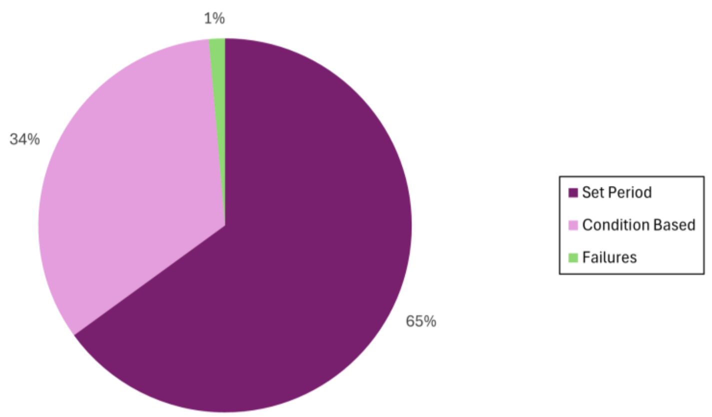

Figure 5.19: Distribution of criteria triggering a renewal of the crossing panel for a renewal strategy considering the SP and CB criteria

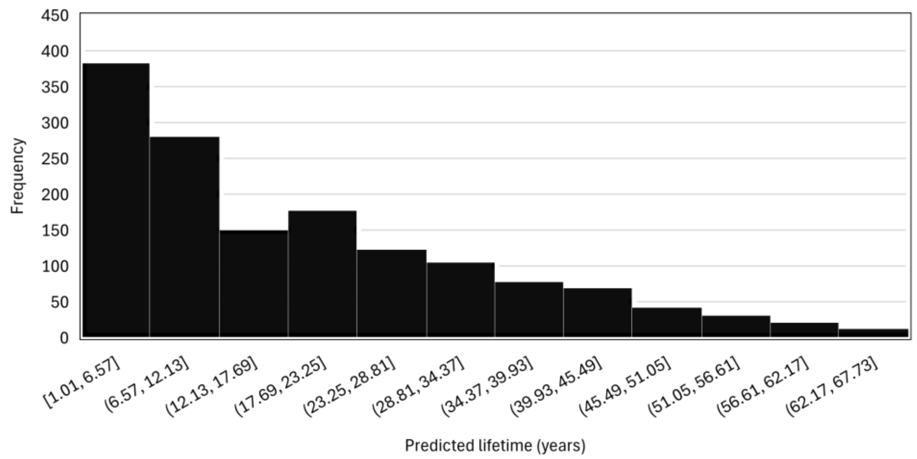

Figure 5.20: Distribution of the predicted lifetimes from simulating the crossing panel PN with the SP and CB criteria considered in the renewal strategy

Naturally, an analysis on all possible combinations of renewal criteria is not possible however the benefits of considering multiple at once has been demonstrated.

5.4.2 Switch panel

5.4.2.1 Validation

Completing the same validation checks described in 5.4.1.1 on the switch panel confirmed that the developed PN was an appropriate modelling type and structure, as the model degraded at a similar rate and due to similar behaviour as to what was seen in the observed data (Tables 5.13 and 5.14). Only the five most common failure types were compared during the validation, as the rest combined were responsible for 3.22% of the failures, hence their impact in isolation is negligible.

Average no. actions	Observed	Simulation
Faults	5.22	4.39
Grinds	4.24	3.51
Welds	0.98	0.91

Table 5.13: Average number of faults, grinding actions, and welding actions occuring on a switch panel unit for the observed data and simulation results

Failure type	Observed	Simulation
Lipping in the machined length of the switch or stock rail	68.27%	69.50%
Gauge 2 contacts below any damage or crack	17.4%	14.10%
Shallow flange contact angle	6.16%	6.60%
Transverse cracking in RCF	3.88%	5.17%
Switch rail height relative to the stock rail	1.07%	1.39%

Table 5.14: Contribution of each failure type to the faults on the switch panel for the observed data and simulated results

The slight discrepancy between number of faults observed compared to what was predicted by the simulation is thought to be a result of the assumption that maintenance always takes the maximum allotted time to complete. As Table 5.14 shows, 70% of all faults can be attributed to lipping in the machined length of the switch or stock rail. This has a maximum permitted time of 182 days to remove the fault by grinding, but observing the maintenance history revealed this correction was typically completed in under 90 days. It is not thought that this assumption will impact the renewal predictions greatly, but it can be rectified if the firing times to complete maintenance are changed to better represent engineers' behaviour. Such analysis was beyond the scope of this research.

Figure 5.21 confirms that the average lifespan found from simulating the PN with every renewal criterion considered separately always converged after 1500 simulations, hence the developed PN was capable of accurately predicting the behaviour of a switch panel.

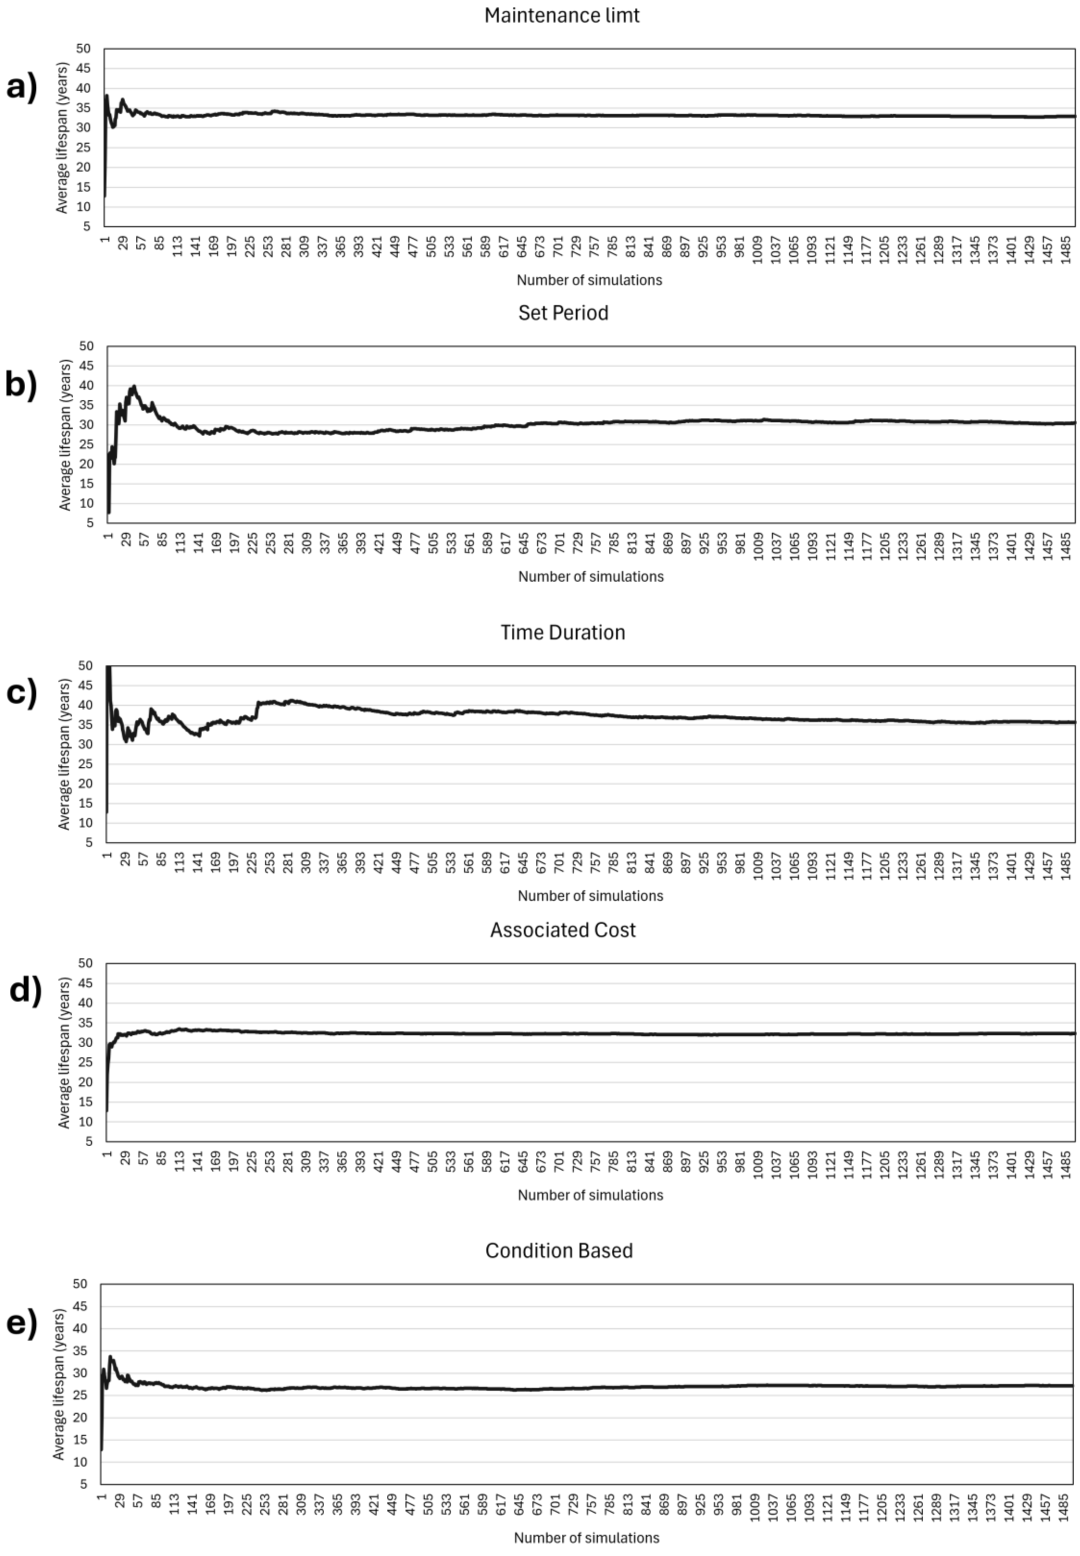

Figure 5.21: Convergence of the mean predicted lifespan for a switch panel unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation

5.4.2.2 Expected lifespan of the switch panel

Simulating the Petri net with the assumed renewal parameters has predicted lifespans between 20-35 years when considering each renewal criterion separately for switch panels on the HS1 line (Table 5.15).

Renewal strategy	**ML**	**SP**	**TD**	**AC**	**CB**
Mean [SD] lifespan (years)	30.90 [10.91]	24.95 [22.57]	27.97 [24.98]	32.30 [6.87]	25.77 [10.60]

Table 5.15: Average expected lifespan of a switch panel unit with each renewal criterion considered in isolation for the renewal strategy

Again, it was found that the results using the Associated Cost criterion had a low SD compared to the mean lifespan, due to little deviation between the predicted values. This behaviour was also seen in Figure 5.21, as the average lifespan converged after only 150 simulations. AC – a limit-based renewal strategy - considers the cost of all operations, including visual inspections which occur at a minimum frequency of every 6 months. This effectively places an upper bound on the maximum lifespan of the unit which the other renewal polices do not have, contributing to the small SD seen in Table 5.15.

The convergence plots in Figure 5.21 also verify the large SD seen with the SP and TD renewal criteria, as neither truly converge until at least 600 simulations. This is more clearly seen in Figure 5.22 which provides a magnified view of the SP, TD, and AC convergence plots in Figure 5.21 over the first 600 simulations. As the figure shows, the average lifespan of the switch panel renewed under the SP and TD criteria still varies slightly at 600 simulations compared to the AC criterion which converged early, resulting in a smoother plot. This behaviour was again due to the early 'settling in' failures on the switch panels causing the renewal thresholds of SP and TD to be met early, with a large right skew to the predictions. The large standard deviations suggest that these criteria are currently unsuitable to be considered in isolation for planning future renewal interventions as there is a high-level of uncertainty, however this may simply be reflective of the short observation period for the data populating the models. Distributions derived from switch panels with a longer observation period that covers a unit's entire lifespan are more likely to indicate wear-out, which is tailored to a frequency-based renewal criteria such as TD.

a) b) c)

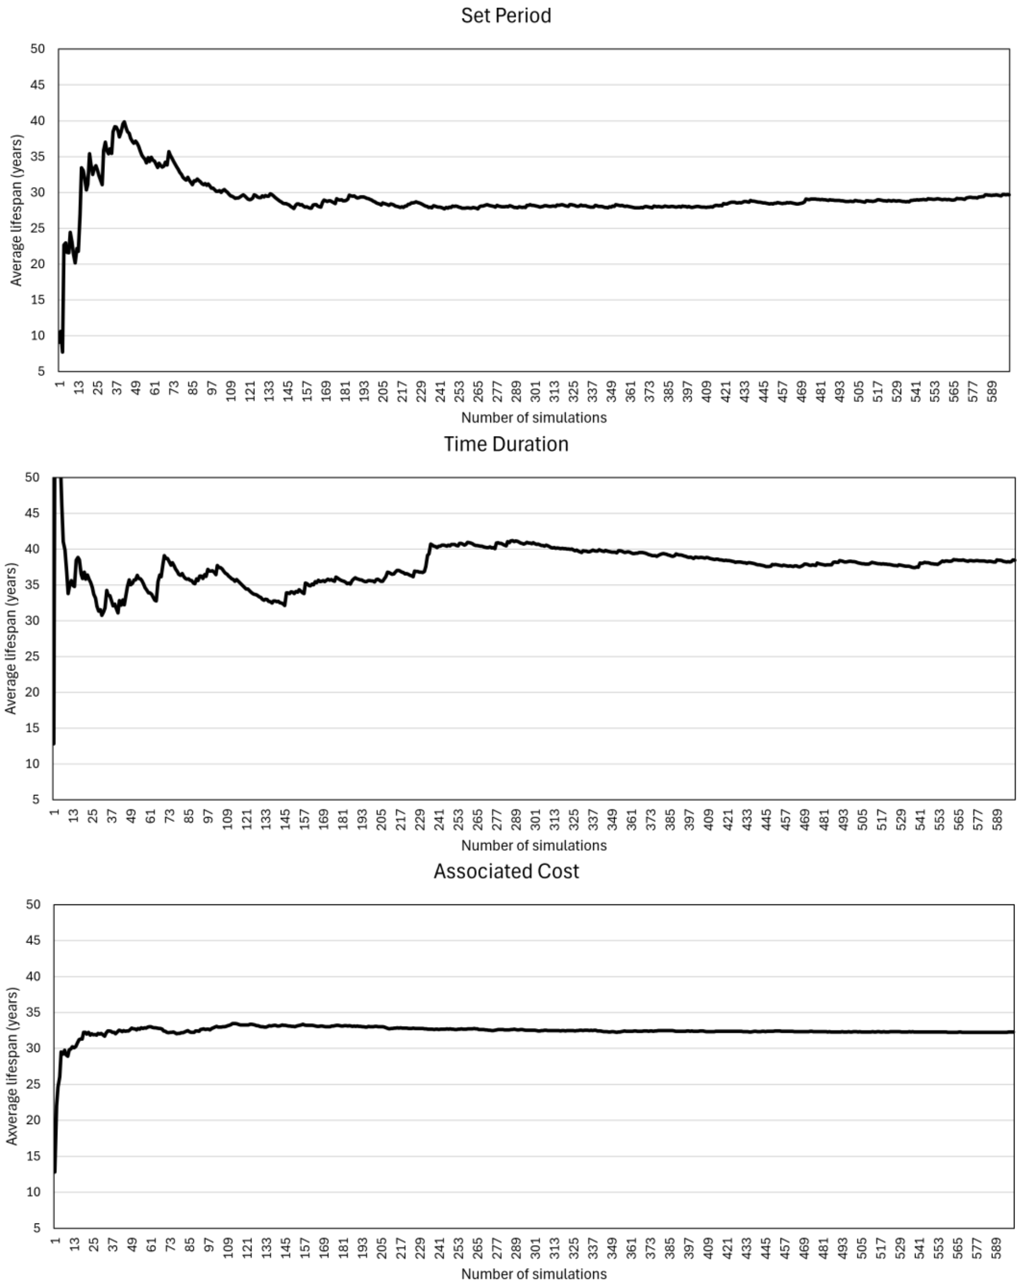

Figure 5.22: Magnified plots of convergence of the mean predicted lifespan for a switch panel unit with a renewal strategy considering the a) SP b) TD c) AC criterion in isolation

5.4.2.3 Expected condition of the switch panel

Implementing solely the AC renewal strategy and terminating the simulation of the PN after 40 years allowed for the average condition of the switch panel over time to be seen (Figure 5.23). The frequent variation of the condition is thought to result from the observations occurring at 6-month intervals.

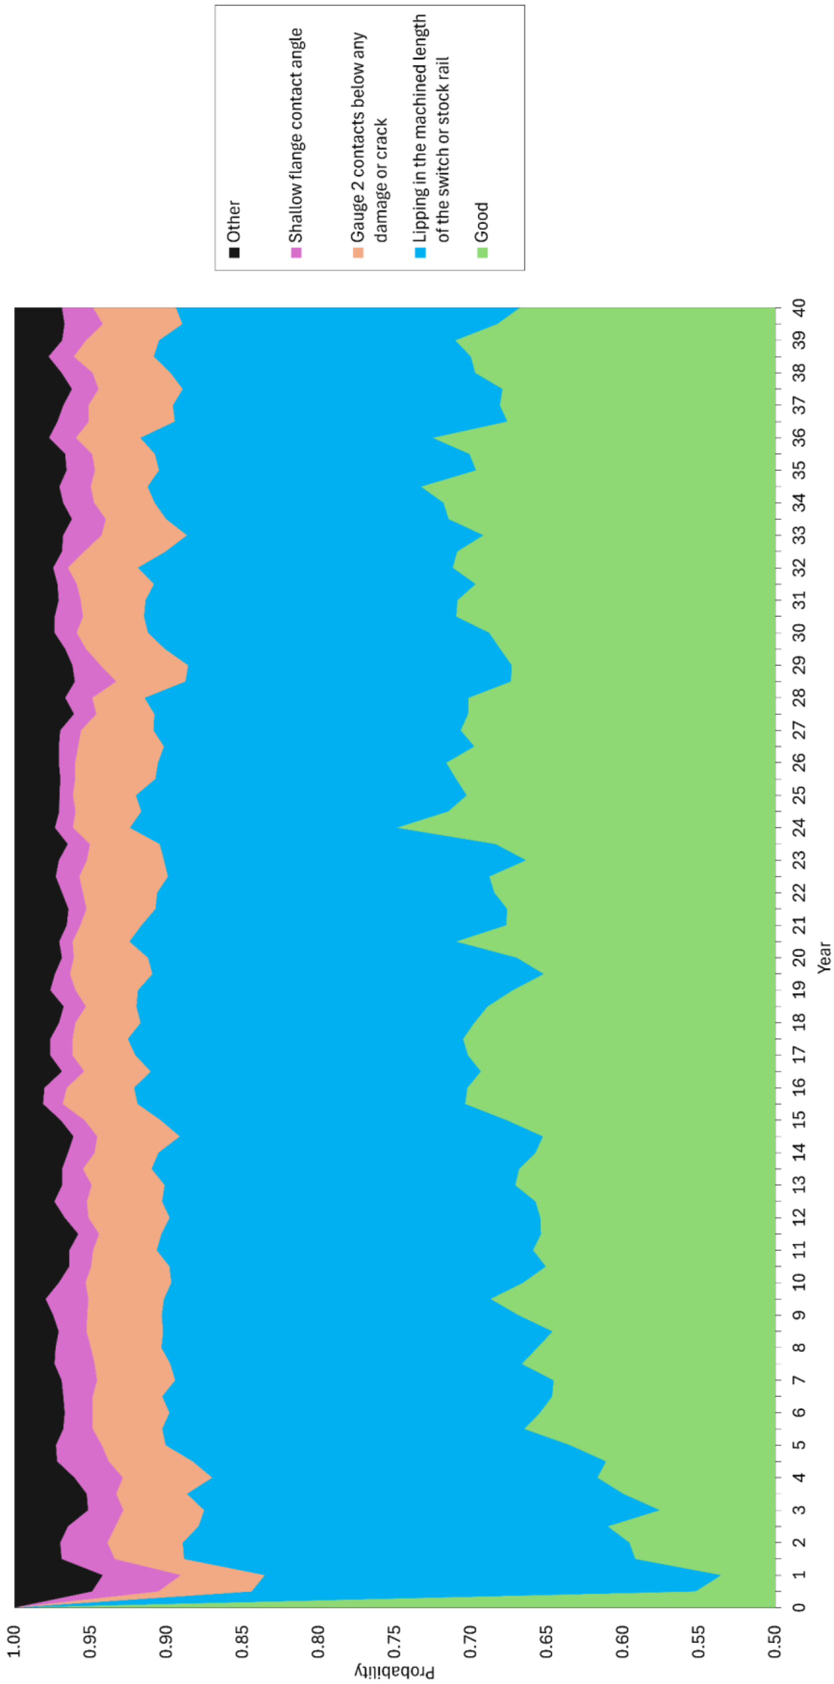

Figure 5.23: Average condition of a switch panel unit over a 40-year observation period with the AC criterion considered in the renewal strategy

Figure 5.23 confirms that lipping is overwhelmingly the most common failure type on a switch panel as it is expected to be present a minimum of 17% of the time at any point over the 40-years. The 'settling in' failures discussed during the analysis of the rail defect history can be seen, as within 1 year of installation there is a 45% chance that the switch panel has experienced at least 1 fault. After the initial three years of operation, the condition slightly improves as the asset enters its useful lifetime, then remains moderately consistent throughout. There was little-to-no evidence for the impact of renewal on the asset's condition observed, which was expected at around 32 years.

Collecting statistics on the average number of corrective and preventative maintenance actions across the 40 years shows that grinding the switch panel is used frequently to rectify its condition and prolong the expected lifespan (Table 5.16). This is again reflective of the high number of lipping faults experienced on a switch panel which can only be corrected by removing a thin layer of the surface of the rail and may suggest a need for more wear resistant materials to prolong the unit's predicted life. In Grossoni et al.'s (2021) study into degraded S&C assets, the benefits of redesigning the S&C, such as with a thicker switch blade, was discussed to combat the effect of wear; the findings in Figure 5.23 and Table 5.16 suggest this may be of benefit for HS1 also.

Action	**Renewal**	**Grinding**	**Welding**	**Ban to-facing moves (daily)**	**Lubricate rail**	**Inspection**
Expected no. over 40-year observation	0.89	11.03	3.06	0.26	0.74	87.64

Table 5.16: Expected maintenance actions completed on a switch panel unit with the AC criterion considered in the renewal strategy over a 40-year observation period

Figure 5.23 has shown the expected condition of the switch panel over time but has been grouped per failure type for easier comprehension. The structure of the degradation sub-model means each individual fault within the failure type is modelled, which allows for the asset behaviour to be examined in more detail. Figure 5.24 describes the convergence for the probability of each fault of the 'shallow flange contact angle' failure type, a measurement of wear on the rail head, occurring over a 40-year simulation. It's clear that with the data currently populating the model, a shallow flange contact angle of less than 60 degrees is expected to occur no more than once during this time, but there is almost no chance of this degrading further to below 55 degrees. This reflects what was seen in the historical failure data, as there was only 1 account of a contact angle for the switch rail recorded at below 55 degrees across 143 units.

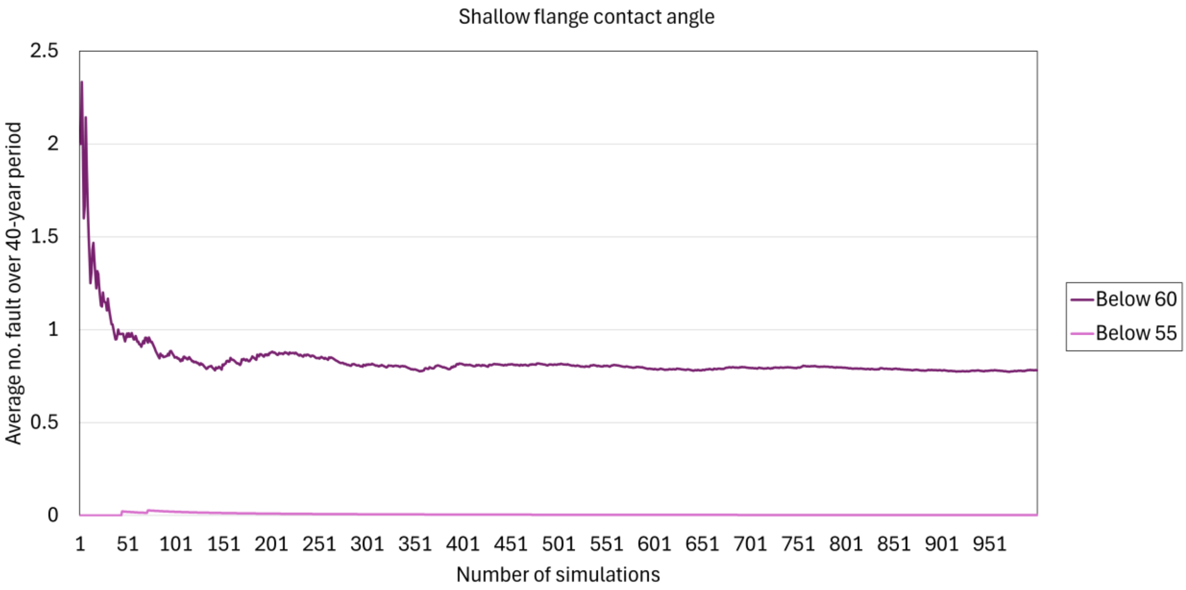

Figure 5.24: Expected no. faults for the shallow flange contact angle for a 40-year observation of a switch panel unit

5.4.2.4 Comparing the expected lifespan of multiple assets

There are very few examples present in the literature focusing on modelling the asset management of an S&C. Litherland and Andrews' PN model (2023b) is currently the most advanced and has been demonstrated to adequately investigate the maintenance procedure, however every component (including the switch blade and crossing nose) was assumed to require a renewal at the same time. Current policy meanwhile is to consider the crossing panel and switch panel as two separate assets, that are replaced independently.

Renewal is an expensive and disruptive process, that requires the line to be blocked whilst the new asset is installed. Replacing two major components of the S&C at the same time would reduce the amount of time the line is unusable, helping to provide a more efficient service. However, it was found that simulating the switch and crossing panel PN with identical renewal parameters would typically result in the switch panel having a much shorter average lifespan, and if this were to trigger the renewal of both assets, the crossing panel would be replaced before it could experience the same degree of wear. The results of Table 5.17 demonstrate that it is correct to treat the two components separately as it maximises the lifespan from both, reducing costs.

	ML	SP	TD	AC	CB
Renewal parameters	6 faults	2 faults within a 2-year review period	3 consecutive faults within 2.5 years of the previous fault	2000 units spent on operations	8mm material removed by grinding or 10% weld coverage
Mean [SD] lifespan switch panel (years)	16.21 [7.00]	11.41 [9.00]	12.59 [10.13]	21.20 [5.32]	25.77 [10.60]
Mean [SD] lifespan crossing panel (years)	25.70 [10.97]	23.93 [21.98]	35.95 [33.47]	22.49 [6.08]	29.17 [14.49]

Table 5.17: Average expected lifespan of the switch and crossing panel unit for identical renewal criterion considered in isolation for the renewal strategy

There is a high disparity between the average lifespan for limit-based renewal strategies – the expected lifetime of the crossing panel is almost 10 years longer than the switch panel for a renewal strategy that replaces the assets when 6 faults have occurred. However, if both S&C components are held to an identical cost threshold, there is just over 1 year difference in the predicted lifespans. This behaviour is explained when simulating both assets over a 30-year period with the AC renewal strategy implemented and observing how maintenance is completed (Table 5.18). The crossing panel experiences a higher number of welds, which are more expensive to complete compared to grinding, and will require a larger number of inspections. The minimum number of inspections over the 30-year period is 60 as they are completed biannually, but the results of Table 5.18 indicate that both components will regularly require extra inspections as a preventative maintenance intervention for a noticed fault.

Asset	Renewal	Appoint watchman	Grind	Weld	Impose speed limits	Ban to-facing moves	Lubricate	Inspect
Switch panel	0.95	N/A	8.31	2.30	N/A	0.13	0.57	65.67
Crossing panel	0.92	0.10	2.09	4.81	0.65	N/A	N/A	75.05

Table 5.18: Expected maintenance actions completed on a switch and crossing panel units with the AC criterion considered in the renewal strategy over a 30-year observation period

5.4.3 Plain line rail

5.4.3.1 Validation

Simulating the plain line rail model for the first 4223 days of operation with no renewal strategy implemented (achieved by inhibiting all transitions leading to renewal via strategic placing of tokens) showed that the simulated PN displayed similar behaviour to what was observed from the rail defect history. For example, failure on the HS1 units were extremely rare, with around 90% having not experienced a single fault, which was reflected in the low average number of failures for the PN simulation (Table 5.19).

Average no. actions	Observed	Simulation
Faults	0.10	0.08
Grinds	0.07	0.06
Welds	0.04	0.02

Table 5.19: Average number of faults, grinding actions, and welding actions occuring on a plain line rail unit for the observed data and simulation results

Following similar validations for the switch and crossing panel (Tables 5.10 and 5.14), the ability of the developed PN to model the degradation of the plain line rail was investigated, by considering the proportion of the 5 most common failures for the observed data and simulation results (Table 5.20). Simulating the PN 1500 times found 'Damage on the running surface' to be three times as common as what was observed on the HS1 line. Whilst the models are not expected to perfectly mimic the observed data, a discrepancy of this size prompted an initial review into the PN structure and transition firing times, which were all confirmed to be implemented correctly.

Instead, it was theorised that this could have resulted from the rare behaviour of the failure type (which had less than 30 recorded faults across 1739 units in the observed data), as the unpredictable nature of rare events means their true frequency is unlikely to be observed if the number of simulations is too low. This was explored by simulating the PN 3000 times and comparing results, finding that the proportion of experienced failures has converged slightly toward what was seen on the actual HS1 line compared to only 1500 simulations (Table 5.20). This is particularly evident for 'Damage on the running surface', for which increasing the number of simulations in the analysis has reduced this failure type from contributing to 15% of all faults experienced on the plain line rail to just 8.33%, far more in line with the 5.95% share observed on the HS1 line.

Following this, it is concluded that the PN can adequately model the degradation of a segment of plain line rail, with the accuracy of any results improving with a higher number of simulations. However, a higher number of simulations will increase the memory allocation and computational time required for the analysis.

Failure type	Observed	Simulation (1500)	Simulation (3000)
Defective aluminothermic weld	44.32%	38.64%	39.29%
Wheelburns	18.91%	15%	19.05%
Transverse defect from a tache ovale, squat, wheelburn, or weld repair	7.56%	15%	11.90%
Damage on the running surface	5.95%	15%	8.33%
Untestable rail	5.95%	2.27%	3.57%

Table 5.20: Contribution of each failure type to the faults on the plain line rail for the observed data and simulated results

Plotting the average lifespan of a plain line rail unit over 1500 simulations for various renewal strategies (Figure 5.25) indicated that any lifespan predictions obtained from the PN model were reliable, as convergence at the 1% level was always reached. Hence, with the validation complete it can be concluded that the developed model was a sufficient representation of the asset management strategy for the plain line rail, and statistics regarding the expected lifespan and overall condition can be collected. Given the discussed issues with higher simulation numbers, and the large amount of analysis completed over this section, all further results are reflective of a PN that was simulated 1500 times.

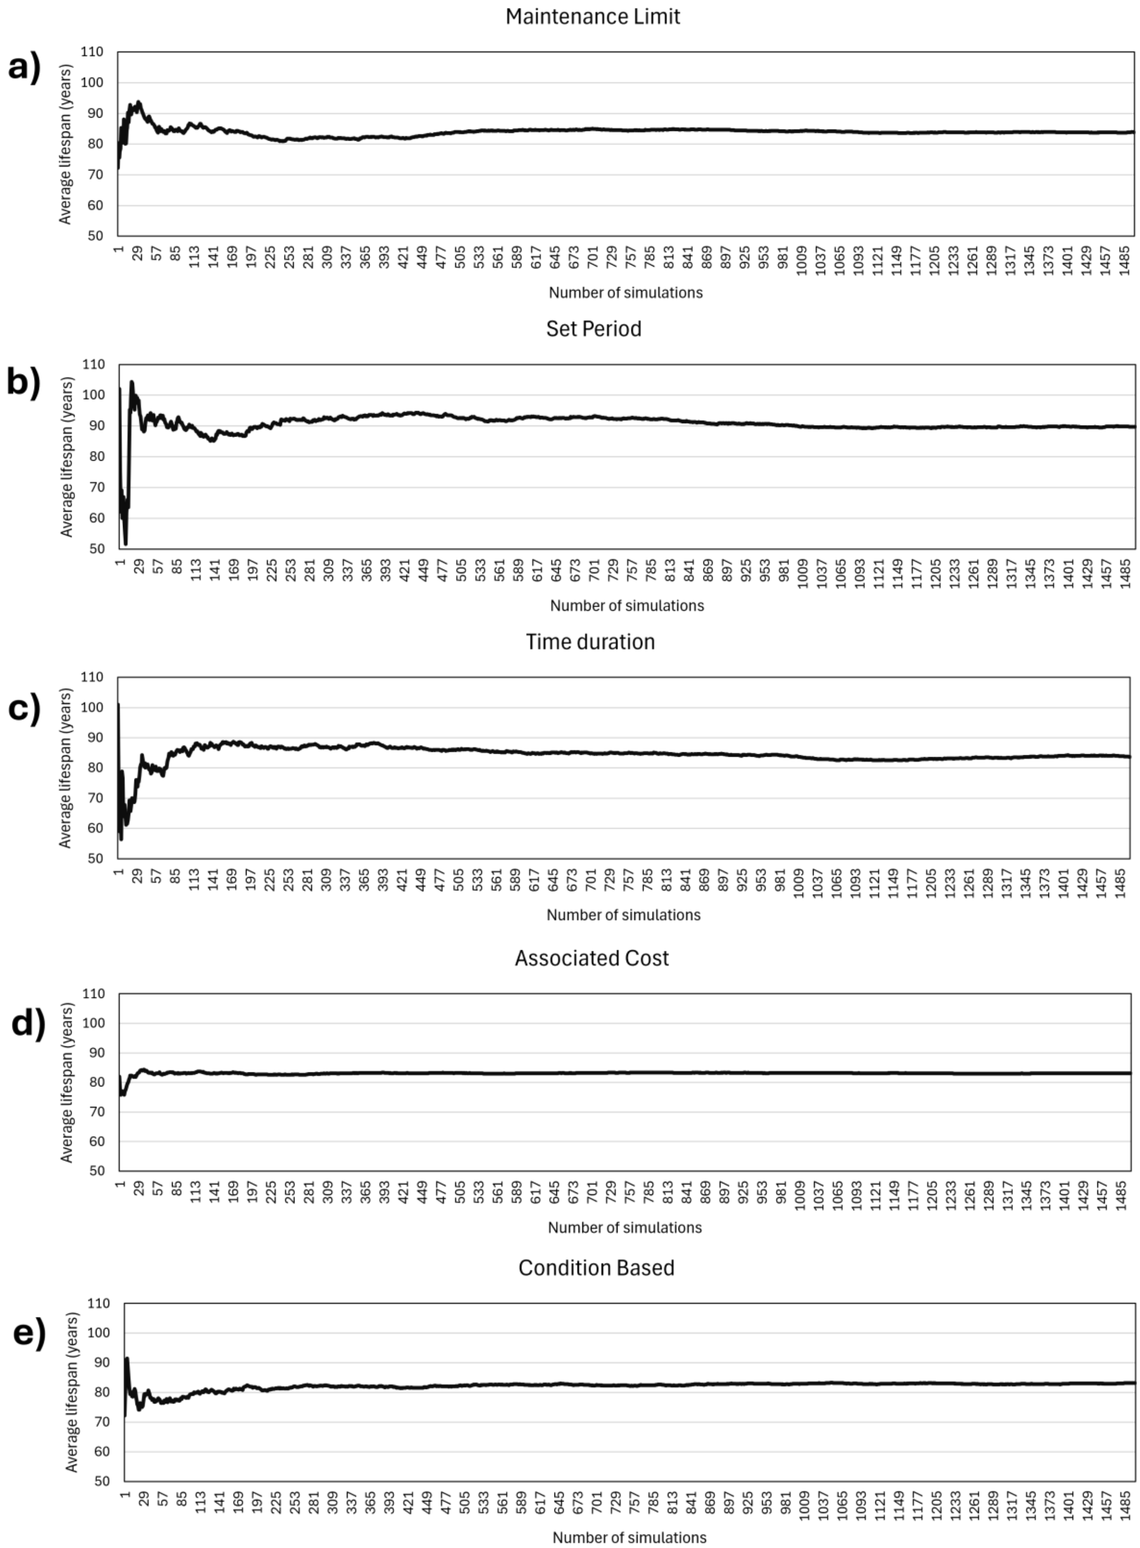

Figure 5.25: Convergence of the mean predicted lifespan for a plain line rail unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation

5.4.3.2 Expected lifespan of the plain line rail

The results from simulating the models with a single criterion considered as the renewal strategy is shown in Table 5.21. The relatively low renewal thresholds (Table 5.8) used for the plain line rail compared to other rail-based assets such as the switch panel is reflective of the behaviour seen during the analysis in Chapter 4. Despite this, the expected lifespans were much larger.

Renewal strategy	ML	SP	TD	AC	CB
Mean [SD] lifespan (years)	60.98 [29.08]	63.95 [30. 31]	64.34 [30.15]	60.17 [22.35]	61.47 [28.39]

Table 5.21: Average expected lifespan of a plain line rail unit with each renewal criterion considered in isolation for the renewal strategy

Given the low amount of faults found in the rail defect history, and because the mean lifespan and associated SD are similar regardless of which criterion was implemented, a further investigation into the behaviour of the rail prior to renewal was required. Figures 5.26A – 5.26E show the impact of the chosen criteria during the simulations – all demonstrate that the plain line rail was largely replaced due to an unrepairable failure rather than the renewal threshold being met. This is particularly evident for the TD criteria, where there were only 13 occurrences in 1500 simulations of 3 faults happening consecutively within 7 years of one another.

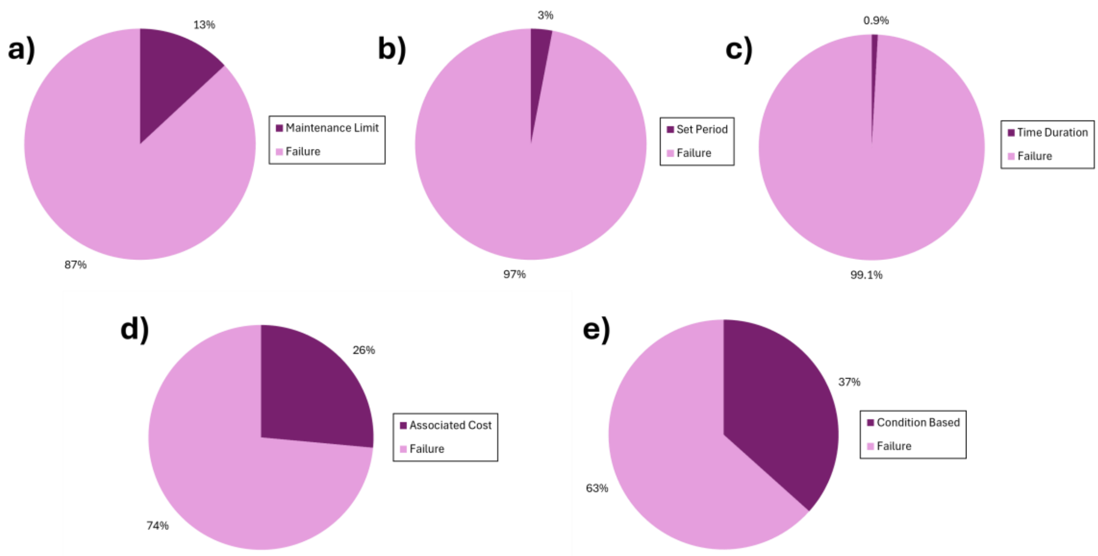

Figure 5.26: Distribution of causes for renewal on the crossing panel for a renewal strategy considering the a) ML, b) SP, c) TD, d) AC, e) CB criterion in isolation

A closer look at the causes of those failures revealed that 98% were due to the vertical wear reaching an unsafe limit, which is likely a reflection of this wear measurement being the only one obtainable from the historical UTU inspections. A better understanding of the degradation of the rail profile, such as its sidewear, will impact the predicted lifetimes but is not expected to change the proportion of replacements due to the chosen renewal strategy – most units will degrade due to wear.

Even with an extremely conservative renewal strategy of replacing the unit of plain line rail every time there is a fault, most renewals were still triggered due to the minimum rail depth having been reached (Figure 5.27). In fact, 63% of plain line rail units were expected to be replaced due to wear before it can experience at least one other fault. This mimics findings by Cannon et al. (2003), who found that the majority of replacements of the rail on a North American railway line was due to rail wear. The results in this section can immediately be used to update NRHS's asset management strategy to maximise life expectancy, as it is proven that the plain line rail renewals are mostly determined by rail wear. Further analysis should be completed to find the optimal frequency of the cyclic grinding campaign, so that faults on the rail are kept to a minimum whilst retaining as much of the rail material as possible. Additionally, it is known that strengthening the rail, such as with heat treatments, can make the rail more resistant to wear (Solano-Alvarez & Bhadeshia, 2024), with the findings in Figure 5.26 suggesting that implementing this after an initial rail renewal would be beneficial to prolong the useful lifetime of the HS1 line.

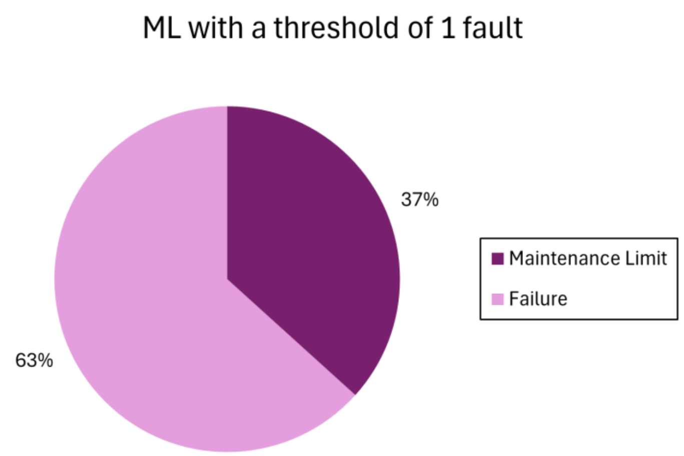

Figure 5.27: Distribution of causes for renewal on the crossing panel for a renewal strategy considering the ML criterion in isolation with a threshold of 1 fault

5.4.3.3 Expected condition of the plain line rail

Changing the simulation time of the PN to terminate after 80 years, with the AC criterion implemented as the sole criteria in the renewal strategy, will model the expected behaviour of a plain line rail unit over its entire operational period. Figure 5.28 and Table 5.22 reflect the low number of failures observed during the data analysis, as the asset is expected to be in a good condition over 90% of the time, and both local grinding or welding are expected to occur less than once over 80 years.

The likelihood of a 'vertical wear' fault – which requires a renewal to rectify – has a very small chance of appearing within the first 4 years of operation, then after 10 years its likelihood steadily increases over time. No evidence could be seen for the impact of renewal, however this is likely due to the large standard deviation from the implemented renewal strategy, with 95% of the plain line rail units expected to require a replacement between 15-105 years of service.

Action	Renewal	Local Grinding	Welding	Cyclic grinding	Fit emergency fishplates	Speed restriction (daily)	Inspection
Expected no. over 80-year observation	0.85	0.39	0.16	38.94	0.33	1.73	634.54

Table 5.22: Expected maintenance actions completed on a plain line rail unit with the AC criterion considered in the renewal strategy over a 80-year observation period

Table 5.22 confirms that localised grinding remains the most popular corrective maintenance action (other than cyclic grinding which is assumed to occur every 2 years regardless of the behaviour of the track) to restore the condition of a rail-based asset. Welding is often considered a higher-risk maintenance action as it is more susceptible to damage due to causing material inhomogeneity and geometric irregularity within the rails (Steenbergen, 2009). The table also shows that over 80 years the modelled plain line rail is expected to spend an average of only 2 days with a speed restriction implemented. This finding is validated by a report completed by the ORR (2024), who found that in 2023/24, only 458 trains (0.76% of all timetabled) were delayed due to issues attributed to HS1. Considering this statistic includes delays resulting from other railway assets, such as signalling issues or earthworks, it is reasonable for the developed model to find the impact of speed restrictions to be negligible on the asset management strategy of plain line rail.

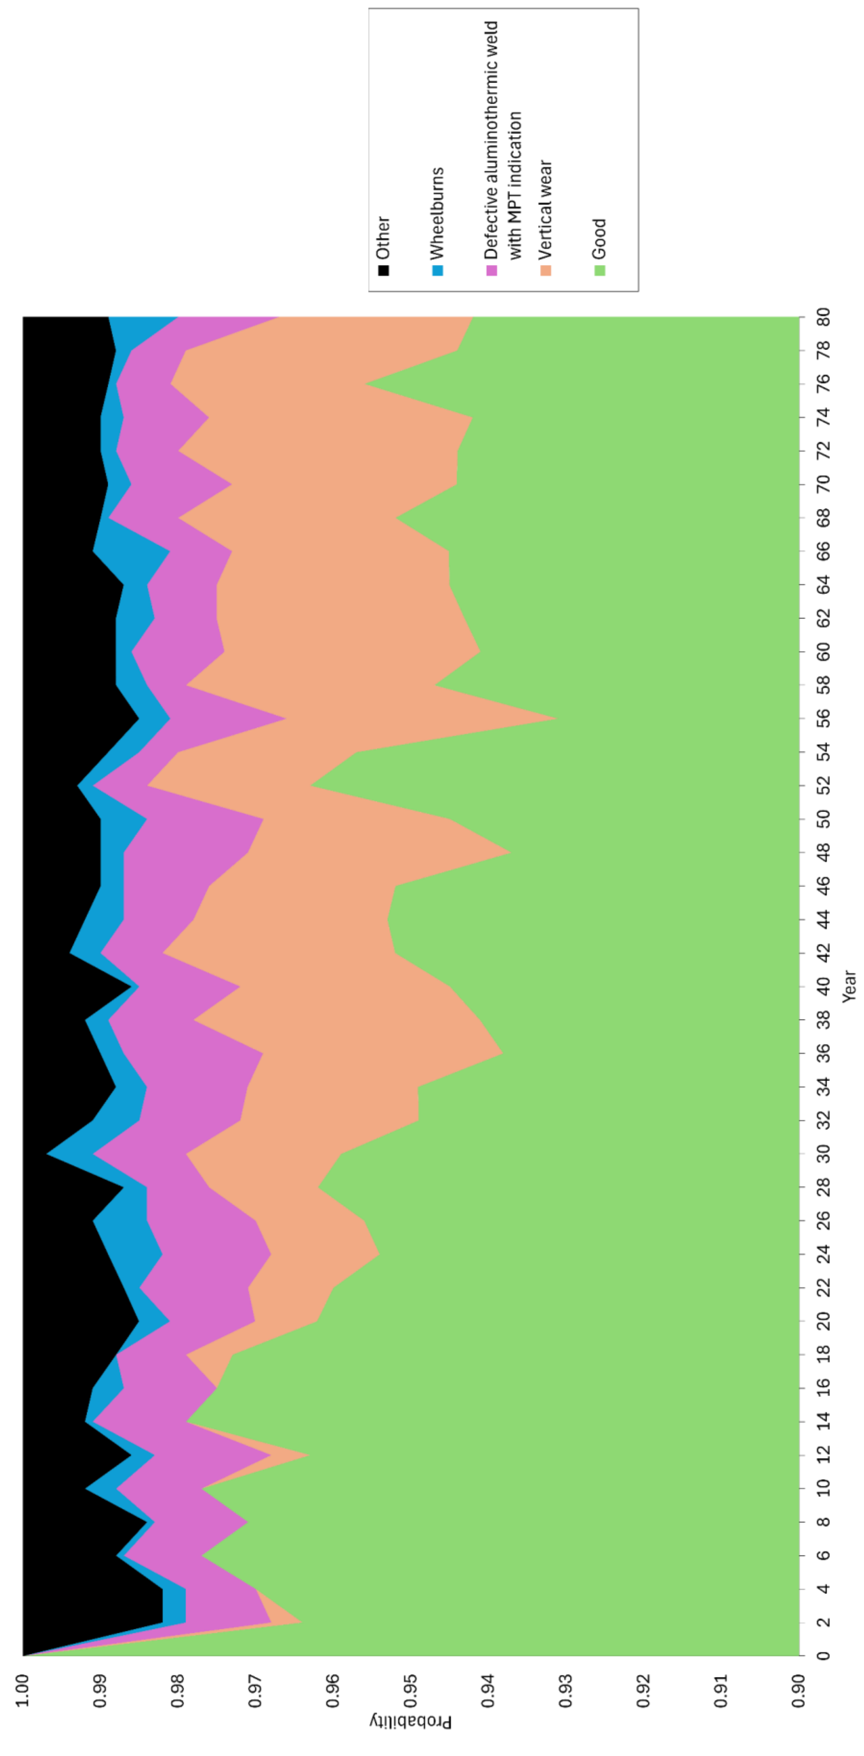

Figure 5.28: Average condition of a plain line rail unit over an 80-year observation period with the AC criterion considered in the renewal strategy

5.4.3.4 Investigating the importance of a renewal threshold

Within the original 40-year renewal strategy implemented at the installation of the HS1 line, it was expected that every track asset would be renewed at least once in that time. The predicted lifespans found throughout this section, though limited by the lack of suitable distributions to model the development of non-vertical wear, suggest that the plain line rail asset is expected to last much longer than 40 years without the need for renewal. Conversations with track engineers confirmed that this was a realistic prediction given the HS1 line's current behaviour but recommended that although the standards require renewal of the asset after 14mm had been lost from the surface of the rail, a replacement would most likely be scheduled after 8mm as it is easier to integrate the unit of rail with the depths of those surrounding it. It is important for the PN model to be adaptable, as renewal thresholds may change and any lifespan predictions made will be unreliable.

As the depth removed by cyclic grinding is based off the vertical wear rate, and Table 5.22 indicates that welding is not expected to occur more than once during a simulation hence is unlikely to trigger the renewal threshold, the CB renewal criterion can be used to update the renewal predictions in line with current practises. By changing the parameter of the renewal criterion to 8mm, the expected lifetime reduced from 61.47 years to 34.75 years. Seamlessly integrating the unit of rail with the depth of the surrounding sections is important as it reduces the forces exerted on the rail by the rolling stock rapidly changing height, minimising localised defects, however this also reduces the useful lifespan of the rail by 27 years. Additional research should be completed regarding the optimal threshold for rail wear depth that maximises its lifespan whilst minimising the negative impact of inconsistent rail heights, however this is currently beyond the scope of this thesis.

Additionally, the original 14mm depth limit is a critical threshold that risks derailment if breached. An efficient renewal schedule will aim to replace the unit before this threshold is met, and as such it is important to know how quickly the wear limit is expected to be reached from previously recorded conditions. For example, if the unit is expected to experience late-stage wear-out, there could be a very short time between a safe condition of 10mm of vertical wear recorded to a critical limit of 14mm of vertical wear reached, requiring an emergency renewal. Changing the threshold for minimum rail depth by 1mm each time for the CB renewal strategy revealed a linear increase in predicted lifespan (Figure 5.29), with an average of 4 years added to the expected lifespan of the rail for each mm added to the renewal threshold. This gives adequate time between 13mm of wear recorded and scheduling a renewal, for the unit to be replaced before passengers on the rail are put at risk.

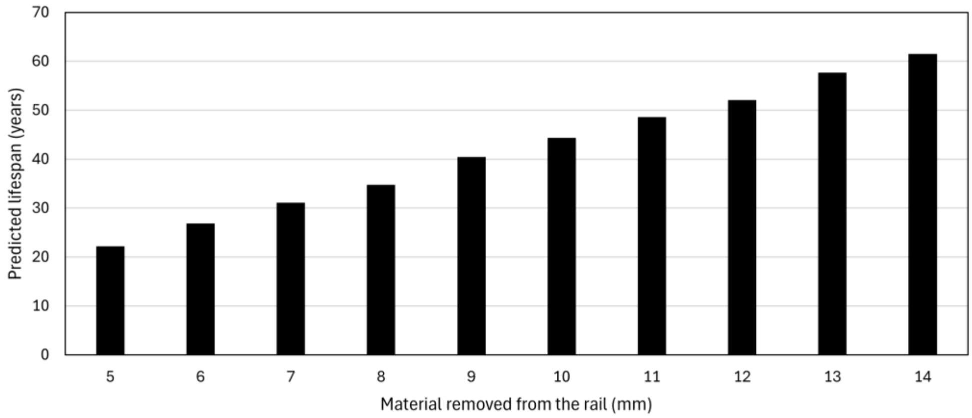

Figure 5.29: Predicted lifespan of the plain line rail achieved by changing the thresholds for a renewal strategy considering solely the CB criteria

5.4.4 Ballast

5.4.4.1 Validation

Due to the assumption that the asset management of the ballast follows the information defined in the standards (hence it is only rectified once the ballast has degraded beyond a 'good' condition), and evidence showing that this is not implemented in practice (Table 5.2), it is impossible to validate the tamping intervention history with what was observed during the simulation of the ballast PN. However, implementing each of the five renewal criteria in isolation and finding the average lifespan after 1500 simulations revealed that the results had all converged (Figure 5.30), hence it is assumed that the model is appropriate for predicting the expected lifespan of a 200m section of ballast.

5.4.4.2 Expected lifespan of the ballast

Simulating the ballast PN model with each criterion considered in isolation as the renewal strategy was used to find the expected lifespan, with the results shown in Table 5.23.

Renewal strategy	**ML**	**SP**	**TD**	**AC**	**CB**
Mean [SD] lifespan (years)	79.36 [27.35]	88.58 [52.89]	78.50 [48.80]	86.74 [9.10]	80.28 [28.63]

Table 5.23: Average expected lifespan of a ballast unit with each renewal criterion considered in isolation for the renewal strategy

The standard deviations for the mean average lifespans were comparatively large, except when the renewal strategy is based solely on the associated cost of operating the ballast unit. This results in a large variation between the predicted lifespans - for example, 95% of the expected lifespans for the ballast with only the ML criteria included in the renewal strategy are expected to lie within a 109-year range. This may result from the data populating the model, as it is drawn from units with varying characteristics (such as the maximum permittable speed that rolling stock can travel over the unit) that are expected to influence the behaviour of the ballast. This causes a wide variety in degradation rates, which is reflected in the predicted lifetimes. In Chapter six, artificial neural networks are introduced to explore whether modelling ballast subgroups will reduce the uncertainty surrounding lifespan predictions.

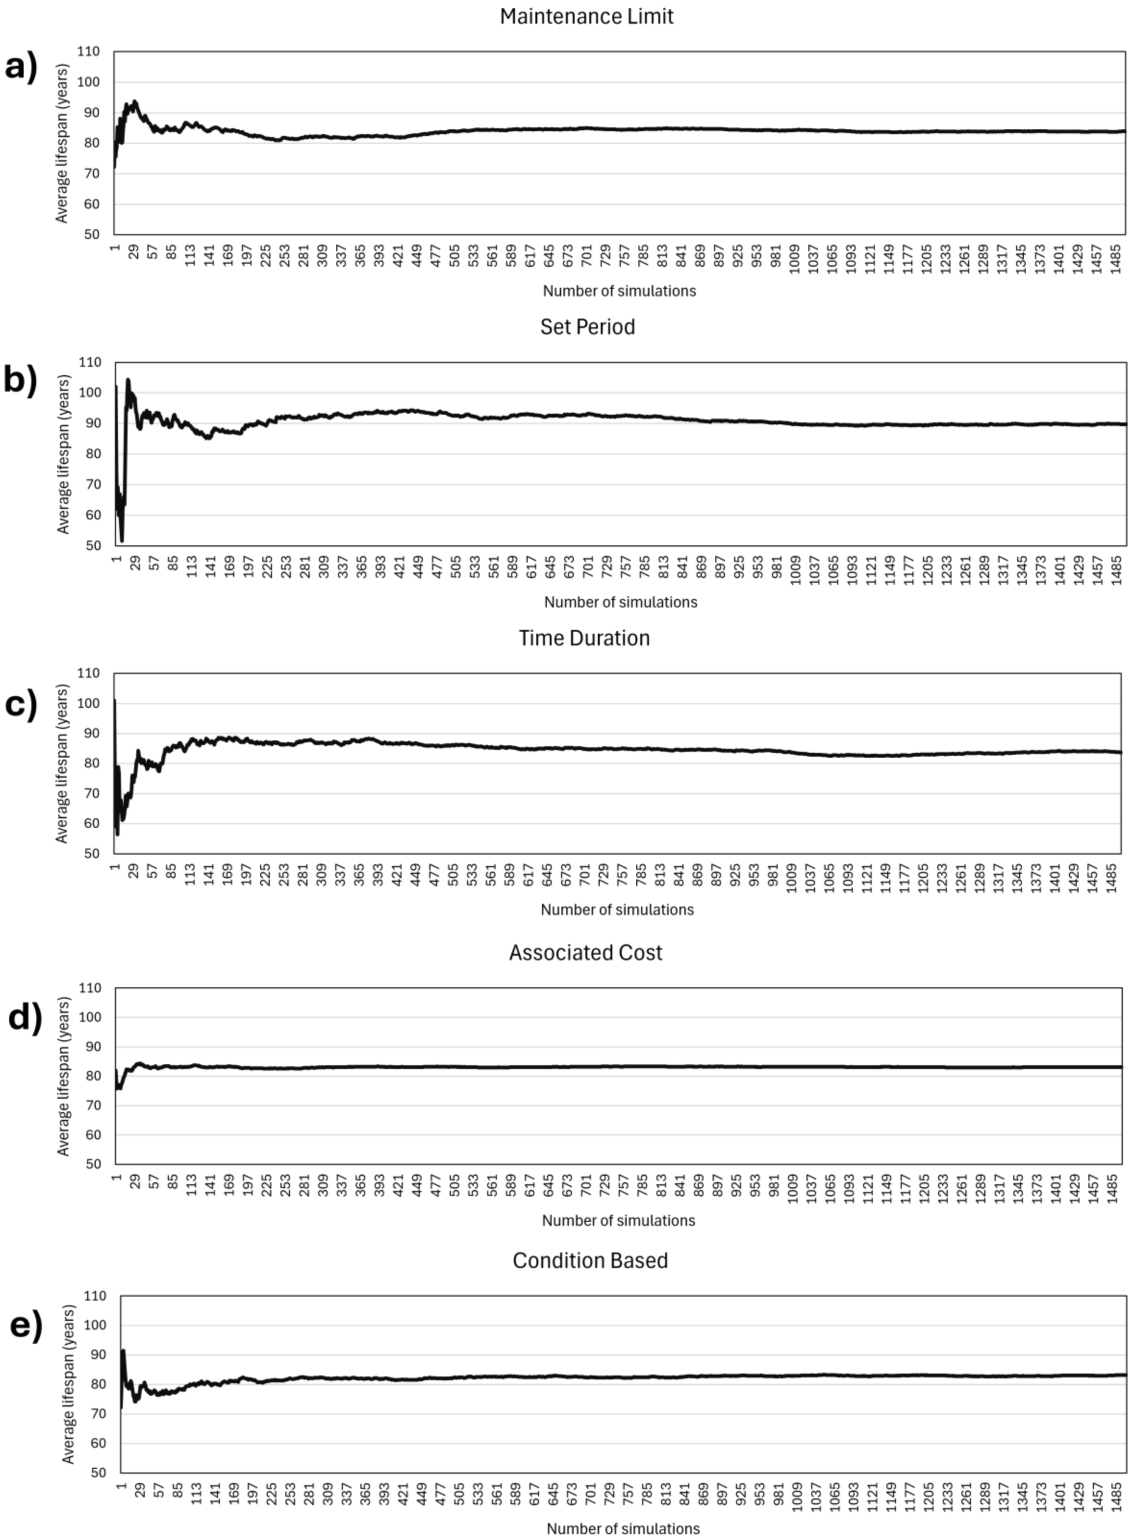

Figure 5.30: Convergence of the mean predicted lifespan for a ballast unit with a renewal strategy considering the a) ML b) SP c) TD d) AC e) CB criterion in isolation

Additionally, the lifetimes obtained from the analysis shown in Table 5.23 (although dependent on the given renewal parameters) are significantly larger than any of the other considered track assets and are over double the length of the original 40-year renewal schedule. These findings are thought to be due to a combination of several factors:

- The natural long-lasting behaviour of the ballast – for example, research in by Zhang et al. (2019) found that less than 20% of the studied high-speed railway ballast aggregate was damaged after 4 million cycles of cyclic loads representing high-speed trains
- The assumption that a ballast unit cannot be tamped whilst in a 'good' condition despite historically this occurring for 66% of all high-output tamping interventions
- The assumption that the aggregate across a 200m section is renewed in isolation, with no opportunity for units to be replaced prematurely via opportunistic renewal

It is acknowledged that an entire line model, that considers the behaviour of multiple adjoining ballast units to influence the frequency of maintenance and renewals performed, may be better suited to address some of these factors, however the developed PN only models a single 200m section of ballast behaving in isolation.

5.4.4.3 Expected condition of the ballast

Changing the PN to terminate the simulation after 100 years instead of when a renewal has occurred, and implementing only the AC criterion in the renewal strategy (as it had the least amount of variance in the predicted lifespans) has found the ballast's expected behaviour over its lifespan. Figure 5.31 shows that the likelihood of the ballast being in a 'good' condition rapidly decreases over the initial 12 years of operation, then remains moderately consistent at around 50% for the next 60 years. The impact of renewal can be observed, with the unit's expected condition improving steadily from around 77 years, where large-scale ballast replacements are expected to start.

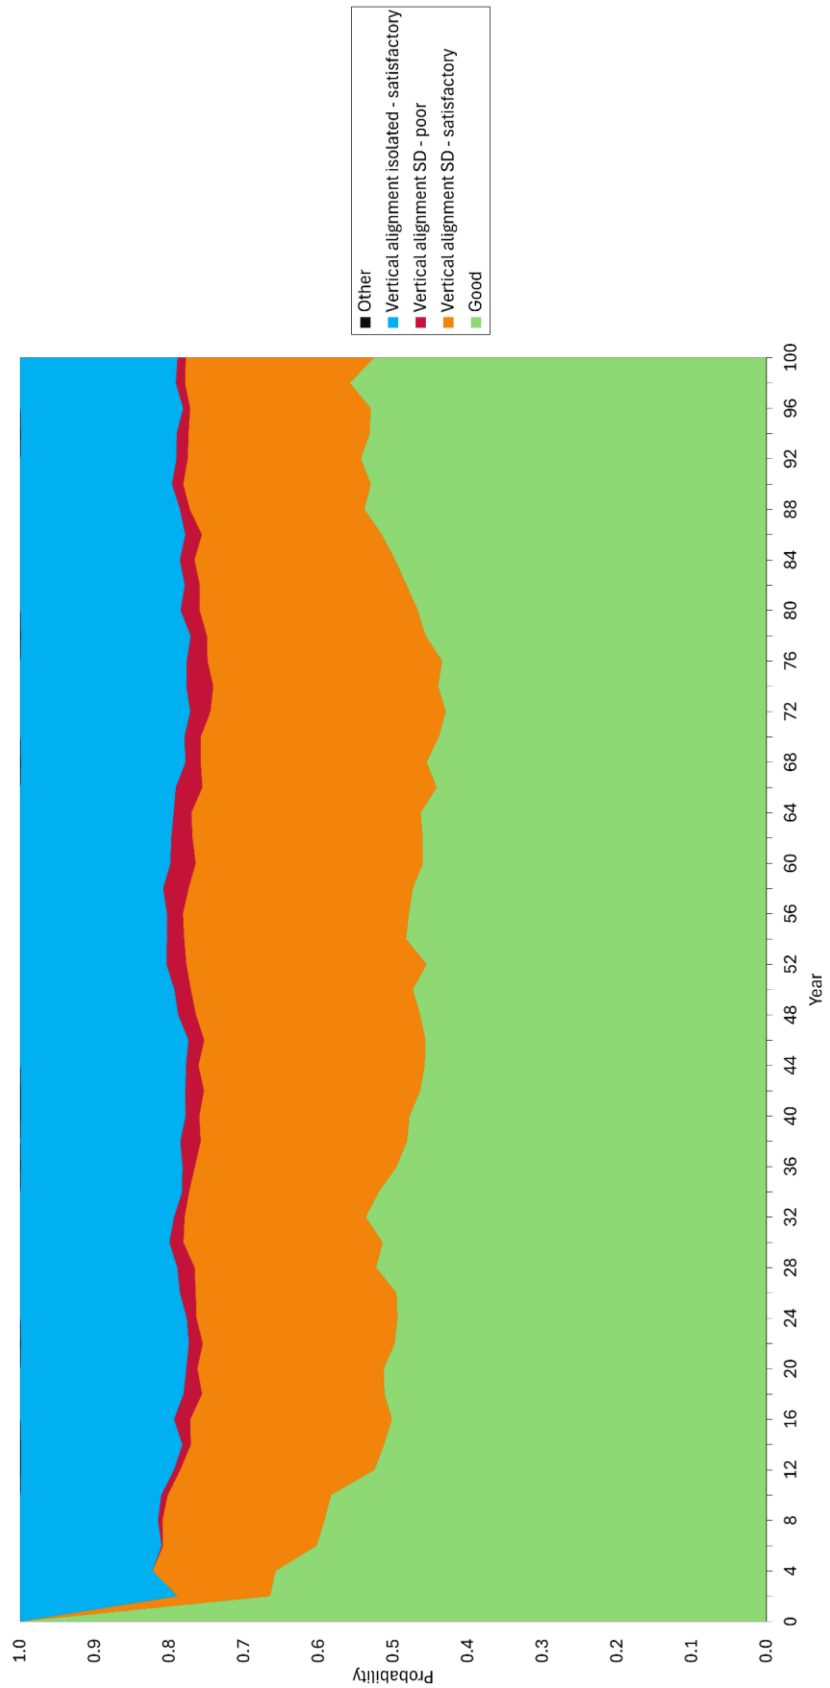

Figure 5.31: Average condition of a ballast unit over a 100-year observation period with the AC criterion considered in the renewal strategy

Table 5.24 shows that high-output tamping (which corrects the entire ballast unit to rectify the SD of the vertical alignment) is a far more common maintenance action than localised tamping which only corrects a small area of ballast for isolated vertical alignment faults. Despite this, the probability of the ballast unit experiencing an isolated vertical alignment fault is consistently around 20% (Figure 5.31). This results from the assumption that over 40% of isolated faults deemed in a 'satisfactory' condition are permitted to degrade further until they are in a more severe condition or are inadvertently rectified by high-output tamping aiming to correct a poor SD. Therefore, despite their lower likelihood of occurrence, once an isolated vertical alignment fault has developed it is likely to be present on the ballast unit for several years.

Action	Renewal	High-output tamping	Local tamping	Block line (daily)	Impose speed restrictions (daily)	Inspection
Expected no. over 100-year observation	0.92	9.49	2.00	0	0	890

Table 5.24: Expected maintenance actions completed on a ballast unit with the AC criterion considered in the renewal strategy over a 40-year observation period

Blocking the line or imposing speed restrictions are only enacted when the isolated vertical alignment failure is in a critical state, but Table 5.24 shows that this did not occur during the simulation of the PN model. This is supported by the historical geometry behaviour, as there was only one occurrence of a critical isolated vertical alignment reading recorded in 186 observations of 849 ballast units.

5.4.4.4 Adapting the asset management strategy

Table 5.2 demonstrated that the ballast is often maintained too early rather than when its condition had reached a 'poor' state, and correcting its condition is a necessity. This is likely because the tamping equipment is hired annually from an external company giving a limited window of time to perform maintenance; tamping is scheduled to be as cost effective as possible by correcting the largest areas within this window. Also, the machines will typically tamp beyond the confines of the required 200m section and correct areas adjacent to degraded sections, especially if located between two poor geometry sections. The purpose of any maintenance is to extend the usable lifespan of an asset, therefore any changes to the way maintenance is performed will subsequentially impact the predicted renewal time. One of the largest benefits of creating the models from a Petri net (instead of relying solely on – for example – reliability analysis of historical datasets) is it is easily adaptable, which allows for optimisation of the asset management strategy. This can be investigated by using an optimisation technique such as the genetic algorithm to quickly investigate multiple strategies, or by changing a single parameter at a time, such as the inspection frequency.

The focus of this research is not centred around investigating the implemented maintenance strategy, but the model can be used to address the previously described issue of over-tamping to show the extended lifespan that can be achieved if a fleet of tamping machines were owned by NRHS, meaning that ballast maintenance can be performed only when needed. This is achieved by changing the probability of maintenance occurring during a section in a 'satisfactory' condition to 0%, so maintenance was only performed if the condition had degraded to a poor (or beyond) condition for either failure type – standard deviation and isolated measurements of the vertical alignment geometry. This change was enacted by using the decision-making strategy described in Figure 5.5 and adjusting the firing time of T2 to zero, so that the ballast is always permitted to degrade until tamping interventions are necessary.

It was found that implementing a renewal policy based on a maintenance limit of 9 interventions increased the lifespan by 53 years to an average of 132 years, but reduced the probability of the time the ballast unit spends in a 'good' condition to an average of just 27% (Figure 5.32).

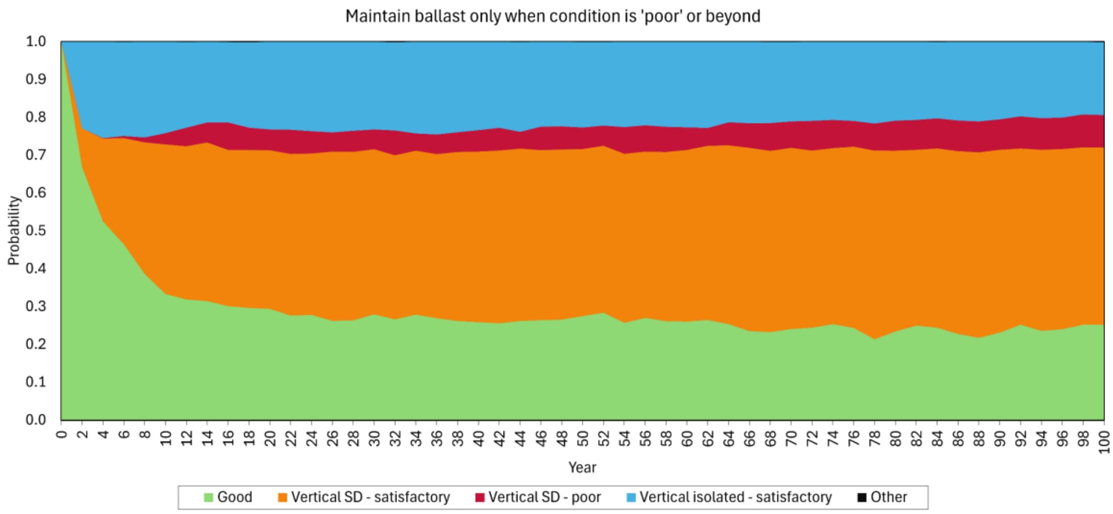

Figure 5.32: Average condition of a ballast unit over a 100-year period with the ML criterion considered in the renewal strategy and a maintenance strategy of only maintaining once the ballast condition is poor or beyond

This may be undesirable as although there is a cost benefit to replacing the ballast less frequently, the ride quality will be poorer and there is emerging evidence to suggest that a poor ballast geometry can increase the degradation of rail-based assets (Ngamkhanong et al., 2021) which may ultimately add costs to the track maintenance budget. Conversely, changing the management strategy to always perform maintenance once the condition of the ballast for either failure type has degraded beyond a 'good' condition will improve the ride quality and ensure it there is never above a 30% chance of the ballast being in degraded state (Figure 5.33), but the unit will require a renewal every 60.43 years.

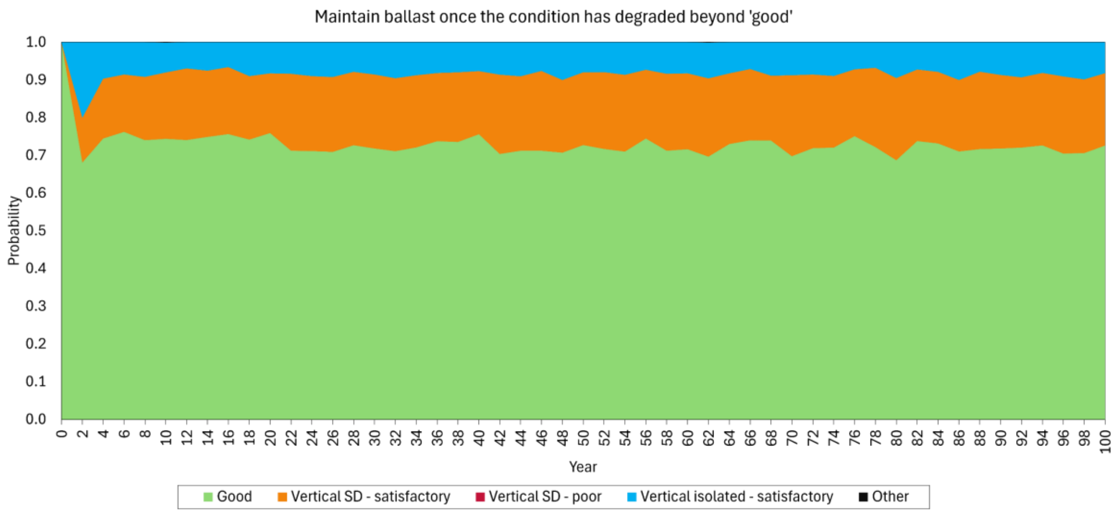

Figure 5.33: Average condition of a ballast unit over a 100-year period with the ML criterion considered in the renewal strategy and a maintenance strategy of only maintaining once the ballast condition has degraded beyond 'good'

5.5 Summary

During this chapter, a methodology has been developed to model the asset management strategy of the switch panel, crossing panel, plain line rail and the ballast. To address the complexities of engineering procedures, additional features – inhibitor arcs, place dependent transitions, and reset transitions - have been added to the original Petri net structure. Considerable effort has been spent to ensure that the models accurately reflect the inspection, degradation, maintenance, and renewal process. Railway standards have detailed the degradation process for a failure and the required maintenance action for each fault. Once built, the model was populated using a mixture of continuous probability distributions found from statistical analysis of failure data, and recommended timescales for maintenance procedures taken from the standards.

An algorithm was built to compute the Petri nets in C++, and they were simulated until convergence was reached at the 1% level. The movement and duration of time a token spent in a place was recorded to collect statistics concerning the asset's behaviour over its lifetime. Given the wide variety of results that can be collected from simulating the model, not every statistic from every single model is discussed, but overall, the results have demonstrated the methodology's ability to:

- Predict the expected lifespan of a railway component given a renewal strategy
- Find the distribution of the lifespan predictions
- Observe the expected condition of an asset over a lifespan
- Breakdown the projected cost per maintenance action over a control period
- Predict the expected number of maintenance actions over a lifespan
- Find the most influential renewal parameter for triggering a renewal
- Consider the likelihood of each fault within a failure type occurring
- Adapt the renewal thresholds to investigate their impact on the predicted lifespans
- Investigate and optimise the impact of the management strategy on lifespan predictions

Given the issues with the low quantity of failure data discussed in Chapter 4, there is high uncertainty surrounding the current lifetime predictions. This is expected to improve once the track has aged and there is a larger failure dataset for analysis and populating the models. Any results discussed are also subject to the implemented renewal parameters which are estimated. Despite this, there were several key findings:

- The switch panel degrades much faster than the crossing panel and the two should be replaced separately
- The plain line rail will largely degrade due to vertical wear instead of other failures
- The ballast is historically over tamped and is expected to be operational for a lot longer if this is rectified
- All assets are expected to be in a good condition with no faults at least 50% of the time

Chapter 6

Developing artificial neural networks to investigate asset degradation

6.1 Introduction

In Chapter five, the capability of Petri nets was demonstrated to model the behaviour of high-speed railway assets, predicting the expected lifetime given a choice of different renewal strategies. However, it was also observed that there was often a high standard deviation for the distribution of lifetime predictions. This may occur for a variety of reasons, including the implemented renewal strategy – for example, a strategy based around the Time Duration criterion considered in isolation typically resulted in an estimated lifetime with a large variation and a heavy right skew. It can also stem from the distributions used to populate the transitions in the model; a dataset of times-to-failure derived from a railway track with varying features can cause a high uncertainty in any fitted distributions which is reflected in the predictions. The aim of this chapter is to address the limitations of the previous reliability analysis by accounting for all the factors that can influence track behaviour when modelling the degradation of an asset type. Aside from being used in the PN models to find expected lifetimes with less variation, understanding the behaviour of different subgroups can be used to inform the order which renewals should be completed, as typically a renewal schedule spans over several years due to budget constraints and to minimise large disruptions to the service. The fastest degrading assets should be targeted first to avoid unnecessary maintenance costs.

The methodology described in Chapter 4 can be applied to investigate an asset subgroup, such as conventional or high-speed track. However, a large percentage of conventional speed areas on the HS1 line are located at St Pancras which experiences higher traffic, and it is impossible to recognise the impact of either speed <u>or</u> traffic on the expected lifetime predictions. Dividing the two groups again by expected traffic per subgroup will reduce the size of the datasets, and the more specific the subgroup, the smaller the available dataset; for example, studying the degradation of a section of plain line rail at high speed, experiencing 33.47MGT of traffic per year, and not present in a tunnel or located on a curve, leaves only 23 out of a total of 1719 units to be considered. This can increase bias and is time consuming as a separate analysis must be completed for each subgroup. To overcome these issues, a different methodology was needed. This was chosen to be artificial neural networks (ANN).

Research by Alwosheel, Van Cranenburgh, and Chorus (2018) states that the size of a dataset used in an ANN analysis should be at least 10-50 times the amount of weights used in the net, for efficient training. This rule of thumb excluded all faults obtained from the rail defect history data due to the low number of recorded faults, but as the standard deviation of the vertical alignment for the ballast and the vertical wear of the plain line rail were obtained from condition data, both had an adequate sample size for training. Given that the majority of all plain line replacements were triggered due to the vertical wear fault, and 80% of all maintenance actions on the ballast were high-output tamping, a response to a high SD of the vertical alignment, it is assumed that considering these faults alone are suitable for recognising the impact of different features on the expected lifetimes of the ballast and plain line rail.

The literature review highlighted different ways in which ANN have been applied to predict the behaviour of track assets. Previous reliability analysis has fitted linear regression lines to the condition data, negating the need for an ANN that predicts the future condition of the asset, such as those used by Shafahi et al. (2008) and Yilboga et al. (2010). Both Guler (2014) and Khajehei et al. (2022) used an ANN effectively by evaluating the known characteristics of the track to predict the gradient of the fitted regression lines, which indicates how quickly the asset is expected to degrade. A similar approach is considered in this chapter.

6.2 Determining track features

Before the ANN can be developed, the various characteristics of the line that can impact degradation must first be identified. There is currently no single index of the HS1 track that lists the exact features of every section, including the infrastructure present and operative features such as the maximum permitted speed for the rolling stock, so one was created. This information was largely gained by observing the schematic, a small segment of which is shown in Figure 6.1.

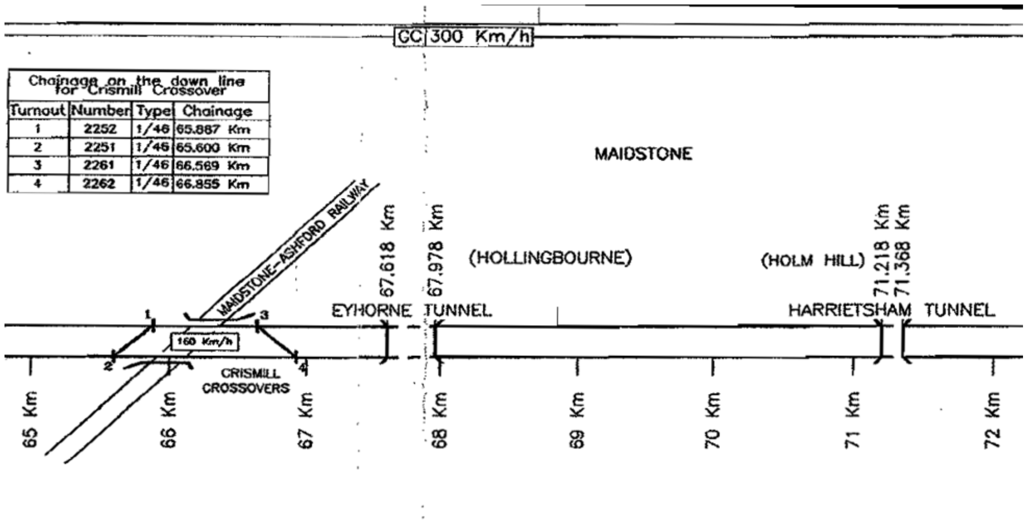

Figure 6.1: A section of the HS1 schematic

6.2.1 Geographical features

Within the schematic, there were seven geographical features identified:

- Cant – the optimal height difference between the elevation of the left and right rail, which is used to indicate where the line is curved.
- Tunnel – the line is in a tunnel.
- Viaduct – the line is over a viaduct
- Elevation - the relative height of the rail in comparison to a baseline of 0m at the St Pancras station.
- Incline - the gradient at which the line increases or decreases in elevation.
- Superstructure - the rolling stock is travelling through S&C or plain line.
- Substructure - the rolling stock is travelling upon ballasted or slab track.

6.2.2 Operational features

Three further operational features were also defined:

- Speed - the maximum permitted speed for the rolling stock travelling over the unit.
- Traffic - the expected yearly tonnage from the rolling stock.
- Age - the age of the track from its installation or renewal.

The speed could also be obtained from the schematic, and the age is considered from when the line opened (28/09/2003 for units located between the beginning of the channel tunnel and Fawkham Junction in Kent, or 14/11/2007 for units located elsewhere on the line) as no renewals had yet taken place. The expected yearly tonnage was found from a report on the amount of tonnage each S&C unit had received in 2017 and was used to approximate the tonnage on the main lines between successive S&C units. It is assumed that every year of operation received similar tonnage distribution to what was observed in 2017.

6.2.3 Environmental features and opposing characteristics

Environmental features could not be obtained but are not expected to vary (or exist in cases such as falling rocks) across the relatively short railway line. Wherever possible, additional information provided by NRHS was used to confirm the location of the considered features, and a full index was created for the HS1 line. The index could then define the features of every unit of the plain line rail and ballast; if a unit contained two opposing characteristics (for example, if a 200m section of ballast extended across track within and outside a tunnel), it was retained if over 70% of the unit existed within one feature, with that feature chosen to represent it, and removed if not.

6.3 Developing the ANN

6.3.1 Considered variables

The features defined in the previous section, as well the number of high-output tamps previously completed on the ballast unit (Tamps) were used as input variables in an ANN, with the intercept and gradient of the linear regression lines found in Chapter 3 used as the outputs (Tables 6.1-6.2). Given that the intercept for the vertical wear of the plain line rail was assumed to be zero as the rail is initially in a perfect condition, this was not considered as an output variable. The predictions made from the ANN can be used to calculate the time at which a failure threshold will be reached, which can be incorporated back into the PN as a fixed firing time for the relevant degradation transition. Alternatively, this time could be inverted and used as the degradation rate parameter for the exponential distribution to retain stochastic behaviour in the Petri nets.

Variable	Unit	Layer
$\boldsymbol{B_1}$ - Speed	Kilometres per hour	Input
$\boldsymbol{B_2}$ - Traffic	Annual KTonnes	Input
$\boldsymbol{B_3}$ – Cant	Millimetres	Input
$\boldsymbol{B_4}$ - Tunnel	1 if located within a tunnel, 0 if not	Input
$\boldsymbol{B_5}$ - Viaduct	1 if located on a viaduct, 0 if not	Input
$\boldsymbol{B_6}$ - Elevation	Metres	Input
$\boldsymbol{B_7}$ - Incline	Percentage	Input
$\boldsymbol{B_8}$ - Age	Days	Input
$\boldsymbol{B_9}$ - Superstructure	1 if beneath a S&C, 0 if beneath plain line	Input
$\boldsymbol{B_{10}}$ – Tamps	Integer	Input
$\boldsymbol{B_{11}}$ – Intercept	Millimetres	Output
$\boldsymbol{B_{12}}$ – Gradient	Millimetres per day	Output

Table 6.1: Variables of the SD of the vertical alignment (ballast) ANN

Variable	Unit	Layer
$\boldsymbol{PL_1}$ - Speed	Kilometres per hour	Input
$\boldsymbol{PL_2}$ - Traffic	Annual KTonnes	Input
$\boldsymbol{PL_3}$ – Cant	Millimetres	Input
$\boldsymbol{PL_4}$ - Tunnel	1 if located within a tunnel, 0 if not	Input
$\boldsymbol{PL_5}$ - Viaduct	1 if located on a viaduct, 0 if not	Input
$\boldsymbol{PL_6}$ - Elevation	Metres	Input
$\boldsymbol{PL_7}$ - Incline	Percentage	Input
$\boldsymbol{PL_8}$ - Age	Days	Input
$\boldsymbol{PL_9}$ – Substructure	1 if located over slab track, 0 if located over ballasted track	Input
$\boldsymbol{PL_{10}}$ – Gradient	Millimetres per day	Output

Table 6.2: Variables of the vertical wear (plain line) ANN

6.3.2 Sufficient variation

It is important to ensure that the inputs of an ANN are relevant to the outputs, to improve the efficiency and performance of the model. A review into feature selection for machine learning was completed (Li et al., 2017), stating that an input variable with low variation will contribute very little to the model's performance, as it will be unable to help discriminate between differences for the output variables. Before the ANN was implemented, descriptive statistics for each input were collected and displayed in Table 6.3. The results show there was reasonable variation for each considered feature and therefore should all be included as inputs to the ANN, to study their impact on the degradation rate.

Feature	Minimum	Maximum	Mean	Standard deviation
Traffic	5276	22127	13310	4005
Speed	160	300	277	34
Cant	0	160	73	63
Elevation	-10	110	60	30
Incline	-2.5	2.5	0.1	1.4
Age	0	6607	2532	1880
Tunnel	0	1	0.30	0.46
Superstructure	0	1	0.07	0.26
Substructure	0	1	0.22	0.41
Tamps	0	18	3.25	3.22

Table 6.3: Variance of the considered features

6.3.3 Training parameters

The algorithm to compute the ANN was developed in C++, and implemented under the following conditions:

- The ReLU activation function (restated below in Equation 6.1) was used within the hidden layers.

$$f(x) = \max(0, x) \tag{6.1}$$

- The Linear activation function (Equation 6.2) was used for the output layer.

$$f(x) = x \tag{6.2}$$

- The learning rate was set to 0.005.
- A 70:30 split of the available data was used for training and testing datasets respectively.
- Max-min normalisation was applied to the datasets before processing.
- The mean squared error was used as the cost function.
- He normal initialisation (He et al., 2015) was used to set the initial value of the weights, and the bias was set to 1.

An example of the structure is given in Figure 6.2 for the SD of the vertical alignment ANN, modelling degradation of the ballast. During the analysis, the number of hidden layers within the structure was varied between 1 and 2 layers, with the number of neurons in each layer varied between 0.5 - 4 times the amount of input variables. Each structure was trained for 100 epochs of the training dataset, and the results from the cost function applied to the testing and training datasets was used to find the optimal structure and number of epochs for model performance. Once established, the prediction ability of the ANN was evaluated by finding the R^2 value between the predictions made from implementing the testing dataset in the model compared to the true output value.

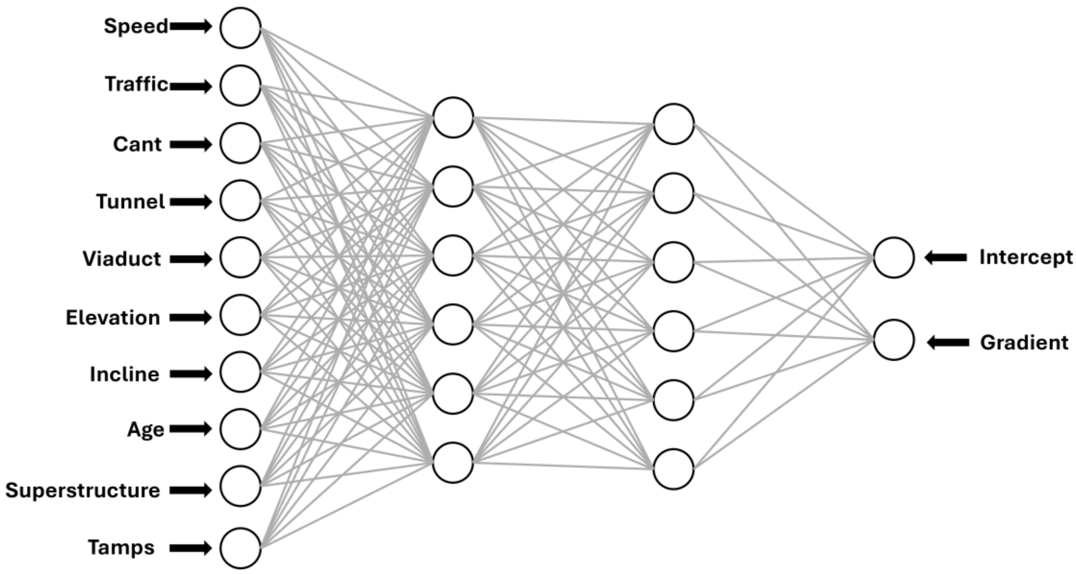

Figure 6.2: Structure of the ANN modelling ballast degradation

Initially, both ANNs were poor at predicting the required outputs for each fault with an optimal R^2 score of less than 0.2. Plotting the residuals revealed clusters around the extremes due to the 'clipping' of data during normalisation, but it was found that removing the clipped data from the analysis and retraining the model had an adverse effect on the performance. Further adaptions were investigated to try improve the model's prediction accuracy, including changing B_{11} (*Intercept*) to an input variable to provide an anchor to the regression line, modifying the learning rate, and changing the error function to MAE. It was theorised that the hidden layer's activation function could be causing the poor performance due to the 'dying ReLU' problem (Singh et al., 2023), where a negative input value to a neuron is processed as 0 after being passed through the activation function, meaning any associated weights are not adjusted during backpropagation. Over time, this can lead to the neuron being unable to learn anything about the data, contributing to the model's poor performance. To rectify this, the ReLU activation function was replaced with 'leaky ReLU' (Equation 6.3):

$$f(x) = \begin{cases} 0.01x \ \ for\ x < 0 \\ x \ \ for\ x \geq 0 \end{cases} \tag{6.3}$$

Including B_{11} as an input variable resulted in a slight but notable improvement in the model's ability, but the other modifications to the structure or training parameters had only marginal impacts on the performance, with the inclusion of a second hidden layer or many hidden neurons often resulting in slower running models with no noticeable improvements in prediction accuracy.

6.3.4 Artificial dataset

The poor performance of the model was unexpected, especially given the variety of training parameters that were investigated. To ensure that the implementation of the methodology was not flawed, and therefore affecting the ANN's ability to make predictions, an artificial dataset with a clear established degradation trend was created and passed through the developed model. The predictions made were studied to see whether the model had recognised the trend, thus evaluating how well the current methodology can be applied to data. A dataset was created using the exact same input values for the ballast ANN (Table 6.1 with B_{11} now considered as an input variable) but changing the output according to Equation 6.4:

$$B_{12} = \log(B_2) + \frac{B_1}{10} + \left(B_7 \times \frac{B_3}{100}\right) + 2B_9 \tag{6.4}$$

Equation 6.4 has a combination of linear and nonlinear transformations of the inputs but crucially there were also variables that had no impact on the degradation rate. To reflect the natural variation and noise that occurs between sections, 20% of the dataset was randomly selected and B_{12} was multiplied by a random factor between 0.5 - 2. Identical training conditions to section 6.3.3 were implemented, except for the activation function for hidden layers which was now as described in Equation 6.3. The model's structure had a single hidden layer and was simulated across a range of hidden neurons for 100 epochs, finding that the MSE of the testing dataset was minimised at 20 neurons (Table 6.4), and although the model did not overtrain, the error had converged after 50 epochs of the training dataset (Figure 6.3).

Neurons in hidden layer	MSE of the testing dataset
5	0.0138
10	0.0098
15	0.0094
20	0.0092
25	0.0094
30	0.0102
35	0.0105

Table 6.4: Prediction accuracy of the ANN with a varying number of hidden neurons

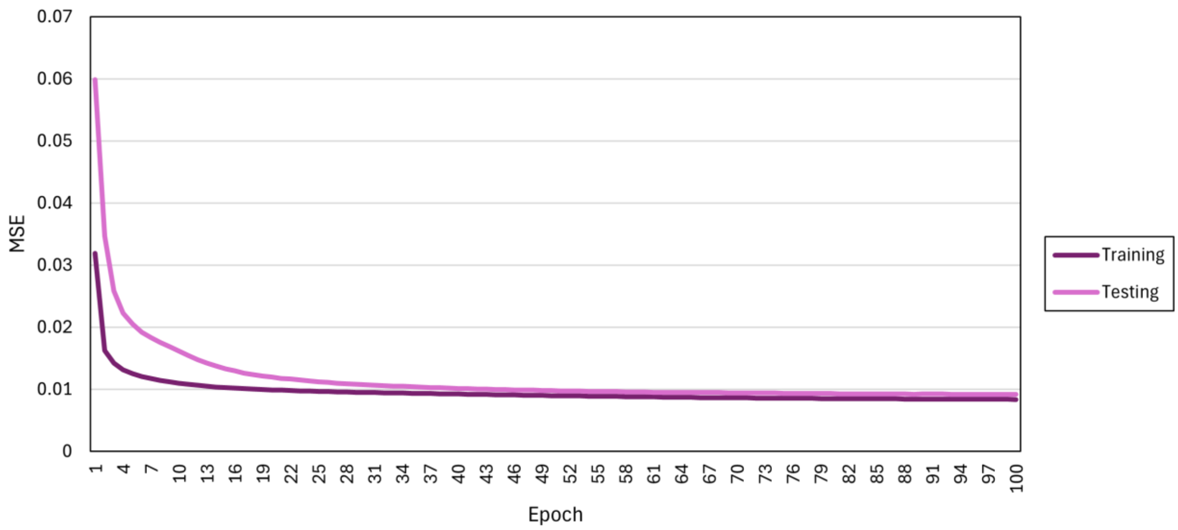

Figure 6.3: Prediction accuracy of the ANN per epoch of the artificial training dataset

The predictions made from the ANN with an artificial dataset could explain 86% of the variation of the actual data (Figure 6.4), even with a moderate amount of noise, demonstrating that the developed models are able to accurately identify the relationship between the inputs and output with the given parameters and dataset size. Knowing that the execution of the model was correct, the investigation instead focused on how the data was collected and the included input variables, to attempt to improve the ANN's accuracy.

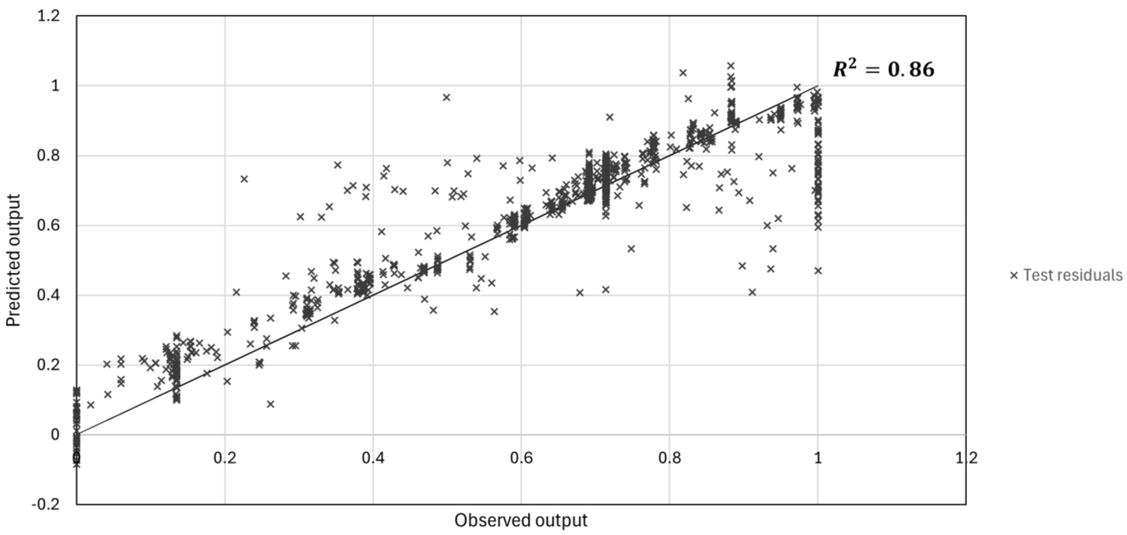

Figure 6.4: Residuals for the artificial testing dataset

6.4 Investigating the ANN's prediction accuracy

6.4.1 Identification of tamping interventions

Previously, it was assumed that unexplained improvements of the track geometry quality were due to unrecorded tamping interventions, and artificial maintenance actions were created to explain this behaviour. Regression lines were fitted to the SD measurements of the vertical alignment between the recorded and assumed tamping events (described further in section 4.3.2), and the ANN was used to predict the gradient of the fitted lines. The high-output tamping history of the ballast is widely considered an important factor for influencing the degradation rate and is included as an input parameter in the ANN; hence it is vital to correctly identify when tamping was completed. To try improve the model's prediction ability, four strategies to recognise tamping interventions were considered with a separate dataset created for each:

- BD1 – a dataset of regression lines fitted to geometry measurements between recorded and artificial tamping interventions.
- BD2 – a dataset of regression lines fitted to geometry measurements between only recorded tamping interventions. If an artificial tamping intervention occurred within that interval, the fitted line was removed from the study.
- BD3 – a dataset of regression lines fitted to geometry measurements between artificial tamping interventions, which were classified according to Equation 4.1. This can include officially recorded tamping events but ignored those which appeared to have no significant impact on ballast quality, in case of incorrectly reported interventions.
- BD4 – it was observed that the ballast's degradation behaviour did not appear to change after an artificial tamping intervention was created, as would be expected. This could suggest a potential issue with the calibration of the measurement equipment, and risks incorrectly labelling data such as in BD1 and BD3, or excluding viable data such as in BD2. The BD4 dataset was therefore comprised of the average regression line between recorded tamping interventions. This was completed using the BD1 dataset, with any fitted line between a recorded tamping interval used to find the average gradient.

6.4.2 Simplifying the model

The poor performance of the ANN may also result from the high variation of the outputs (the gradient of the fitted regression lines), demonstrated in Table 4.4, which can impact the ability to establish a trend in the data. Considerable effort was made to simplify the dataset and reduce the variation, including grouping, averaging, and incorporating additional variables. Four models were applied to each dataset:

- M1 – The model predicted the gradient of each fitted regression line.
- M2 – The model predicted the average gradient of the regression lines fitted to a 200m ballast unit.
- M3 – The model predicted the gradient of each fitted regression line within a tamping subgroup (e.g. all regression lines fitted after 1 tamp had been completed but before 2 tamps)
- M4 – The model predicted the gradient of the fitted regression lines, using the gradient of the line fitted to measurements between installation and the first recorded tamp on the ballast section as an input.

The considered variables per model (where B_{13} refers to the initial gradient, and n is the number of high-output tamps completed over the unit) were as detailed in Table 6.5:

Model	**Variables**
M1	Inputs: $B_1, B_2, B_3, B_4, B_5, B_6, B_7, B_8, B_9, B_{10}, B_{11}$
	Output: B_{12}
M2	Inputs: $B_1, B_2, B_3, B_4, B_5, B_6, B_7, B_9$
	Output: $\frac{1}{n}\sum_0^n B_{12}$
M3	Inputs: $B_1, B_2, B_3, B_4, B_5, B_6, B_7, B_8, B_9, B_{11}$
	Output: B_{12}
M4	Inputs: $B_1, B_2, B_3, B_4, B_5, B_6, B_7, B_8, B_9, B_{10}, B_{11}, B_{13}$
	Output: B_{12}

Table 6.5: Variables included per model

6.4.3 Applying the ANN to the ballast

To exhaustively investigate whether the SD of the vertical alignment was correlated with any of the track features, each dataset (BD1, BD2, BD3, BD4) was split into training:testing datasets by a ratio of 70:30, and applied to each model (M1, M2, M3, M4) over a range of training parameters. Following the initial investigation of the BD1 dataset implemented into M1 in section 6.3.3, a single hidden layer was used in the structure, with the number of hidden neurons in the layer varied between

$$\left[\frac{n}{2}, 2n\right] \tag{6.5}$$

where n is the number of input variables

The activation function remained leaky ReLU to process information entering the hidden layer, with Linear used for the output layer. A range of learning rates [0.001, 0.005, 0.01], and two cost functions (MAE and MSE) were used to investigate the model's performance, trained over 100 epochs of the training data. Table 6.6 details the optimal combination of parameters for the predictive ability of each model across the four datasets, evaluated by its associated R^2 score.

	BD1		**BD2**		**BD3**		**BD4**	
M1	0.36	MAE	0.36	MAE	0.35	MSE	0.29	MAE
		0.005		0.005		0.005		0.001
		79		91		80		97
		12		19		19		18
M2	0.26	MAE	0.14	MSE	0.21	MAE	0.22	MAE
		0.01		0.01		0.01		0.005
		98		100		49		94
		15		13		16		16
M3 (Between installation and 1 tamp)	0.28	MSE	0.25	MSE	0.26	MAE	0.30	MAE
		0.005		0.005		0.005		0.01
		90		66		96		82
		17		17		17		16
M4	0.33	MAE	0.40	MAE	0.37	MSE	0.40	MAE
		0.01		0.01		0.01		0.01
		97		96		99		93
		12		23		9		23

Table 6.6: R^2 score given the optimal training parameters: cost function, learning rate, training epochs, and hidden neurons for models predicting the degradation rate for SD of the vertical alignment

Table 6.6 shows that the models were unable to find a significant correlation between the features of the HS1 line and the gradient (or average gradient) of linear regression lines fitted to the SD of the vertical alignment. The joint best performing model used the gradient of the regression line fitted to measurements between installation to the first tamping intervention as an input variable (M4), with the implemented dataset only recognising officially recorded tamping interventions (BD2). The associated R^2 value of 0.4 suggests the model has a weak ability to understand degradation trends, but plotting the residuals for the predictions made from the testing and training datasets show there is still a high disparity between what was observed and what was obtained from the ANN (Figure 6.5).

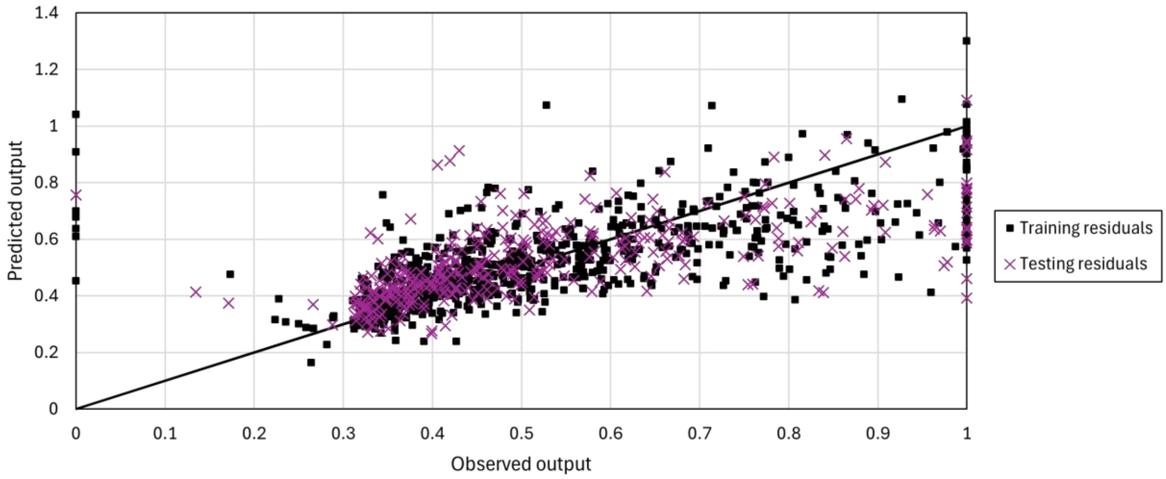

Figure 6.5: Plot of residuals for the BD2 training and testing datasets implemented in the M4 model

The models that considered the previous maintenance history as an input variable (M1 & M4) outperformed those that didn't, confirming that high-output tamping does have an impact on the degradation rate of the ballast. However, there was no clear best way to distinguish when a tamping intervention had occurred, with each of the four datasets varying in their prediction ability depending on what model was used. To improve the usefulness of the ANN, stricter recordings of interventions on the ballast should be employed, with further research needed if the dataset still shows unexplained improvements in quality. Table 6.6 highlights the mean absolute error as the most appropriate cost function for this research, as this was largely utilised within the best performing model compared to the MSE, most likely due to its robustness against outliers and the noisy nature of the datasets. The optimal learning rate was typically higher, with a rate of 0.001 only selected as the most appropriate once.

Finally, the optimal training epoch was often close to, or at, 100 which was the upper limit considered. More epochs give the model a greater chance to learn the underlying trends but increases its complexity. An investigation into the MSE of the M2 model implemented with the BD2 datasets (shown in Figure 6.6) revealed that including more training epochs hardly reduced the error from the testing dataset, which had largely converged after 100 epochs, but did have a negative impact on the computation time hence the upper epoch limit was kept to 100.

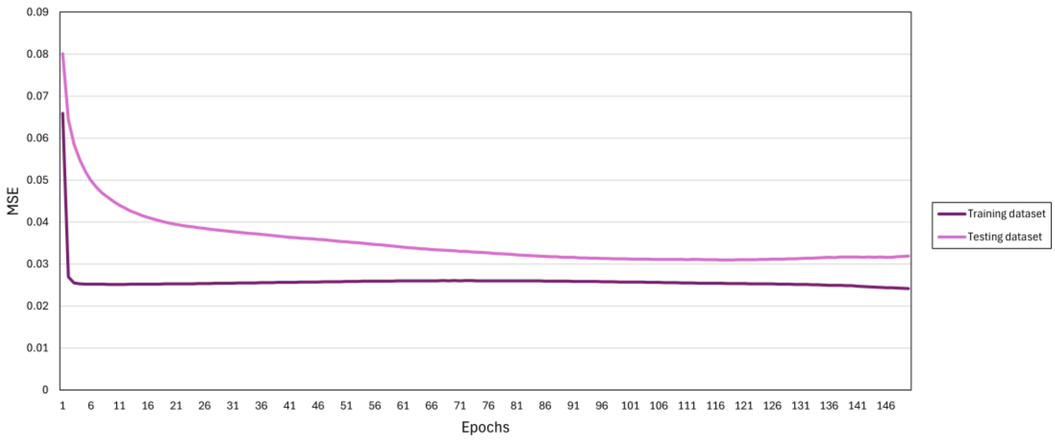

Figure 6.6: MSE of the training and testing BD2 datasets implemented in the M4 model

6.4.4 Applying the ANN to the plain line rail

A similar approach to sections 6.4.1-6.4.3 was applied to improve the ANN for the vertical wear on the plain line rail, but as there have been no restorative maintenance interventions, only a single linear regression line was fitted per plain line unit. As such, no data manipulations like grouping or averaging could be considered in an attempt to simplify the model. Nevertheless, as the fitted lines were related to measurements from a specific UTU vehicle (described further in section 4.3.3), this could be exploited to investigate if the method of data collection is causing the low prediction accuracy of the models. In total, three datasets were created:

- PLD1 – a dataset of regression lines fitted to wear measurements from the UTU vehicle with the most inspections over that unit.
- PLD2 – a dataset of regression lines fitted to wear measurements from the UTU vehicle with the highest gradient from that unit.
- PLD3 - a dataset of the average gradient of the regression lines fitted to wear measurements from each UTU vehicle per unit.

Identical parameters as those described in section 6.4.3 were used, with the optimal results shown in Table 6.7:

<table>
<tr><th colspan="2">PLD1</th><th colspan="2">PLD2</th><th colspan="2">PLD3</th></tr>
<tr><td rowspan="4">0.05</td><td>MAE</td><td rowspan="4">0.21</td><td>MSE</td><td rowspan="4">0.36</td><td>MSE</td></tr>
<tr><td>0.01</td><td>0.001</td><td>0.001</td></tr>
<tr><td>3</td><td>100</td><td>100</td></tr>
<tr><td>6</td><td>12</td><td>12</td></tr>
</table>

Table 6.7: R^2 score given the optimal training parameters: cost function, learning rate, training epochs, and hidden neurons for models predicting the degradation rate for vertical wear

Using the regression lines calculated from measurements taken by the UTU vehicle with the most inspections over the unit (PLD1) had a very low R^2 value (0.05), with the mean gradient calculated across all units being as good a prediction of future degradation rates as anything that can be obtained from the model. It is thought that finding the average gradient of an unit from inspections by all three UTU vehicles (PLD3) may have improved the model's performance as it reduces the impact of the known inaccuracies of the ultrasonic testing units for measuring vertical wear, with the PLD1 dataset suffering from the 'garbage in, garbage out' effect (Kilkenny & Robinson, 2018).

Predicting the degradation rate for the vertical wear on plain line rail units was most effective using an ANN with the mean squared error used as the cost function, a high number of training epochs, and a low learning rate. Despite the exhaustive approach to finding the optimal training parameters, and the proven ability of the model using an artificial dataset, both the ballast and plain line rail still had low R^2 values, with less than 40% of the variation of the data being explained by the developed ANN. It is clear that the methodology is currently not appropriate to be applied to the studied railway line, but additional research is needed to explain why.

6.5 Evaluating the model's performance given the behaviour of the HS1 line

The results in section 6.4. were reflective of the optimal combination of four training parameters – training epochs, cost function, number of hidden neurons, and learning rate – but there are further alterations of artificial neural networks that can be investigated, such as using an adaptive learning rate, or a different backpropagation algorithm. Given that other similar research has performed well with similar parameters, the R^2 score was found to improve by an average of less than 0.1 between the worst and best performing combination of parameters, and the implemented parameters were shown to be effective with an artificial dataset in section 6.3, it is believed that the high prediction errors result instead from the behaviour of the studied railway line.

6.5.1 Random behaviour within track subgroups

Varied behaviour within subgroups is evident when observing the linear regression lines fitted in Chapter 3. Figure 6.7 shows the development of the SD of the vertical alignment immediately after installation for three consecutive ballast units between 63.229-63.829km on the 'down' line of the track, all with identical features:

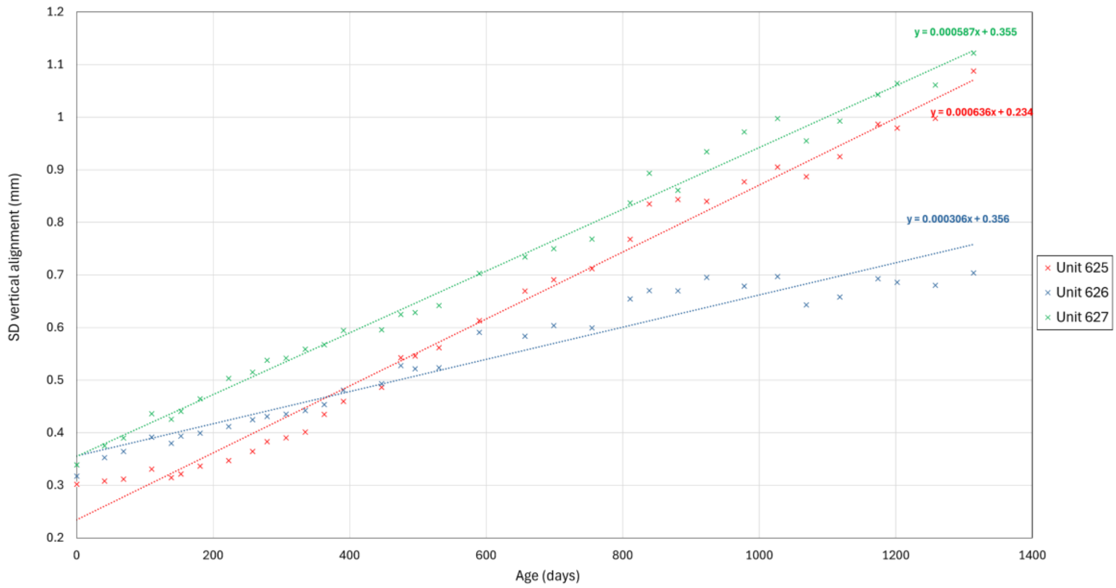

Figure 6.7: Fitted regression lines to SD vertical alignment measurements for consecutive units with identical features

Natural variation is expected but the fitted lines have contrasting gradients, with no apparent association to their shared features. The Pearson's correlation coefficient (Sedgwick, 2012) has frequently been used to measure the strength of the linear relationship between the variables of an ANN (Alqahtani & Ray, 2023; Kumar et al., 2022), however completing this analysis on the BD1 and PLD1 datasets showed at best a weak association between any of the track features and the gradients of the fitted regression lines (Table 6.8), but mostly no correlation given how close the values were to zero. Sensitivity analyses in Khajehei et al.'s (2022) and Zhang et al.'s (2025) research identified the previous maintenance history, initial degradation level, and traffic load as the most influential factors for ballast degradation, and the curve radius and superelevation (both of which were explained via the cant variable in this research) as the most influential for rail degradation, respectively. Their relative unimportance to the degradation rate for each asset on the HS1 line likely contributed to the model's inability to detect an underlying trend.

The lack of association between any track features and degradation was unexpected but may reflect the current high quality of the HS1 line. For example, the track is young, with the results found in Chapter 5 suggesting that the ballast and plain line rail were less than 30% into their operational period at the time of this research and expected to be in a 'good' condition at least 50% of the time. Retraining the models when the track has aged enough to distinguish consistently low performing areas, such as when restorative actions like renewals or ballast cleaning have been completed, may increase its performance.

Variable	Gradient (SD of the vertical alignment on the ballast)	Gradient (vertical wear on the plain line rail)
Speed	0.211	-0.092
Traffic	-0.029	0.313
Cant	-0.064	0.115
Tunnel	0.028	0.084
Viaduct	-0.033	-0.035
Elevation	-0.005	-0.052
Incline	0.022	-0.095
Age	0.003	0.108
Superstructure	-0.012	N/A
Substructure	N/A	0.124
Tamps	0.198	N/A
Intercept	0.344	N/A

Table 6.8: Applying the Pearson's correlation coefficient to the track features and gradients of the fitted regression lines

6.5.2 Omitted track features

An artificial neural network can only detect an underlying trend if the relevant parameters are included as input variables. The schematic and other relevant information have been used to identify as many track features as possible, but some influential degradation parameters may have been omitted – for example, a frequent area of research for wear on the railway line is focused on the wheel/rail interaction, with multiple studies focusing on how the wheelset design (Magel & Kalousek, 2017) or condition of the wheels (Biam et al., 2013) can have a subsequent impact on how quickly the rail degrades. Including information surrounding the wheelset as variables in the ANN may improve the model's prediction ability, but as no such information was available, this could not be considered. Additionally, despite Table 6.3 showing reasonable variation between the included track features, the HS1 line is still largely homogeneous compared to other railways as identical materials and rail design were used throughout. This homogeneity may potentially limit the extent to which variation in degradation behaviour can be detected.

6.5.3 Shared behaviour between track subgroups

Although section 6.5.1 found a low association between the behaviour of sections with identical features, a lot of the HS1 line still degraded similarly regardless of its characteristics. This can be seen in Figure 6.5, with a large cluster of residuals around 0.3-0.5 for the observed outputs. Similar issues were found in a paper written by Popov et al. (2022), who also applied the ballast dataset used here to an ANN. Although there are fundamental differences in the use of the model– geometry recordings across 50m segments beneath a single rail were used to classify whether the segment had deteriorated, improved, or not changed from two consecutive inspections- the study also encountered issues resulting from the uniformity of the data. Roughly 1500 training samples had to be removed as the track was largely in a good quality and the majority were classified as 'not changed', preventing the model from training adequately. Unfortunately, a similar approach could not be applied for this research as there was no input variable combination that was a clear majority, causing the clustering of data. The authors also stated that output processing was required due to the high sensitivity of the model applied to the testing data, noting that a dataset exhibiting more degradation may perform better.

Ultimately, it was concluded that the young age and excellent overall condition of the HS1 line meant that the ballast and plain line rail did not show any shared behaviour associated to their considered features, causing random degradation patterns between units. This issue, which was observed in other research using identical databases, caused the developed models to make unreliable predictions for future degradation events.

6.6 Summary

This chapter has explored utilising a different methodology to model the degradation of a high-speed track asset. Unlike a reliability analysis, an ANN can consider multiple factors influencing degradation at once and make predictions across different subgroups without the need for a completely new analysis or reducing the dataset. Once trained, an examination of the input variables can be used to highlight the features of the track that are expected to degrade faster and therefore inform how renewals should be scheduled in the most efficient manner. However, as demonstrated by the discussed results, ANN's are difficult to execute. They require a large dataset and an optimal combination of many parameters to train effectively, with their performance dependant on the quality of the dataset used to populate them. An exhaustive approach to developing and implementing an ANN found a weak ability to predict the degradation rate of a fault given the features of the track, but ultimately it was concluded that the HS1 line has not degraded enough for an underlying trend to be recognised. The current predictions made from the models have a lot of uncertainty and incorporating them back into the Petri nets is unlikely to help reduce the variation in lifetime predictions discussed at the beginning of this chapter.

Nevertheless, the work in this chapter has successfully developed a methodology for future application. This required the development of a track index, consideration of ways to produce the datasets, and creation of an algorithm to compute the ANN. The work in section 6.4 has highlighted the optimal parameters – such as the cost function, or training epochs – that maximises the model's prediction accuracy, leading to a more efficient training. Returning to the research when the track exhibits a higher degree of degradation, such as when it has aged or under a different maintenance strategy, means that the developed models can be quickly implemented to investigate whether there is now an established trend for the development of faults, leading to more accurate and specific renewal prediction.

Chapter 7

Conclusions

7.1 Summary

Rail travel is an increasingly popular method of transport, with over 60,000 trains timetabled across the HS1 line in 2023/24 (Office of Rail and Road, 2024). As the assets of the track age and wear, maintenance can be performed to rectify its condition but eventually a renewal is scheduled, which requires adequate funding to be provided when needed. However, if the expected frequency of renewal is not known there is a risk of overcharging railway operators to ensure funds are available which has a subsequent impact on the price passengers pay to use the service, hence accurate prediction of when a track component's life has been exhausted is vital. The aim of the research presented in this thesis is to predict when renewal is required by developing a model incorporating the inspection, degradation, maintenance and renewal processes of a track asset. Those considered are the crossing and switch panel of the S&C, the ballast, and the plain line rail.

Following a literature review to evaluate the most appropriate methodology for this research, an analysis of failure data obtained from the HS1 line was used to fit suitable continuous probability distributions and establish degradation trends. Then, a framework was developed to investigate the influence of the asset management strategy on the expected lifespan of a railway component, allowing for a renewal to be scheduled based on a combination of five criteria: a threshold on the amount of maintenance that could be completed, a limit on the amount of failures experienced within a set period of time, the length of time between successive maintenance interventions, a threshold on the cost of operations, and the condition of the asset unit after maintenance has been completed. Finally, the importance of the physical, geographical, and operational features of the track in relation to how each attribute can affect the degradation rate of a fault - and hence the predicted lifespan of an asset - was investigated. This research has made several contributions to the understanding of degradation and asset management modelling, and the developed models can be used as a decision-making tool for track engineers, assisting to amend the current renewal schedule in a way that maximises asset usefulness thus minimising the costs charged to the rail operators.

7.2 Reflection on the initial objectives

Chapter one identified seven objectives to achieve the aim. This section will summarise and evaluate how the work completed for this thesis has accomplished each objective:

1. *Review the surrounding literature regarding existing models that predict track degradation and asset management*

   In Chapter two, an extensive literature review was completed by assessing the current degradation and asset management models for track components. By appraising over 30 studies, the review was able to assess the current knowledge surrounding track asset renewals and identify areas of the implemented methodologies that aided in making lifespan predictions. The review also highlighted limitations in the current literature, namely the need for a developed degradation model to consider every different fault type and severity impacting the track, and for an asset management model to include a range of renewal criteria based on the condition of the asset, that can be implemented simultaneously. The exclusions of these factors can cause overly simplistic models and unrealistic renewal predictions, creating a clear need for a new approach.

2. *Decide on an appropriate methodology to achieve the aim*

   During the literature review, the considered research for the degradation models was derived from one of three methods – deterministic models, reliability-based models, and artificial neural networks - and the asset management models were implemented using either Markov or Petri net models. The strengths and limitations of each methodology were evaluated. From this, the structure of the research was developed, as it was understood that deterministic models were too limited to represent the stochastic behaviour of degradation, and the accompanying limitations of Markov models meant that they were deemed unsuitable for the complexities of asset management processes. The evaluation identified Petri nets populated with results from either reliability-based models or ANN as an appropriate methodology to predict when a renewal is necessary.

3. *Understand the types of degradation and resulting maintenance actions defined by the railway standards*

   Following a thorough scrutiny of the provided railway standards, the various failure types and corresponding maintenance were defined in Table 4.2 and Table 5.1. Five geography parameters were identified for measuring the condition of the ballast, but because a single ballast fault can be associated with multiple parameters, the vertical alignment was chosen as an overall representation of the ballast condition. This examination of the standards was used to inform the structure for the asset management models in Chapter five.

4. *Process and clean the historical data from the HS1 line, and use it to predict future degradation events*

   In Chapter four, the maintenance history obtained from physical inspections of the asset was cleaned to remove repeat observations and note right censored failures, creating a timeline of degradation for each unit of an asset type. The condition history, collected by regular monitoring by the TRV and UTU vehicles, had a linear regression line fitted to the measurements, which was extrapolated to find the time when a failure threshold was met. Samples were created for the times-to-failure (measured in days) of each fault. Maximum likelihood estimation and the AIC and BIC were used to evaluate which continuous probability distribution was the most appropriate to explain the behaviour of the sample, which could be used to predict future degradation events and inform of when failure is expected relative to the life stage of an asset. This analysis was a novel approach to predicting degradation as it considered every failure type and fault that could impact the railway, but it was noted that the small sample sizes for many faults caused high uncertainty in the fitted distributions. Additionally, it was observed that the UTU vehicles were flawed when measuring the profile of the rail, meaning that only the impact of vertical wear could be effectively examined.

5. *Establish the factors that influence the wear of the considered assets with use, such as the number of train journeys passing over the asset per annum*

   In Chapter six, an artificial neural network was created to consider the importance of each track feature when predicting the expected degradation rate of the vertical wear of the plain line rail and standard deviation of the vertical alignment of the ballast. The different factors were identified from a developed track index and used as input values in the ANN. After an exhaustive approach to finding the optimal training parameters, the predictive ability of the model was still poor, and an analysis of the dataset concluded that the line was too young to establish the relationship between different factors and degradation. The Pearson's correlation coefficient was able to identify that the speed, previous maintenance history, and intercept of the regression line had the strongest linear correlation to ballast degradation, and the traffic and cant influenced rail wear the most (Table 6.8), but these associations were still very weak.

6. *Develop a model that can incorporate asset management decisions, and populate with the degradation predictions to predict when a renewal is necessary*

   Chapter five introduced a methodology for the asset management modelling of the considered track components via a Petri net. Separate sub-models were created for the inspection, degradation, maintenance, and renewal processes, with an individual degradation sub-model developed for each failure type. To address the known limitations of previous models, the renewal sub-model allows for a renewal strategy to be implemented based on a combination of five criteria, with strategic placing of tokens around the net allowing for a total of 31 separate strategies to be considered. The firing times for the transitions relating to degradation was sampled from the distributions obtained from the reliability analysis, and the developed model was solved using Monte Carlo simulation. The ability of the methodology was demonstrated by exploring the predicted lifespan of an asset, with the results indicating that the switch panel typically reaches the renewal threshold faster than the crossing panel, and under the assumed parameters, the ballast and plain line rail were predicted to be operational for significantly longer than their expected lifetime stated at installation. The PN's were quick to run and computationally inexpensive, but several assumptions were required (described further in section 5.2.6) to make the model functional that could limit its ability to accurately represent asset management, particularly the assumption that maintenance is always 100% effective.

7. *Investigate the different renewal criteria to evaluate their impact on the lifespan and condition of an asset*

   The capabilities of the Petri net models were demonstrated during the latter half of Chapter five. Tables 5.11, 5.15, 5.21, and 5.23 describe the expected lifespan of an asset, given a renewal strategy with each criterion considered in isolation. The results in these tables describe how the Associated Cost criterion typically led to lifespan predictions with a small standard deviation, and plotting the condition over time for an asset under this renewal strategy revealed they were expected to remain in a 'good' condition at least 50% of the time. In section 5.4.3.4, the impact of changing the renewal thresholds within a criterion were investigated for the plain line rail, finding that the lifespan of the rail was positively correlated with the permitted depth of vertical wear. Additionally, the benefits of considering multiple criteria within a single renewal strategy was discussed for the crossing panel in section 5.4.1.4, finding that implementing all five criteria simultaneously had a negative impact on the expected lifespan but can be used to indicate which measurement of wear is driving the replacements. The objective was successfully achieved by manually adapting the PN, such as changing the multiplicity of an arc or the location of tokens, to investigate the impact on the simulated results, but it is acknowledged this could be streamlined by implementing a genetic algorithm. This would investigate a large search space of renewal criteria and thresholds almost instantaneously, to find the renewal strategies that optimise variables such as the amount spent operating the asset, expected lifespan, and time spent in an undesirable condition. The benefits of utilising the genetic algorithm for this objective is it can quickly inform track engineers of a need to adapt the current asset management strategy, helping to provide a more efficient service and potentially reduce costs passed on to the passengers. However, its 'black box' approach, that removes the ability to apply context during calculation, means that the proposed optimal renewal strategies may be unsuitable or impractical to implement. Additionally, the genetic algorithm can suffer from premature convergence (Katoch et al., 2021), potentially resulting in a limited investigation of the different renewal criteria compared to the longer but more rigorous approach of manually changing the PN and analysing the results.

7.3 Limitations and future work

Upon completion of this thesis there were several limitations and areas of future work identified:

- **Assuming the maintenance restores to 'as good as new'** – It was necessary to assume that maintenance is always 100% effective due to the limited available data, however further research is needed to investigate the likelihood of degradation after repairs, and how this may influence when the renewal criteria is met.

- **No consideration of the interactions between faults –** Currently, there is no consideration of whether the occurrence of one fault can impact the risk of another fault developing. Doing so may lead to a more efficient management strategy that prolongs the lifespan of the rail by removing potential risk factors.

- **Using an older dataset** - The limitations of the current dataset have been emphasized throughout, as in many cases the observation period extended across less than half of an assets expected lifetime. This is particularly evident by the lack of the expected 'wear out' seen across all four components, which may impact the time that the renewal criteria thresholds are met. Using a failure database from an older track that has experienced at least one full renewal could improve the accuracy of the lifespan predictions. However, a larger dataset can still be limited if the older failure data is not meticulously recorded in the manner described by the railway standards used to build the structure of the PNs, as the lifespan predictions rely heavily upon the modelled behaviour of individual, defined faults. Additionally, the results found from an older track may not be representative of high-speed conditions which are typically newer railways kept to a higher standard.

- **Developing an entire system model –** Currently, each asset unit is considered in isolation, but further research should be centred around creating an entire system model that considers the behaviour of multiple track components concurrently. This would allow for investigation into the interactions between assets – for example, if a degraded ballast causes an increase in defects on the rail - and the implementation of opportunistic renewals, with a focus on how this impacts the expected lifespan.

- **Applying an optimisation technique –** the genetic algorithm optimisation technique was briefly described in Chapter five, and a similar approach could help to expand on the final objective of this research, by applying the technique to find the renewal strategies that best maximise the lifespan whilst minimising the associated costs and length of time the railway spends in an unsuitable condition.

References

Ahsanullah, M., & Hamedani, G. (2010). *Exponential Distribution - Theory and Methods*. (1st ed.). Nova Science Publishers, Incorporated.

Ait Ali, A., Odolinski, K., Pålsson, B., & Tortensson, P. (2024). Evaluating the mix of maintenance activities on railway crossings with respect to life-cycle costs. *European Journal of Transport and Infrastructure Research, 24*(1), 1-29.

Alqahtani, H., & Ray, A. (2023). Feature selection of surface topography parameters for fatigue-damage detection using pearson correction method and neural network analysis. *Fatigue & Fracture of Engineering Materials & Structures, 46*(5), 1810-1820.

Alwosheel, A., Van Cranenburgh, S., & Chorus, C. (2018). Is your dataset big enough? Sample size requirements when using artificial neural networks for discrete choice analysis. *Journal of choice modelling, 28*, 167-182.

Andrade, A., & Teixeira, P. (2011). Uncertainty in rail-track geometry degradation: Lisbon-Oporto line case study. *Journal of transportation engineering, 137*(3), 193-200.

Andrews, J. (2013). A modelling approach to railway track asset management. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 227*(1), 56-73.

Andrews, J., Prescott, D., & Rozières, F. (2014). A stochastic model for railway track asset management. *Reliability Engineering & System Safety, 130*, 76-84.

Audley, M. (2014). *Rail Track Geometry Degradation and Maintenance Decision Making*. [PhD Thesis, University of Nottingham]. Nottingham eTheses.

Audley, M., & Andrews, J. (2013a). The effects of tamping on railway track geometry degradation. *Proceedings of the Institution of Mechanical Engineering Part F: Journal of Rail and Rapid Transit, 227*(4), 376-391.

Audley, M., & Andrews, J. (2013b). A Petri-Net Modelling Approach to Rail Track Geometry Maintenance and Inspection. In L. Jackson, & J. Andrews (Eds.), *Proceedings of AR2TS Seminar: Advances in Risk and Reliability Technology Symposium* (pp. 230-243). Loughborough University.

Aursudkij, B. (2007). *A laboratory study of railway ballast behaviour under traffic loading and tamping maintenance.* [PhD Thesis, University of Nottingham]. Nottingham eTheses.

Bai, W., Liu, R., An, R., Wang, F., & Sun, Q. (2019). Renewal management framework for urban rail transit assets. *Transport, 34*(1), 9-18.

Bakar, N., & Tahir, I. (2009). Applying multiple linear regression and neural network to predict bank performance. *International Business Research, 2*(4), 176-183.

Biam, J., Gu, Y., & Murray, M. (2013). A dynamic wheel-rail impact analysis of railway track under wheel flat by finite element analysis. *Vehicle System Dynamics, 51*(6), 784-797.

Bryant, L., & Andrews, J. (2015). Modelling railway bridge degradation based on historical maintenance data. *Safety and Reliability, 35*(2), 32-55.

Cannon, D., Edel, K., Grassie, S., & Sawley, K. (2003). Rail defects: an overview. *Fatigue & Fracture of Engineering Materials & Structure, 26*(10), 865-886.

Chattopadhyay, G., & Kumar, S. (2009). Parameter estimation for rail degradation model. *International Journal of Performability Engineering , 5*(2), 119.

Christensen, S., & Hansen, N. (1993). Coloured Petri nets extended with place capabilities, test arcs and inhibitor arcs. In M. Ajmone Marson (Eds.), *Application and Theory of Petri Nets 1993* (pp. 186-205). Springer Berlin Heidelberg.

Cornish, A. (2014). *Life-time monitoring of in service switches and crossings through field experimentation.* [PhD Thesis, Imperial College London]. Spiral.

Cornish, A., Smith, R., & Dear, J. (2016). Monitoring of strain of in-service railway switch rails through field experimentation. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 230*(5), 1429-1439.

Falamarzi, A., Moridpour, S., & Nazem, M. (2019). A review of rail track degradation prediction models. *Australian Journal of Civil Engineering, 17*(2), 152-166.

Fecarotti, C., & Andrews, J. (2020). A Petri net approach to assess the effects of railway maintenance on track availability. *Infrastructure Asset Management, 7*(3), 201-220.

Ginos, B. (2009). *Parameter Estimation for the Lognormal Distribution.* [Master's Dissertation, Brigham Young University]. BYU ScholarsArchive.

Goodarzi, S., Kashani, H., Saeedi, A., Oke, J., & Ho, C. (2023). Stochastic analysis for estimating track geometry degradation rates based on GPR and LiDAR data. *Construction and Building Materials, 369*, 130591.

Goutsias, J., & Jenkinson, G. (2013). Markovian dynamics on complex reaction networks. *Physics Reports, 529*, 199-264.

Grossoni, I., Hughes, P., Bezin, Y., Bevan, A., & Jaiswal, J. (2021). Observed failures at railway turnouts: Failure analysis, possible causes and links to current and future research. *Engineering Failure Analysis , 119*, 104987.

Guler, H. (2013). Decision support system for railway track maintenance and renewal management. *Journal of Computing in Civil Engineering, 27*(3), 292-306.

Guler, H. (2014). Prediction of railway track geometry deterioration using artificial neural networks: a case study for Turkish state railways. *Structure and Infrastructure Engineering, 10*(5), 614-626.

Guler, H., Jovanovic, S., & Evren, G. (2011). Modelling railway track geometry deterioration. *Proceedings of the Institution of Civil Engineers-Transport, 164*(2), 65-75.

Haji, S., & Abdulazeez, A. (2021). Comparison of optimisation techniques based on gradient descent algorithm: A review. *PalArch's Journal of Archaeology of Egypt/Egyptology, 18*(4), 2715-2743.

Hamadache, M., Dutta, S., Olaby, O., Ambur, R., Stewart, E., & Dixon, R. (2019). On the fault detection and diagnosis of railway switch and crossing systems: An overview. *Applied Sciences, 9*(23), 5129.

He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision* (pp. 1026-1034). IEEE.

He, Q., Li, H., Bhattacharjya, D., Parikh, D., & Hampapur, A. (2015). Track geometry defect rectification based on track deterioration modelling and derailment risk assessment. *Journal of the Operational Research Society, 66*(3), 392-404.

Hokstad, P., Langseth, H., Lindqvist, B., & Vatn, J. (2005). Failure Modeling and Maintenance Optimization for a Railway Line. *International Journal of Performability Engineering, 1*(1), 51-64.

Ibe, O. (2013). *Markov Processes for Stochastic Modeling.* (2nd ed.). Elsevier.

Indrayan, A., & Mishra, A. (2021). The importance of small samples in medical research. *Journal of Postgraduate Medicine, 67*(4), 219-223.

Jalili Hassankiadeh, S. (2011). *Failure analysis of railway switches and crossings for the purpose of preventative maintenance.* [Master Degree Project, KTH Royal Institution of Technology]. DiVA portal.

Jun, H., Seo, J., Jeon, I., Lee, S., & Chang, Y. (2016). Fracture and fatigue crack growth analyses on a weld-repaired railway rail. *Engineering failure analysis, 59*, 478-492.

Kassa, E., Andersson, C., & Nielsen, J. (2006). Simulation of dynamic interaction between train and railway turnout. *Vehicle system dynamics, 44*(3), 247-258.

Katoch, S., Chauhan, S., & Kumar, V. (2021). A review on genetic algorithm: past, present, and future. *Multimedia tools and applications, 80*(5), 8091-8126.

Khajehei, H., Ahmadi, A., Soleimanmeigouni, I., & Nissen, A. (2019). Allocation of effective maintenance limit for railway track geometry. *Structure and Infrastructure Engineering, 15*(12), 1597-1612.

Khajehei, H., Ahmadi, A., Soleimanmeigouni, I., Haddadzade, M., Nissen, A., & Latifi Jebelli, M. (2022). Prediction of rail track geometry degradation using artificial neural network: a case study. *International Journal of Rail Transportation, 10*(1), 24-43.

Kilkenny, M., & Robinson, K. (2018). Data quality: "Garbage in-garbage out". *Health Information Management Journal, 47*(3), 103-105.

Kittaneh, O., Ammach, S., Helal, S., Barkat, E., & Majid, M. (2022). Comparitive Study Between Lognormal and Weibull Distributions in Modeling Commercial Concentrator III-V Triple Junction Solar Cells Lifetimes. *International Journal of Renewable Energy Research, 12*(1), 547-556.

Kramer, O. (2017). *Genetic Algorithm Essentials.* (1st ed.). Springer International Publishing AG.

Kuha, J. (2004). AIC and BIC: Comparisons of assumptions and performance. *Sociological methods & research, 33*(2), 188-229.

Kumar, A., Yin, B., Shaikh, A., Ali, M., & Wei, W. (2022). CorrNet: pearson correlation based pruning for efficient convolutional neural networks. *International Journal of Machine Learning and Cybernetics, 13*(12), 3773-3783.

Kumar, S. (2006). *A study of the rail degradation process to predict rail breaks.* [Licentiate Thesis, Luleå University of Technology]. DiVA portal.

Kumari, J. (2019). *Asset Management of Railway Tracks Using Stochastic Petri Nets.* [Master Degree Project, Luleå University of Technology]. DiVA portal.

Lewis, R. (2011). Track Geometry Recording and Usage. Notes for a lecture to Network Rail.

Li, J., Cheng, K., Wang, S., Morstatter, F., Trevino, R., Tang, J., & Liu, H. (2017). Feature selection: A data perspective. *ACM computing surveys (CSUR), 50*(6), 1-45.

Li, Z. (2009). Squats on railway rails. In R. Lewis, & R. Olofsson (Eds.), *Wheel-rail interface handbook* (pp. 409-436). Woodhead Publishing Limited.

Litherland, J., & Andrews, J. (2023a). A reliability study of railway switch and crossing components. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 237*(2), 205-217.

Litherland, J., & Andrews, J. (2023b). A Petri net asset management framework for railway switches and crossings. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 237*(3), 322-334.

Litherland, J., & Andrews, J. (2025). System Petri net modelling for railway infrastructure asset management. *Infrastructure Asset Management, 12*(2), 110-128.

London St Pancras Highspeed. (2025). *Our Highspeed Route [Online Image].* Retrieved from https://stpancras-highspeed.com/highspeed-stations/

London St. Pancras Highspeed (Formerly High Speed One). (2024). *HS1 Asset Management Annual Statement 2023/24.* Retrieved from https://stpancras-highspeed.com/wp-content/uploads/2025/01/2023-24-final-amas-040624.pdf

Lyngby, N., Hokstad, P., & Vatn, J. (2008). RAMS Management of Railway Tracks. In K. Misra (Eds.), *Handbook of Performability Engineering* (pp. 1123-1145). Springer.

Magel, E., & Kalousek, J. (2017). Designing and assessing wheel/rail profiles for improved rolling contact fatigue and wear performance. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 231*(7), 805-818.

Marsaglia, G. (2004). Evaluating the normal distribution. *Journal of Statistical Software, 11*(4), 1-11.

Michas, G. (2012). *Slab Track Systems for High-Speed Railways.* [Master Degree Project, KTH Royal Institute of Technology]. DiVA portal.

Mooney, C. (1997). *Monte Carlo simulation.* SAGE.

Network Rail (High Speed) Ltd. (2020a). *Inspection, maintenance and repair procedures for cast, welded and fabricated crossings in the track (C-03-TK-35-2012).* (Internal use only - copies may be available upon request to the relevant controlled publication distributor).

Network Rail (High Speed) Ltd. (2020b). *Inspection and repair to control the risk of derailment at switches (C-03-TK-35-2009).* (Internal use only - copies may be available upon request to the relevant controlled publication distributor).

Network Rail (High Speed) Ltd. (2020c). *Track Geometry, Gauging, and Clearances (C-03-TK-35-2004).* (Internal use only - copies may be available upon request to the relevant controlled publication distributor).

Network Rail (High Speed) Ltd. (2020d). *Rail Management (C-03-TK-35-2003).* (Internal use only - copies may be available upon request to the relevant controlled publication distributor).

Ngamkhanong, C., Kaewunruen, S., & Baniotopoulos, C. (2021). Influences of ballast degradation on railway track buckling. *Engineering Failure Analysis, 122*, 105252.

Nguyen, Q., Ly, H., Ho, L., Al-Ansari, N., Le, H., Tran, V., . . . Pham, B. (2021). Influence of data splitting on performance of machine learning models in prediction of shear strength of soil. *Mathematical Problems in Engineering, 2021*(1), 4832864.

Nielsen, M. (2011). *Parameter Estimation for the Two-Parameter Weibull Distribution.* [Master's Dissertation, Brigham Young University]. BYU ScholarsArchive.

Novak, P., Danhel, M., Blazek, R., Kohlik, M., & Kubatova, H. (2017). Predicting the life expectancy of railway fail-safe signalling systems using dynamic models with censoring. In *2017 IEEE International Conference on Software Quality, Reliability and Security (QRS)* (pp. 329-339). IEEE.

Office of Rail and Road. (2021). *Passenger Rail Usage 2021-22 Quarter 1*. Retrieved from https://dataportal.orr.gov.uk/media/2010/passenger-rail-usage-2021-22-q1.pdf

Office of Rail and Road. (2024). *Annual report on HS1 Ltd 2023 to 2024.* Retrieved from https://www.orr.gov.uk/annual-report-hs1-ltd-2023-2024/asset-management

Popov, K., De Bold, R., Chai, H., Forde, M., Ho, C., Hyslip, J., . . . Hsu, S. (2022). Big-data driven assessment of railway track and maintenance efficiency using Artificial Neural Networks. *Construction and Building Materials, 349*, 128786.

Prescott, D., & Andrews, J. (2015). Investigating railway track asset management using a Markov analysis. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 229*(4), 402-416.

Rama, D., & Andrews, J. (2013). A reliability analysis of railway switches. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of rail and rapid transit, 227*(4), 344-363.

Reisig, W. (2012). *A primer in Petri net design.* Springer Science & Business Media.

ReliaSoft. (2020). *Weibull++ Software*. (Version 2020) [Computer software]. HBK.

Rempelos, G., Ognibene, G., Le Pen, L., Blainey, S., Preston, J., & Powrie, W. (2024). Railway Track Deterioration Models: A Review of the State of the Art. *Transportation Geotechnics*, 101377.

Robinson, S. (1970). A short proof of Cramer's rule. *Mathematics Magazine, 43*(2), 94-95.

Sachan, S., & Donchak, N. (2020). Generalized stochastic Petri-net algorithm with fuzzy parameters to evaluate infrastructure asset management policy. In *IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)* (pp. 1-8). IEEE.

Sadeghi, J., & Askarinejad, H. (2012). Application of neural networks in evaluation of railway track quality condition. *Journal of mechanical science and technology, 26*(1), 113-122.

Saleh, A., Prescott, D., Remenyte, R., & Chiachio, M. (2023). An optimized asset management petri net model for railway sections. In N. Rajic, W. Chiu, M. Veidt, A. Mita, & N. Takeda (Eds.), *Structural Health Monitoring: 9APWSHM* (Vol. 27, pp. 285-291). Materials Research Forum LLC.

Sato, Y. (1995). Japanese studies on deterioration of ballasted track. In J. Elkins (Eds.), *Interaction of Railway Vehicles with the Track and Its Substructure* (Vol. 24, pp. 197-208). Routledge.

Saunders, L., Russell, R., & Crabb, D. (2012). The coefficient of determination: what determines a useful R2 statistic? *Investigative opthalmology & visual science, 53*(11), 6830-6832.

Schneider, A., Hommel, G., & Blettner, M. (2010). Linear regression analysis: part 14 of a series on evaluation of scientific publications. *Deutsches Ärzteblatt International, 107*(44), 776-782.

Sedgwick, P. (2012). Pearson's correlation coefficient. *BMJ (Online), 345*(1), 4483.

Shafahi, Y., & Hakhamaneshi, R. (2009). Application of a maintenance management model for Iranian railways based on the Markov chain and probabilistic dynamic programming. *Scientia Iranica, 16*(1), 87-97.

Shafahi, Y., Masoudi, P., & Hakhamaneshi, R. (2008). Track degradation prediction models, using Markov Chain, artificial neural and neuro-fuzzy network. *8th World Congress on Railway Research: WCRR 2008*, (pp. 1-9).

Sharma, S., Sharma, S., & Athaiya, A. (2020). Activation Functions in Neural Networks. *International Journal of Engineering Applied Sciences and Technology, 4*(12), 310-316.

Singh, Y., Saini, M., & Savita. (2023). Impact and Performance Analysis of Various Activation Functions for Classification Problems. *2023 IEEE International Conference on Contemporary Computing and Communications (InC4). 1*, pp. 1-7. Bangalore: IEEE.

Sola, J., & Sevilla, J. (1997). Importance of Input Data Normalisation for the Application of Neural Networks to Complex Industrial Problems. *IEEE Transactions on Nuclear Science, 44*(3), 1464-1468.

Solano-Alvarez, W., & Bhadeshia, H. (2024). Steels for rails. *Progress in Materials Science, 146*, 101313.

Soleimanmeigouni, I., Ahmadi, A., Arasteh Khouy, I., & Letot, C. (2018). Evaluation of the effect of tamping on the track geometry condition: A case study. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 232*(2), 408-420.

Soleimanmeigouni, I., Ahmadi, A., Nissen, A., & Xiao, X. (2020). Prediction of railway track geometry defects: a case study. *Structure and Infrastructure Engineering, 16*(7), 987-1001.

Steenbergen, M. (2009). Rail welds. In R. Lewis, & U. Olofsson (Eds.), *Wheel-Rail Interface Handbook* (pp. 377-408). Woodhead Publishing Limited.

Sun, Y., Fidge, C., & Ma, L. (2011). Reliability prediction of long-lived linear assets with incomplete failure data. In *2011 International Conference on Quality, Reliability, Risk, Maintenance, and Safety Engineering* (pp. 143-147). IEEE.

Symonds, M., & Moussalli, A. (2011). A brief guide to model selection, multimodal inference and model averaging in behavioural ecology using Akaike's information criterion. *Behavioural ecology and sociobiology, 65*(1), 13-21.

Thurlby, R. (2013). Managing the asset time bomb: a system dynamics approach. *Proceedings of the Institution of Civil Engineers-Forensic Engineering, 166*(3), 134-142.

Veit, P., & Marschnig, S. (2010). Sustainability in Track: A Precondition for High Speed Traffic. In *ASME/IEEE Joint Rail Conference* (Vol. 2, pp. 349-355). IEEE.

Walfish, S. (2006). A review of statistical outlier methods. *Pharmaceutical technology, 30*(11), 82.

Weman, K. (2003). *Welding processes handbook*. Woodhead Publishing Limited.

Widyaningsih, Y., Rizka, H., & Siswantining, T. (2023). Performance comparison between maximum likelihood estimation and variational method for estimating simple linear regression parameter. In D. Aldila, Z. Zainal Abidin, M. Imran, & J. Widakdo (Eds.), *The 9th International Symposium on Current Progress in Mathematics and Sciences (The 9th ISCPMS 2023)* (Vol. 61, p. 01010). EDP Sciences.

Yilboga, H., Eker, O., Güçlü, A., & Camci, F. (2010). Failure prediction on railway turnouts using time delay neural networks. In *2010 IEEE International conference on computational intelligence for measurement systems and applications* (pp. 134-137). IEEE.

Yong, T. (2004). *Extended Weibull distributions in reliability engineering.* [PhD Thesis, National University of Singapore]. ScholarBank@NUS.

Zhang, D., Hu, H., & Roberts, C. (2017). Rail maintenance analysis using Petri nets. *Structure and Infrastructure Engineering, 13*(6), 783-793.

Zhang, H., Shuai, W., Liu, L., Zhang, P., Zhang, K., Lin, H., . . . Li, W. (2025). Prediction of Rail Wear Under Different Railway Track Geometries Using Artificial Neural Networks. *Infrastructures, 10*(7), 154.

Zhang, X., Zhao, C., Zhai, W., Shi, C., & Feng, Y. (2019). Investigation of track settlement and ballast degradation in the high-speed railway using a full-scale laboratory test. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 233*(8), 869-881.

Zhang, Y., Murray, M., & Ferreira, L. (2000). Modelling Rail Track Performance: An Integrated Approach. In *Proceedings of the Institution of Civil Engineers - Transport* (Vol. 141, pp. 187-194). Thomas Telford Limited.

Zhao, J., Chan, A., Roberts, C., & Stirling, A. (2006). Assessing the economic life of rail using a stochastic analysis of failures. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 220*(2), 103-111.

Zou, J., Han, Y., & So, S. (2008). Overview of Artificial Neural Networks. *Artificial Neural Networks, 458*, 14-22.

Zurawski, R., & Zhou, M. (1994). Petri nets and industrial applications: A tutorial. *IEEE Transactions on industrial electronics, 41*(6), 567-583.

Zwanenburg, W. (2006). Degradation Processes of Switches & Crossings. In *IET International Conference On Railway Condition Monitoring* (pp. 115-119). IET.

APPENDIX A

The structure for the models described in Chapter five are too large for graphical representation. Appendix A details the various aspects of the Petri net for each track asset, including the connections and firing time for each transition, initial marking of the net, and firing rules for place dependent transitions.

A.1 Switch panel

A.1.1 Transitions

The information for transitions in the switch panel Petri net, including the number of the transition, input and output places connecting to the transition with the weight of the associated arc, the firing time in days, which places the transition is inhibited by, and all places and transitions that are reset when the transition is fired are described below:

Sub-model type	Transition number	Input places (multiplicity)	Output places (multiplicity)	Firing time in days	Inhibited by	Resets
Inspection	T0	P0	P1, P16 (20)	182	-	-
	T1	P1	P0	1	-	-
Maintenance	T2	P2	P6, P8, P12, P16 (130)	0	-	-
	T3	P3	P6, P8, P12, P16 (200)	0	-	-
	T4	P4	P16 (40)	0	-	P4
	T5	P5	P16 (3)	0	-	P5
Renewal	T6	P6 (6)	P21	0	P7	-
	T7	P8 (2), P10	P10, P21	0	P11	-
	T8	P9	P10	720	-	-
	T9	P10	P9	1	-	P8
	T10	P12 (2)	P12, P14	0	-	T11
	T11	P12	P13	913	-	P12
	T12	P14 (2)	P21	0	P15	-
	T13	P16 (2000)	P21	0	P17	-
	T14	P18 (8)	P21	0	P20	-
	T15	P19 (5)	P21	0	P20	-
	T16	P21	P22	14	-	P0-P101, T0-T121
Degradation (Head worn stock rail)	T17	P23	P24	~E(0.00000166)	-	-
	T18	P24	P25	~E(1)	-	-
	T19	P1, P24	P1, P24, P26	0	P26	-
	T20	P1, P25	P1, P27	0	-	-
	T21	P26	P2, P18	91	P27	P23-P27, T17-T22
	T22	P27	P2, P18 (2)	14	-	P23-P27, T17-T22
Degradation (Wheelburn)	T23	P28	P29	~W(1.92, 21876)	-	-
	T24	P1, P29	P1, P30	0	-	-
	T25	P30	P31	91	-	P28-P30, T23-T25
	T26	P31	P2, P18	~U(0, 1)	-	-
	T27	P31	P3, P19	0.56	-	-
Degradation (Switch rail height relative to the stock rail)	T28	P32	P33	~E(0.0000142)	-	-
	T29	P33	P34	~E(0.00000172)	-	-
	T30	P1, P33	P1, P33, P35	0	P35	-
	T31	P1, P34	P1, P36	0	-	-
	T32	P35	P37	7	P36	P32-P36, T28-T33
	T33	P36	P37	2	-	P32-P36, T28-T33
	T34	P36	P4, P36	1	-	-
	T35	P37	P2, P18 (3)	~U(0, 1)	-	-
	T36	P37	P3, P19 (3)	0.56	-	-
Degradation (Vertical cracking on the switch rail)	T37	P38	P39	~E(0.00000166)	-	-
	T38	P1, P39	P1, P40	0	-	-
	T39	P40	P41	91	-	P38-P41, T37-T39
	T40	P41	P2, P18	~U(0, 1)	-	-
	T41	P41	P3, P19	0.56	-	-
Degradation	T42	P42	P43	~E(0.00000166)	-	-
	T43	P43	P44	~E(1)	-	-

(Sidewear on the stock rail)	T44	P44	P45	~E(1)	-	-
	T45	P1, P43	P1, P43, P46	0	P46	-
	T46	P1, P44	P1, P44, P47	0	P47	-
	T47	P1, P45	P1, P48	0	-	-
	T48	P47	P3, P5, P19	56	P48	P42-P48, T42-T50
	T49	P48	P22	7	-	P0-P101, T0-T121
	T50	P48	P4, P48	1	-	-
Degradation (Transverse cracking in RCF)	T51	P49	P50	~L(11.05, 2.69)	-	-
	T52	P1, P50	P1, P51	0	-	-
	T53	P51	P52	182	-	P49-P51, T51-T55
	T54	P52	P2, P18	~U(0, 1)	-	-
	T55	P52	P3, P19	0.56	-	-
Degradation (Lipping in the machined length of the switch or stock rail)	T56	P53	P54	~L(6.48, 1.29)	-	-
	T57	P1, P54	P1, P55	0	-	-
	T58	P55	P2, P18	91	-	P53-P55, T56-T58
Degradation (Lipping in the machined length of the switch or stock rail preventing switch toe closing)	T59	P56	P57	~E(0.00000166)	-	-
	T60	P1, P57	P1, P58	0	-	-
	T61	P58	P59	28	-	P56-P58, T59-T63
	T62	P59	P2, P18 (2)	~U(0, 1)	-	-
	T63	P59	P3, P19 (2)	0.56	-	-
Degradation (Wheel strike of switch toe)	T64	P60	P61	~E(0.00000166)	-	-
	T65	P1, P61	P1, P62	0	-	-
	T66	P62	P2, P18 (3)	7	-	P60-P62, T64-T66
Degradation (Ballast pitting)	T67	P63	P64	~L(26.32, 8.79)	-	-
	T68	P1, P64	P1, P65	0	-	-
	T69	P65	P66	91	-	P63-P65, T67-T69
	T70	P66	P2, P18	~U(0, 1)	-	-
	T71	P66	P3, P19	0.56	-	-
Degradation (Shallow flange contact angle)	T72	P67	P68	~W(0.53, 42416)	-	-
	T73	P68	P69	~W(2.28, 561054)	-	-
	T74	P1, P68	P1, P68, P70	0	P70	-
	T75	P1, P69	P1, P71	0	-	-
	T76	P70	P5, P72	91	P71	P67-P71, T72-T78
	T77	P71	P73	7	-	P67-P71, T72-T78
	T78	P70	P4, P70	1	-	-
	T79	P72	P2, P18	~U(0, 1)	-	-
	T80	P72	P3, P19	0.56	-	-
	T81	P73	P2, P18 (3)	~U(0, 1)	-	-
	T82	P73	P3, P19 (3)	0.56	-	-
Degradation **(**Defective weld repair)	T83	P74	P75	~L(12.99, 3.03)	-	-
	T84	P1, P75	P1, P76	0	-	-
	T85	P76	P77	91	-	P74-P76, T83-T85
	T86	P77	P2, P18	~U(0,1)	-	-
	T87	P77	P3, P19	0.56	-	-
	T88	P19	P19, P74	0	P74-P76	-
Degradation (Damage at the switch toe)	T89	P78	P79	~E(0.00000166)	-	-
	T90	P79	P80	~E(1)	-	-
	T91	P1, P79	P1, P79, P81	0	P81	-
	T92	P1, P80	P1, P82	0	-	-
	T93	P81	P3, P19 (3)	7	P82	P78-P82, T89-T94
	T94	P82	P3, P19 (3)	2	-	P78-P82, T89-T94
Degradation (Gauge 2 contacting at any damage associated with the switch toe)	T95	P83	P84	~E(0.00000166)	-	-
	T96	P84	P85	~E(1)	-	-
	T97	P1, P84	P1, P84, P86	0	P86	-
	T98	P1, P85	P1, P87	0	-	-
	T99	P86	P3,P5,P19 (3)	7	P87	P83-P87, T95-T101
	T100	P87	P3, P19 (3)	2	-	P83-P87, T95-T101
	T101	P87	P4, P87	1	-	-
Degradation (Gauge 2 contacts below any damage or crack)	T102	P88	P89	~E(0.000160)	-	-
	T103	P89	P90	~E(0.0000345)	-	-
	T104	P89	P91	~E(0.0000543)	-	-
	T105	P90	P91	~E(0.00000176)	-	-
	T106	P1, P89	P1, P89, P92	0	P92	-
	T107	P1, P90	P1, P90, P93	0	P93	-
	T108	P1, P91	P1, P94	0	-	-
	T109	P92	P3, P19	91	P93, P94	P88-P94, T102-T111
	T110	P93	P3, P19 (2)	28	P94	P88-P94, T102-T111
	T111	P94	P3, P19 (2)	28	-	P88-P94, T102-T111
Degradation (Gauge 2 contacts above any damage or crack)	T112	P95	P96	~E(0.00000166)	-	-
	T113	P96	P97	~E(1)	-	-
	T114	P96	P98	~E(1)	-	-
	T115	P97	P98	~E(1)	-	-
	T116	P1, P96	P1, P96, P99	0	P99	-
	T117	P1, P97	P1, P97, P100	0	P100	-
	T118	P1, P98	P1, P101	0	-	-

	T119	P99	P3, P19 (2)	28	P100, P101	P95-P101, T112-T121
	T120	P100	P3, P19 (3)	7	P101	P95-P101, T112-T121
	T121	P101	P3, P19 (3)	7	-	P95-P101, T112-T121

A.1.2 Initial marking

The initial number of tokens in each place of the switch panel Petri net is described below:

Initial marking value	**Place no.**
1	P0, P9, P23, P28, P32, P38, P42, P49, P53, P56, P60, P63, P67, P74, P78, P83, P88, P95
0	P1, P2, P3, P4, P5, P6, P7, P8, P10, P11, P12, P13, P14, P15, P16, P17, P18, P19, P20, P21, P22, P24, P25, P26, P27, P29, P30, P31, P33, P34, P35, P36, P37, P39, P40, P41, P43, P44, P45, P46, P47, P48, P50, P51, P52, P54, P55, P57, P58, P59, P61, P62, P64, P65, P66, P68, P69, P70, P71, P72, P73, P75, P76, P77, P79, P80, P81, P82, P84, P85, P86, P87, P89, P90, P91, P92, P93, P94, P96, P97, P98, P99, P100, P101

A.1.3 Place dependent transitions

The rules for the place dependent transitions of the switch panel are described below.

If a token is present in any of the 'Marked place' places, all the corresponding dependent transitions have their firing time modified to the time (in days) described in the last column of the table, for the duration of the time a token is present in the marked place. After a token is removed, the firing time returns to what was described in the *Transitions* table.

Marked place	**Dependent transition**	**Adjusted firing time (days)**
P51, P93, P94, P99	T0	7
P26, P47	T0	28
P46	T0	56
P36, P82, P87	T21, T22, T25, T32, T39, T48, T49, T53, T58, T61, T66, T69, T76, T77, T85, T93, T99, T109, T110, T111, T119, T120, T121	2
P35, P48, P62, P71, P81, P86, P100, P101	T21, T22, T25, T39, T48, T53, T58, T61, T69, T76, T85, T109, T110, T111, T119	7
P27	T21, T25, T39, T48, T53, T58, T61, T69, T76, T85, T109, T110, T111, T119	14
P58, P93, P94, P99	T21, T25, T39, T48, T53, T58, T69, T76, T85, T109	28
P47	T21, T25, T39, T53, T58, T69, T76, T85, T109	56
P26, P30, P40, P55, P65, P70, P76, P92	T53	91

A.2 Crossing panel

A.2.1 Transitions

Sub-model type	Transition number	Input places	Output places	Firing time (days)	Inhibited by	Resets
Inspection	T0	P0	P1, P16 (20)	182	-	-
	T1	P1	P0	1	-	-
Maintenance	T2	P2	P6, P8, P12, P16 (130)	0	-	-
	T3	P3	P6, P8, P12, P16 (200)	0	-	-
	T4	P4	P16 (15)	0	-	P4
	T5	P5	P16 (20)	0	-	P5
Renewal	T6	P6 (6)	P21	0	P7	-
	T7	P8 (2), P10	P10, P21	0	P11	-
	T8	P9	P10	720	-	-
	T9	P10	P9	1	-	P8
	T10	P12 (2)	P12, P14	0	-	-
	T11	P12	P13	913	-	P12
	T12	P14 (2)	P21	0	P15	-
	T13	P16 (2000)	P21	0	P17	-
	T14	P18 (8)	P21	0	P20	-
	T15	P19 (5)	P21	0	P20	-
	T16	P21	P22	14	-	P0-P198, T0-T275
Degradation (Longitudinal vertical cracking of casting)	T17	P23	P24	~E(0.00000166)	-	-
	T18	P24	P25	~E(1)	-	-
	T19	P1, P24	P1, P24, P26	0	P26	-
	T20	P1, P25	P1, P27	0	-	-
	T21	P26	P28	91	P27	P23-P27, T17-T22
	T22	P27	P28	28	-	P23-P27, T17-T22
	T23	P28	P2, P18	~U(0, 1)	-	-
	T24	P28	P3, P19	0.58	-	-
Degradation (Longitudinal crack in the apron)	T25	P29	P30	~E(0.00000166)	-	-
	T26	P30	P31	~E(1)	-	-
	T27	P1, P30	P1, P30, P32	0	P32	-
	T28	P1, P31	P1, P33	0	-	-
	T29	P32	P3, P19 (3)	7	P33	P29-P33, T25-T31
	T30	P33	P3, P4, P19 (3)	2	-	P29-P33, T25-T31
	T31	P33	P5, P33	1	-	-
Degradation (Defects associated with the TMZ)	T32	P34	P35	~E(0.00000166)	-	-
	T33	P35	P36	~E(1)	-	-
	T34	P1, P35	P1, P35, P37	0	P37	-
	T35	P1, P36	P1, P38	0	-	-
	T36	P37	P2, P18 (3)	7	P38	P34-P38, T32-T38
	T37	P38	P2, P18 (3)	2	-	P34-P38, T32-T38
	T38	P38	P5, P38	1	-	-
Degradation (Transverse cracking of casting in the crossing vee to a depth below the running surface)	T39	P39	P40	~W(1.46,17042)	-	-
	T40	P40	P41	~W(3.61, 577435)	-	-
	T41	P41	P42	~W(1, 1)	-	-
	T42	P42	P43	~W(1, 1)	-	-
	T43	P43	P44	~W(1, 1)	-	-
	T44	P1, P40	P1, P40, P45	0	P45	-
	T45	P1, P41	P1, P41, P46	0	P46	-
	T46	P1, P42	P1, P42, P47	0	P47	-
	T47	P1, P43	P1, P43, P48	0	P48	-
	T48	P1, P44	P1, P49	0	-	-
	T49	P45	P3, P19	91	P46-P49	P39-P49, T39-T54
	T50	P46	P3, P19 (2)	28	P47-P49	P39-P49, T39-T54
	T51	P47	P3, P19 (2)	14	P48, P49	P39-P49, T39-T54
	T52	P48	P3, P4, P19 (3)	2	P49	P39-P49, T39-T54
	T53	P49	P3, P4, P19 (3)	2	-	P39-P49, T39-T54
	T54	P49	P5, P49	1	-	-
Degradation (Transverse cracking of casting in the wing rail to a depth below the running surface)	T55	P50	P51	~E(0.0000150)	-	-
	T56	P51	P52	~E(0.00000187)	-	-
	T57	P52	P53	~E(1)	-	-
	T58	P53	P54	~E(1)	-	-
	T59	P1, P51	P1, P51, P55	0	P55	-
	T60	P1, P52	P1, P52, P56	0	P56	-
	T61	P1, P53	P1, P53, P57	0	P57	-
	T62	P1, P54	P1, P58	0	-	-
	T63	P55	P3, P19	91	P56-P58	P50-P58, T55-T67
	T64	P56	P3, P18 (2)	28	P57, P58	P50-P58, T55-T67
	T65	P57	P3, P18 (3)	7	P58	P50-P58, T55-T67
	T66	P58	P3, P18 (3)	2	-	P50-P58, T55-T67

	T67	P58	P5, P58	1	-	-
Degradation (Transverse cracking of a casting elsewhere in the box section from the running surface to a length)	T68	P59	P60	~E(0.00000166)	-	-
	T69	P60	P61	~E(1)	-	-
	T70	P61	P62	~E(1)	-	-
	T71	P1, P60	P1, P60, P63	0	63	-
	T72	P1, P61	P1, P61, P64	0	64	-
	T73	P1, P62	P1, P65	0	-	-
	T74	P63	P3, P19 (2)	28	P64, P65	P59-P65, T69-T76
	T75	P64	P3, P19 (3)	7	P65	P59-P65, T69-T76
	T76	P65	P3, P19 (3)	7	-	P59-P65, T69-T76
Degradation (Horizontal cracking in the upper or lower fillet radius of a casting extending to a position)	T77	P66	P67	~E(0.00000166)	-	-
	T78	P67	P68	~E(1)	-	-
	T79	P68	P69	~E(1)	-	-
	T80	P1, P67	P1, P67, P70	0	P70	-
	T81	P1, P68	P1, P68, P71	0	P71	-
	T82	P1, P69	P1, P72	0	-	-
	T83	P71	P4, P73	7	P72	P66-P72, T77-T86
	T84	P72	P4, P73	2	-	P66-P72, T77-T86
	T85	P71	P5, P71	1	-	-
	T86	P72	P5, P72	1	-	-
	T87	P73	P2, P18 (3)	~U(0,1)	-	-
	T88	P73	P3, P19 (3)	0.58	-	-
Degradation (Squat on a casting surface length)	T89	P74	P75	~L(8.26, 1.83)	-	-
	T90	P75	P76	~L(9.13, 2.23)	-	-
	T91	P76	P77	~L(14.86, 1.59)	-	-
	T92	P1, P75	P1, P75, P78	0	P78	-
	T93	P1, P76	P1, P76, P79	0	P79	-
	T94	P1, P77	P1, P80	0	-	-
	T95	P78	P3, P19	365	P79, P80	P74-P80, T89-T97
	T96	P79	P3, P19	182	P80	P74-P80, T89-T97
	T97	P80	P3, P19	91	-	P74-P80, T89-T97
Degradation (Horizontal cracking of the flangeway wall or wing rail)	T98	P81	P82	~E(0.00000166)	-	-
	T99	P82	P83	~E(1)	-	-
	T100	P83	P84	~E(1)	-	-
	T101	P1, P82	P1, P82, P85	0	P85	-
	T102	P1, P83	P1, P83, P86	0	P86	-
	T103	P1, P84	P1, P87	0	-	-
	T104	P85	P3, P19 (2)	14	P86, P87	P81-P87, T98-T106
	T105	P86	P3, P19 (3)	7	P87	P81-P87, T98-T106
	T106	P87	P3, P19 (3)	2	-	P81-P87, T98-T106
Degradation (Metal breaking out in the wheel transfer area)	T107	P88	P89	~E(0.00000166)	-	-
	T108	P89	P90	~E(1)	-	-
	T109	P90	P91	~E(1)	-	-
	T110	P91	P92	~E(1)	-	-
	T111	P92	P93	~E(1)	-	-
	T112	P93	P94	~E(1)	-	-
	T113	P89	P95	~E(1)	-	-
	T114	P90	P96	~E(1)	-	-
	T115	P91	P97	~E(1)	-	-
	T116	P92	P98	~E(1)	-	-
	T117	P93	P99	~E(1)	-	-
	T118	P94	P100	~E(1)	-	-
	T119	P95	P101	~E(1)	-	-
	T120	P96	P102	~E(1)	-	-
	T121	P97	P103	~E(1)	-	-
	T122	P98	P104	~E(1)	-	-
	T123	P99	P105	~E(1)	-	-
	T124	P100	P106	~E(1)	-	-
	T125	P1, P89	P1, P89, P107	0	P107	-
	T126	P1, P90	P1, P90, P108	0	P108	-
	T127	P1, P91	P1, P91, P109	0	P109	-
	T128	P1, P92	P1, P92, P110	0	P110	-
	T129	P1, P93	P1, P93, P111	0	P111	-
	T130	P1, P94	P1, P94, P112	0	P112	-
	T131	P1, P95	P1, P95, P113	0	P113	-
	T132	P1, P96	P1, P96, P114	0	P114	-
	T133	P1, P97	P1, P97, P115	0	P115	-
	T134	P1, P98	P1, P98, P116	0	P116	-
	T135	P1, P99	P1, P99, P117	0	P117	-
	T136	P1, P100	P1, P100, P118	0	P118	-
	T137	P1, P101	P1, P101, P119	0	P119	-
	T138	P1, P102	P1, P102, P120	0	P120	-
	T139	P1, P103	P1, P103, P121	0	P121	-
	T140	P1, P104	P1, P104, P122	0	P122	-
	T141	P1, P105	P1, P105, P123	0	P123	-
	T142	P1, P106	P1, P106, P124	0	P124	-

	T143	P107	P3, P19 (3)	7	P108-P124	P88-P124, T107-T169
	T144	P108	P3, P19 (3)	7	P109-P124	P88-P124, T107-T169
	T145	P109	P3, P19 (3)	7	P110-P124	P88-P124, T107-T169
	T146	P110	P3, P19 (3)	7	P111-P124	P88-P124, T107-T169
	T147	P111	P3, P19 (3)	7	P112-P124	P88-P124, T107-T169
	T148	P112	P3, P4, P19 (3)	2	P113-P124	P88-P124, T107-T169
	T149	P112	P5, P112	1	-	-
	T150	P113	P3, P19 (3)	7	P114-P124	P88-P124, T107-T169
	T151	P114	P3, P4, P19 (3)	2	P115-P124	P88-P124, T107-T169
	T152	P115	P3, P4, P19 (3)	2	P116-P124	P88-P124, T107-T169
	T153	P116	P3, P4, P19 (3)	2	P117-P124	P88-P124, T107-T169
	T154	P117	P3, P4, P19 (3)	2	P118-P124	P88-P124, T107-T169
	T155	P118	P22	1	P119-P124	P0-P198, T0-T275
	T156	P114	P5, P114	1	-	-
	T157	P115	P5, P115	1	-	-
	T158	P116	P5, P116	1	-	-
	T159	P117	P5, P117	1	-	-
	T160	P119	P3, P4, P19 (3)	2	P120-P124	P88-P124, T107-T169
	T161	P120	P3, P4, P19 (3)	2	P121-P124	P88-P124, T107-T169
	T162	P121	P3, P4, P19 (3)	2	P122-P124	P88-P124, T107-T169
	T163	P122	P3, P4, P19 (3)	2	P123, P124	P88-P124, T107-T169
	T164	P123	P22	1	P124	P0-P198, T0-T275
	T165	P124	P22	1	-	P0-P198, T0-T275
	T166	P119	P5, P119	1	-	-
	T167	P120	P5, P120	1	-	-
	T168	P121	P5, P121	1	-	-
	T169	P122	P5, P122	1	-	-
Degradation (Wheelburn on a casting)	T170	P125	P126	~E(0.00000828)	-	-
	T171	P1, P126	P1, P127	0	-	-
	T173	P127	P3, P19	91	-	P125-P127, T170-T173
Degradation (Lipping on casting)	T174	P128	P129	~E(0.000189)	-	-
	T175	P1, P129	P1, P130	0	-	-
	T176	P130	P2, P18	91	-	P128-P130, T174-T176
Degradation (Wear/damage in wheel transfer area)	T177	P131	P132	~E(0.0000423)	-	-
	T178	P1, P132	P1, P133	0	-	-
	T179	P133	P3, P19	182	-	P131-P133, T177-T179
Degradation (Transverse defect from RCF in a casting)	T180	P134	P135	~L(9.08, 2.41)	-	-
	T181	P1, P135	P1,P136	0	-	-
	T182	P136	P3, P19	91	-	P134-P136, T180-T182
Degradation (Transverse cracking of casting extending from running surface to a position)	T183	P137	P138	~E(0.00000166)	-	-
	T184	P128	P139	~E(1)	-	-
	T185	P139	P140	~E(1)	-	-
	T186	P140	P141	~E(1)	-	-
	T187	P141	P142	~E(1)	-	-
	T188	P1, P138	P1, P138, P143	0	P143	-
	T189	P1, P139	P1, P139, P144	0	P144	-
	T190	P1, P140	P1, P140, P145	0	P145	-
	T191	P1, P141	P1, P141, P146	0	P146	-
	T192	P1, P142	P1, P147	0	-	-
	T193	P143	P3, P19 (2)	28	P144-P147	P137-P147, T183-T197
	T194	P144	P3, P19 (3)	7	P145-P147	P137-P147, T183-T197
	T195	P145	P3, P4, P19 (3)	2	P146, P147	P137-P147, T183-T197
	T196	P146	P3, P4, P19 (3)	2	P147	P137-P147, T183-T197
	T197	P147	P3, P4, P19 (3)	2	-	P137-P147, T183-T197

	T198	P145	P5, P145	1	-	-
	T199	P146	P5, P146	1	-	-
	T200	P147	P5, P147	1	-	-
Degradation (Flaking or shelling of the wheel contact surface)	T201	P148	P149	~E(0.0000878)	-	-
	T202	P1, P149	P1, P150	0	-	-
	T203	P150	P3, P19	91	-	P148-P150, T201-T203
Degradation (Defective weld repair)	T204	P151	P152	~E(0.0000117)	-	-
	T205	P1, P152	P1, P153	0	-	-
	T206	P153	P154	91	-	-
	T207	P154	P2, P18	~U(0, 1)	-	P151-P154, T204-T206
	T208	P154	P3, P19	0.58	-	P151-P154, T204-T206
	T209	P19	P19, P151	0	P151-P153	-
Degradation (Ballast pitting)	T210	P155	P156	~E(0.00000670)	-	-
	T211	P1, P156	P1, P157	0	-	-
	T212	P157	P158	91	-	-
	T213	P158	P2, P18	~U(0, 1)	-	P155-P158, T210-T212
	T214	P158	P3, P19	0.58	-	P155-P158, T210-T212
Degradation (Rail end batter)	T215	P159	P160	~E(0.00000166)	-	-
	T216	P1, P160	P1, P161	0	-	-
	T217	P161	P3, P19	91	-	P159-P161, T215-T217
Degradation (Cracking in web of casting not covered elsewhere)	T218	P162	P163	~E(0.00000166)	-	-
	T219	P1, P163	P1, P164	0	-	-
	T220	P164	P165	91	-	-
	T221	P165	P2, P18	~U(0, 1)	-	P162-P165, T218-T220
	T222	P165	P3, P19	0.58	-	P162-P165, T218-T220
Degradation (Crack on foot extending across full width of casting and to a position on any one side)	T223	P166	P167	~E(0.00000166)	-	-
	T224	P167	P168	~E(1)	-	-
	T225	P168	P169	~E(1)	-	-
	T226	P1, P167	P1, P167, P170	0	P170	-
	T227	P1, P168	P1, P168, P171	0	P171	-
	T228	P1, P169	P1, P172	0	-	-
	T229	P170	P173	7	P171, P172	-
	T230	P171	P4, P173	7	P172	-
	T231	P172	P4, P173	2	-	-
	T234	P171	P5, P171	1	-	-
	T235	P172	P5, P172	1	-	-
	T236	P173	P2, P18 (3)	~U(0, 1)	-	P166-P173, T223-T235
	T237	P173	P3, P19 (3)	0.58	-	P166-P173, T223-T235
Degradation (Transverse crack in the foot extending to a position)	T238	P174	P175	~E(0.00000166)	-	-
	T239	P175	P176	~E(1)	-	-
	T240	P176	P177	~E(1)	-	-
	T241	P178	P179	~E(1)	-	-
	T242	P175	P180	~E(1)	-	-
	T243	P180	P181	~E(1)	-	-
	T244	P181	P182	~E(1)	-	-
	T245	P1, P175	P1, P175, P183	0	P183	-
	T246	P1, P176	P1, P176, P184	0	P184	-
	T247	P1, P177	P1, P177, P185	0	P185	-
	T248	P1, P178	P1, P186	0	-	-
	T249	P1, P180	P1, P180, P187	0	P187	-
	T250	P1, P181	P1, P181, P188	0	P188	-
	T251	P1, P182	P1, P189	0	-	-
	T252	P183	P190	28	P184-P189	-
	T253	P184	P191	7	P185-P189	-
	T254	P185	P4, P191	2	P186-P189	-
	T255	P186	P4, P191	2	P187-P189	-
	T256	P187	P191	7	P188, P189	-
	T257	P188	P4, P191	7	P189	-
	T258	P189	P4, P191	2	-	-
	T259	P185	P5, P185	1	-	-
	T260	P186	P5, P186	1	-	-
	T261	P188	P5, P188	1	-	-
	T262	P189	P5, P189	1	-	-

	T263	P190	P2, P18 (2)	~U(0, 1)	-	P174-P191, T238-T262
	T264	P190	P3, P19 (2)	0.58	-	P174-P191, T238-T262
	T265	P191	P2, P18 (3)	~U(0, 1)	-	P174-P191, T238-T262
	T266	P191	P3, P19 (3)	0.58	-	P174-P191, T238-T262
Degradation (Longitudinal cracking of the base of the flangeway)	T267	P192	P193	~E(0.00000166)	-	-
	T268	P1, P193	P1, P194	0	-	-
	T269	P194	P195	91	-	-
	T270	P195	P2, P18	~U(0, 1)	-	P192-P195, T267-T269
	T271	P195	P3, P19	0.58	-	P192-P195, T267-T269
Degradation (Depression of the running surface)	T272	P196	P197	~E(0.00000166)	-	-
	T273	P1, P197	P1, P198	0	-	-
	T274	P198	P3, P4, P19 (3)	7	-	P196-P198, T272-T275
	T275	P198	P5, P198	1	-	-

A.2.2 Initial marking

Initial marking value	**Place no.**
1	P0, P9, P23, P29, P34, P39, P50, P59, P66, P74, P81, P88, P125, P128, P131, P134, P137, P148, P151, P155, P159, P162, P166, P174, P192, P196
0	P1, P2, P3, P4, P5, P6, P7, P8, P10, P11, P12, P13, P14, P15, P16, P17, P18, P19, P20, P21, P22, P24, P25, P26, P27, P28, P30, P31, P32, P33, P35, P36, P37, P38, P40, P41, P42, P43, P44, P45, P46, P47, P48, P49, P51, P52, P53, P54, P55, P56, P57, P58, P60, P61, P62, P63, P64, P65, P67, P68, P69, P70, P71, P72, P73, P75, P76, P77, P78, P79, P80, P82, P83, P84, P85, P86, P87, P89, P90, P91, P92, P93, P94, P95, P96, P97, P98, P99, P100, P101, P102, P103, P104, P105, P106, P107, P108, P109, P110, P111, P112, P113, P114, P115, P116, P117, P118, P119, P120, P121, P122, P123, P124, P126, P127, P128, P129, P130, P132, P133, P135, P136, P138, P139, P140, P141, P142, P143, P144, P145, P146, P147, P149, P150, P152, P153, P154, P156, P157, P158, P160, P161, P163, P164, P165, P167, P168, P169, P170, P171, P172, P173, P175, P176, P177, P178, P179, P180, P181, P182, P183, P184, P185, P186, P187, P188, P189, P190, P191, P193, P194, P195, P197, P198

A.2.3 Place dependent transitions

Marked place	**Dependent transition**	**Adjusted firing time (days)**
P27, P47, P57, P65, P70, P87, P107, P108, P109, P110, P111, P113, P144	T0	1
P46, P56, P64, P86, P143	T0	4
P26, P32, P45, P55, P63, P85, P150	T0	7
P33, P38, P48, P49, P58, P72, P87, P112, P114, P115, P116, P117, P119, P120, P121, P122, P145, P146, P147, P172, P185, P186, P189	T21, T22, T29, T36, T49, T50, T51, T63, T64, T65, T74, T75, T76, T83, T96, T97, T104, T105, T143, T144, T145, T146, T147, T150, T173, T176, T179, T182, T193, T194, T203, T206, T212, T217, T220, T229, T230, T252, T253, T256, T257, T269, T274	2
P32, P37, P57, P64, P65, P71, P86, P107, P108, P109, P110, P111, P113, P144, P170, P171, P184, P187, P188, P198	T21, T22, T49, T50, T51, T63, T64, T74, T96, T97, T104, T173, T176, T179, T182, T193, T203, T206, T212, T217, T220, T252, T269	7
P47, P85	T21, T22, T49, T50, T63, T64, T74, T96, T97, T173, T176, T179, T182, T193, T203, T206, T212, T217, T220, T252, T269	14
P27, P46, P56, P63, P143, P183	T21, T49, T63, T96, T97, T173, T176, T179, T182, T203, T206, T212, T217, T220, T269	28
P26, P45, P55, P80, P127, P130, P136, P150, P153, P157, P161, P164, P194	T96, T179	91
P79, P133	P78	182

A.3 Plain line rail

A.3.1 Transitions

Sub-model type	Transition number	Input places	Output places	Firing time (days)	Inhibited by	Resets
Inspection	T0	P0	P1, P17 (20)	182	-	-
	T1	P1	P0	1	-	-
Maintenance	T2	P2	P10, P11, P14, P17 (130)	0	-	-
	T3	P3	P10, P11, P14, P17 (200)	0	-	-
	T4	P4	P17 (40)	0	-	P4
	T5	P5	P17 (3)	0	-	P5
	T6	P6	P7, P20 (40)	720	-	-
	T7	P7	P6	1	-	P7
	T8	P8	P8, P9	T21/140	-	-
	T9	P9	P22	0	-	-
Renewal	T10	P10 (2)	P25	0	P11	-
	T11	P12 (2), P15	P15, P25	0	P13	-
	T12	P14	P15	720	-	-
	T13	P15	P14	1	-	P12
	T14	P16 (2)	P16, P18	0	-	T15
	T15	P16	P17	913	-	P18
	T16	P18 (2)	P25	0	P19	-
	T17	P20 (2000)	P25	0	P21	-
	T18	P22 (80)	P25	0	P24	-
	T19	P23 (5)	P25	0	P24	-
	T20	P25	P26	14	-	P0-P224, T0-T313
Degradation (Vertical wear)	T21	P27	P28	~L(10.02, 0.54)	-	-
	T22	P1, P28	P1, P29	0	-	-
	T23	P29	P26	14	-	P0-P224, T0-T313
Degradation (Loss of rail foot width)	T24	P30	P31	~E(0.00000014)	-	-
	T25	P31	P32	~E(1)	-	-
	T26	P32	P33	~E(1)	-	-
	T27	P33	P34	~E(1)	-	-
	T28	P1, P31	P1, P31, P35	0	P35	-
	T29	P1, P32	P1, P32, P36	0	P36	-
	T30	P1, P33	P1, P33, P37	0	P37	-
	T31	P1, P34	P1, P38	0	-	-
	T32	P35	P26	91	P36-P38	P0-P224, T0-T313
	T33	P36	P26	28	P37, P38	P0-P224, T0-T313
	T34	P37	P26	7	P38	P0-P224, T0-T313
	T35	P38	P26	2	-	P0-P224, T0-T313
	T36	P36	P4, P36	1	-	-
	T37	P37	P4, P37	1	-	-
	T38	P38	P4, P38	1	-	-
Degradation (Sidewear)	T39	P39	P40	~E(0.00000014)	-	-
	T40	P1, P40	P1, P41	0	-	-
	T41	P41	P26	91	-	P0-P224, T0-T313
	T42	P41	P4, P41	1	-	-
Degradation (Cyclic sidewear)	T43	P42	P43	~E(0.00000014)	-	-
	T44	P1, P43	P1, P44	0	-	-
	T45	P44	P26	91	-	P0-P224, T0-T313
Degradation (Maximum permissible rail gall)	T46	P45	P46	~E(0.00000014)	-	-
	T47	P1, P46	P1, P47	0	-	-
	T48	P47	P26	91	-	P0-P224, T0-T313
Degradation (Rail foot arcing damage)	T49	P48	P49	~E(0.00000014)	-	-
	T50	P49	P50	~E(1)	-	-
	T51	P50	P51	~E(1)	-	-
	T52	P1, P49	P1, P49, P52	0	P52	-
	T53	P1, P50	P1, P50, P53	0	P53	-
	T54	P1, P51	P1, P54	0	-	-
	T55	P52	P5, P26	91	P53, P54	P0-P224, T0-T313
	T56	P53	P5, P26	2	P54	P0-P224, T0-T313
	T57	P54	P26	2	-	P0-P224, T0-T313
	T58	P53	P4, P53	1	-	-
	T59	P54	P4, P54	1	-	-
Degradation (Vertical longitudinal split in head or web)	T60	P55	P56	~E(0.00000014)	-	-
	T61	P56	P57	~E(1)	-	-
	T62	P57	P58	~E(1)	-	-
	T63	P58	P59	~E(1)	-	-
	T64	P1, P56	P1, P56, P60	0	P60	-
	T65	P1, P57	P1, P57, P61	0	P61	-

	T66	P1, P58	P1, P58, P62	0	P62	-
	T67	P1, P59	P1, P63	0	-	-
	T68	P60	P64	182	P61-P63	P55-P66, T60-T71
	T69	P61	P64	91	P62, P63	P55-P66, T60-T71
	T70	P62	P65	28	P63	P55-P66, T60-T71
	T71	P63	P66	7	-	P55-P66, T60-T71
	T72	P64	P2, P18 (10)	~U(0, 1)	-	-
	T73	P64	P3, P19	0.77	-	-
	T74	P65	P2, P18 (20)	~U(0, 1)	-	-
	T75	P65	P3, P19 (2)	0.77	-	-
	T76	P66	P2, P18 (30)	~U(0, 1)	-	-
	T77	P66	P3, P19 (3)	0.77	-	-
Degradation (Transverse defect from a tache ovale, squat, wheelburn, or weld repair)	T78	P67	P68	~E(0.0000016)	-	-
	T79	P68	P68, P69	~E(0.00000029)	P69	-
	T80	P69	P70	~E(0.00000028)	-	-
	T81	P70	P71	~E(1)	-	-
	T82	P1, P69	P1, P69, P72	0	P72	-
	T83	P1, P70	P1, P70, P73	0	P73	-
	T84	P1, P71	P1, P74	0	-	-
	T85	P72	P82	91	P73, P74, P78, P79, P81	P67-P81, T78-T102
	T86	P73	P83	7	P74, P78, P79	P67-P81, T78-T102
	T87	P74	P83	7	P78, P79	P67-P81, T78-T102
	T88	P68	P68, P75	~E(0.00000014)	P75	-
	T89	P75	P76	~E(1)	-	-
	T90	P76	P77	~E(1)	-	-
	T91	P1, P76	P1, P76, P78	0	P78	-
	T92	P1, P77	P1, P79	0	-	-
	T93	P78	P83	7	-	P67-P81, T78-T102
	T94	P79	P83	7	-	P67-P81, T78-T102
	T95	P1, P68	P1, P68, P80	0	P80	-
	T96	P7, P80	P7	0	P69-P79	P67-P81, T78-T102
	T97	P1, P75	P1, P75, P81	0	P81	-
	T98	P81	P82	28	-	P67-P81, T78-T102
	T99	P73	P4, P73	1	-	-
	T100	P74	P4, P74	1	-	-
	T101	P78	P4, P78	1	-	-
	T102	P79	P4, P79	1	-	-
	T103	P82	P2, P18 (10)	~U(0, 1)	-	-
	T104	P82	P3, P19	0.77	-	-
	T105	P83	P2, P18 (30)	~U(0, 1)	-	-
	T106	P83	P3, P19 (3)	0.77	-	-
Degradation (Horizontal and vertical cracking of rail head)	T107	P84	P85	~E(0.00000014)	-	-
	T108	P85	P85, P86	~E(1)	P86	-
	T109	P86	P86, P87	~E(1)	P87	-
	T110	P85	P85, P88	~E(1)	P88	-
	T111	P86	P86, P89	~E(1)	P89	-
	T112	P1, P85	P1, P85, P90	0	P90	-
	T113	P1, P86	P1, P86, P91	0	P91	-
	T114	P1, P87	P1, P92	0	-	-
	T115	P7, P90	P7	0	-	P84-P94, T107-T123
	T116	P91	P95	28	P92-P94	P84-P94, T107-T123
	T117	P92	P96	7	P93, P94	P84-P94, T107-T123
	T118	P1, P88	P1, P88, P93	0	P93	-
	T119	P1, P89	P1, P89, P94	0	P94	-
	T120	P93	P95	28	P94	P84-P94, T107-T123
	T121	P94	P96	7	-	P84-P94, T107-T123
	T122	P92	P4, P92	1	-	-
	T123	P94	P4, P94	1	-	-
	T124	P95	P2, P18 (20)	~U(0, 1)	-	-
	T125	P95	P3, P19 (2)	0.77	-	-
	T126	P96	P2, P18 (30)	~U(0, 1)	-	-
	T127	P96	P3, P19 (3)	0.77	-	-
Degradation (Horizontal defect in the head of the rail)	T128	P97	P98	~E(0.00000014)	-	
	T129	P98	P98, P99	~E(1)	P99	
	T130	P99	P99, P100	~E(1)	P100	
	T131	P99	P99, P101	~E(1)	P101	
	T132	P100	P102	~E(1)	-	
	T133	P1, P98	P1, P98, P103	0	P103	
	T134	P1, P99	P1, P99, P104	0	P104	

	T135	P1, P100	P1, P100, P105	0	P105	
	T136	P1, P101	P1, P106	0	-	
	T137	P1, P102	P1, P107	0	-	
	T138	P7, P103	P7	0	P104-P107	P97-P107, T128-T143
	T139	P104	P108	28	P105-P107	P97-P107, T128-T143
	T140	P105	P109	7	P106, P107	P97-P107, T128-T143
	T141	P106	P109	7	P107	P97-P107, T128-T143
	T142	P107	P109	7	-	P97-P107, T128-T143
	T143	P107	P4, P107	1	-	-
	T144	P108	P2, P18 (20)	~U(0, 1)	-	-
	T145	P108	P3, P19 (2)	0.77	-	-
	T146	P109	P2, P18 (30)	~U(0, 1)	-	-
	T147	P109	P3, P19 (3)	0.77	-	-
Degradation (Horizontal defect in the web of the rail)	T148	P110	P111	~E(0.00000014)	-	-
	T149	P1, P111	P1, P112	0	-	-
	T150	P112	P113	7	-	P110-P112, T148-T151
	T151	P112	P4, P112	1	-	-
	T152	P113	P2, P18 (30)	~U(0, 1)	-	-
	T153	P113	P3, P19 (3)	0.77	-	-
Degradation (Crushing)	T154	P114	P115	~E(0.0000012)	-	-
	T155	P1, P115	P1, P116	0	-	-
	T156	P116	P117	7	-	P114-P116, T154-T157
	T157	P116	P4, P116	1	-	-
	T158	P117	P2, P18 (30)	~U(0, 1)	-	-
	T159	P117	P3, P19 (3)	0.77	-	-
Degradation (Flaking)	T160	P118	P119	~E(0.00000027)	-	-
	T161	P1, P119	P1, P120	0	-	-
	T162	P120	P121	7	-	P118-P120, T160-T163
	T163	P120	P4, P120	1	-	-
	T164	P121	P2, P18 (30)	~U(0, 1)	-	-
	T165	P121	P3, P19 (3)	0.77	-	-
Degradation (Damage on the running surface)	T166	P122	P123	~E(0.0000015)	-	-
	T167	P123	P124	~E(0.00000015)	-	-
	T168	P124	P125	~E(1)	-	-
	T169	P1, P123	P1, P123, P126	0	P126	-
	T170	P1, P124	P1, P124, P127	0	P127	-
	T171	P1, P125	P1, P128	0	-	-
	T172	P7, P126	P7	0	P127, P128	P122-P128, T166-T174
	T173	P127	P5, P129	91	P128	P122-P128, T166-T174
	T174	P128	P5, P130	7	-	P122-P128, T166-T174
	T175	P129	P2, P18 (10)	~U(0, 1)	-	-
	T176	P129	P3, P19	0.77	-	-
	T177	P130	P2, P18 (30)	~U(0, 1)	-	-
	T178	P130	P3, P19 (3)	0.77	-	-
Degradation (Wheelburns)	T179	P131	P132	~E(0.0000048)	-	-
	T180	P1, P132	P1, P133	0	-	-
	T181	P133	P134	91	-	P131-P133, T179-T181
	T182	P134	P2, P18 (10)	~U(0, 1)	-	-
	T183	P134	P3, P19	0.77	-	-
Degradation (Transverse crack in rail foot**)**	T184	P135	P136	~E(0.00000028)	-	-
	T185	P1, P136	P1, P137	0	-	-
	T186	P137	P138	7	-	P135-P137, T184-T187
	T187	P137	P4, P137	1	-	-
	T189	P138	P2, P18 (30)	~U(0, 1)	-	-
	T190	P138	P3, P19 (3)	0.77	-	-
Degradation **(**Crack at hole)	T191	P139	P140	~E(0.00000014)	-	-
	T192	P1, P140	P1, P141	0	-	-
	T193	P141	P142	7	-	P135-P137, T191-T194
	T194	P141	P4, P141	1	-	-
	T195	P142	P2, P18 (30)	~U(0, 1)	-	-
	T196	P142	P3, P19 (3)	0.77	-	-
Degradation (Untestable rail)	T197	P143	P144	~L(19.37, 4.42)	-	-
	T198	P1, P144	P1, P145	0	-	-

	T199	P145	P146	91	-	P131-P133, T197-T199
	T200	P146	P2, P18 (10)	~U(0, 1)	-	-
	T201	P146	P3, P19	0.77	-	-
Degradation (Visible crack not covered by other tables)	T202	P147	P148	~L(24.59, 6.01)	-	-
	T203	P1, P148	P1, P149	0	-	-
	T204	P149	P150	7	-	P147-P149, T202-T205
	T205	P149	P4, P149	1	-	-
	T206	P150	P2, P18 (30)	~U(0, 1)	-	-
	T207	P150	P3, P19 (3)	0.77	-	-
Degradation (Hydrogen shatter cracking)	T208	P151	P152	~E(0.00000014)	-	
	T209	P152	P153	~E(1)	-	
	T210	P1, P152	P1, P152, P154	0	P154	
	T211	P1, P153	P1, P155	0	-	
	T212	P7, P154	P7	0	P155	P151-P155, T208-T213
	T213	P155	P156	91	-	P151-P155, T208-T213
	T214	P156	P2, P18 (10)	~U(0, 1)	-	-
	T215	P156	P3, P19	0.77	-	-
Degradation (Field side rolling contact fatigue)	T216	P157	P158	~E(0.00000014)	-	-
	T217	P158	P158, P159	~E(1)	P159	-
	T218	P158	P158, P160	~E(1)	P160	-
	T219	P160	P161	~E(1)	-	-
	T220	P161	P162	~E(1)	-	-
	T221	P159	P163	~E(1)	-	-
	T222	P1, P158	P1, P158, P164	0	P164	-
	T223	P1, P160	P1, P160, P165	0	P165	-
	T224	P1, P161	P1, P161, P166	0	P166	-
	T225	P1, P162	P1, P167	0	-	-
	T226	P1, P159	P1, P159, P168	0	P168	-
	T227	P1, P63	P1, P169	0	-	-
	T228	P7, P164	P7	0	P165-P169	P157-P169, T216-T233
	T229	P165	P170	91	P166-P169	P157-P169, T216-T233
	T230	P166	P171	28	P167-P169	P157-P169, T216-T233
	T231	P167	P172	7	P168, P169	P157-P169, T216-T233
	T232	P168	P171	28	P169	P157-P169, T216-T233
	T233	P169	P172	7	-	P157-P169, T216-T233
	T234	P170	P2, P18 (10)	~U(0, 1)	-	-
	T235	P170	P3, P19	0.77	-	-
	T236	P171	P2, P18 (20)	~U(0, 1)	-	-
	T237	P171	P3, P19 (2)	0.77	-	-
	T238	P172	P2, P18 (30)	~U(0, 1)	-	-
	T239	P172	P3, P19 (3)	0.77	-	-
Degradation (RCF in full section rail)	T240	P173	P174	~E(0.00000014)	-	-
	T241	P174	P175	~E(1)	-	-
	T242	P175	P176	~E(1)	-	-
	T243	P176	P177	~E(1)	-	-
	T244	P1, P174	P1, P174, P178	0	P178	-
	T245	P1, P175	P1, P175, P179	0	P179	-
	T246	P1, P176	P1, P176, P180	0	P180	-
	T247	P1, P177	P1, P181	0	-	-
	T248	P7, P178	P7	0	P179-P181	P173-P181, T240-T252
	T249	P179	P182	91	P180, P181	P173-P181, T240-T252
	T250	P180	P183	7	P181	P173-P181, T240-T252
	T251	P181	P183	7	-	P173-P181, T240-T252
	T252	P181	P4, P181	1	-	-
	T253	P182	P2, P18 (10)	~U(0, 1)	-	-
	T254	P182	P3, P19	0.77	-	-
	T255	P183	P2, P18 (30)	~U(0, 1)	-	-
	T256	P183	P3, P19 (3)	0.77	-	-
Degradation (RCF in reduced section rail)	T257	P184	P185	~E(0.00000014)	-	-
	T258	P1, P185	P1, P186	0	-	-
	T259	P7, P186	P7	0	-	P184-P186, T257-T259
	T260	P187	P188	~E(0.00000509)	-	-

Degradation (Transverse cracking from a defective arc weld repair)	T261	P1, P188	P1, P189	0	-	-
	T262	P189	P190	91	-	P187-P189, T260-T262
	T263	P190	P2, P18 (10)	~U(0, 1)	-	-
	T264	P190	P3, P19	0.77	-	-
	T265	P19	P19, P187	0	P187-P190	-
Degradation (Arc weld repair with an ultrasonically detected defect)	T266	P191	P192	~E(0.00006615)	-	-
	T267	P192	P193	~E(0.00000553)	-	-
	T268	P1, P192	P1, P192, P194	0	P194	-
	T269	P1, P193	P1, P195	0	-	-
	T270	P194	P196	91	P195	P191-P195, T266-T271
	T271	P195	P196	28	-	P191-P195, T266-T271
	T272	P196	P2, P18 (10)	~U(0, 1)	-	-
	T273	P196	P3, P19	0.77	-	-
	T274	P19	P19, P191	0	P191-P196	-
Degradation (Defective aluminothermic weld with MPT indication)	T275	P197	P198	~E(0.000012)	-	-
	T276	P1, P198	P1, P199	0	-	-
	T277	P199	P200	91	-	P197-P199, T275-T277
	T278	P200	P2, P18 (10)	~U(0, 1)	-	-
	T279	P200	P3, P19	0.77	-	-
Degradation (Defective aluminothermic weld with a horizontal/ transverse UT defect)	T280	P201	P202	~E(0.00000014)	-	-
	T281	P202	P203	~E(1)	-	-
	T282	P203	P204	~E(1)	-	-
	T283	P204	P205	~E(1)	-	-
	T284	P1, P202	P1, P202, P206	0	P206	-
	T285	P1, P203	P1, P203, P207	0	P207	-
	T286	P1, P204	P1, P204, P208	0	P208	-
	T287	P1, P205	P1, P209	0	-	-
	T288	P7, P206	P7	0	P207-P209	P201-P209, T280-T291
	T289	P207	P210	91	P208, P209	P201-P209, T280-T291
	T290	P208	P211	7	P209	P201-P209, T280-T291
	T291	P209	P211	7	-	P201-P209, T280-T291
	T292	P210	P2, P18 (10)	~U(0, 1)	-	-
	T293	P210	P3, P19	0.77	-	-
	T294	P211	P2, P18 (30)	~U(0, 1)	-	-
	T295	P211	P3, P19 (3)	0.77	-	-
Degradation (Crack in weld)	T296	P212	P213	~E(0.00000014)	-	-
	T297	P213	P214	~E(1)	-	-
	T298	P214	P215	~E(1)	-	-
	T299	P1, P213	P1, P213, P216	0	P216	-
	T300	P1, P214	P1, P214, P217	0	P217	-
	T301	P1, P215	P1, P218	0	-	-
	T302	P7, P216	P7	0	P219, P220	P212-P219, T296-T304
	T303	P217	P219	91	P220	P212-P219, T296-T304
	T304	P218	P220	7	-	P212-P219, T296-T304
	T305	P219	P2, P18 (10)	~U(0, 1)	-	-
	T306	P219	P3, P19	0.77	-	-
	T307	P220	P2, P18 (30)	~U(0, 1)	-	-
	T308	P220	P3, P19 (3)	0.77	-	-
Degradation (Transversal cracking in welds in rail foot)	T309	P221	P222	~E(0.00000014)	-	-
	T310	P1, P222	P1, P223	0	-	-
	T311	P223	P224	7	-	P221-P224, T309-T312
	T312	P223	P4, P223	1	-	-
	T313	P224	P2, P18 (30)	~U(0, 1)	-	-
	T314	P224	P3, P19 (3)	0.77	-	-

A.3.2 Initial marking

Initial marking value	Place no.
1	P0, P6, P8, P14, P27, P30, P39, P42, P45, P48, P55, P67, P84, P97, P110, P114, P118, P122, P131, P135, P139, P143, P147, P151, P157, P173, P184, P197, P201, P212, P221
0	P1, P2, P3, P4, P5, P7, P9, P10, P11, P12, P13, P15, P16, P17, P18, P19, P20, P21, P22, P23, P24, P25, P26, P28, P29, P31, P32, P33, P34, P35, P36, P37, P38, P40, P41, P43, P44, P46, P47, P49, P50, P51, P52, P53, P54,P56, P57, P58, P59, P60, P61, P62, P63, P64, P65, P66, P68, P69, P70, P71, P72, P73, P74, P75, P76, P77, P78, P79, P80, P81, P82, P83, P85, P86, P87, P88, P89, P90, P91, P92, P93, P94, P95, P96, P98, P99, P100, P101, P102, P103, P104, P105, P106, P107, P108, P109, P111, P112, P113, P115, P116, P117, P119, P120, P121, P123, P124, P125, P126, P127, P128, P129, P130, P132, P133, P134, P136, P137, P138, P140, P141, P142, P144, P145, P146, P148, P149, P150, P152, P153, P154, P155, P156, P158, P159, P160, P161, P162, P163, P164, P165, P166, P167, P168, P169, P170, P171, P172, P174, P175, P176, P177, P178, P179, P180, P181, P182, P183, P185, P186, P187, P188, P189, P190, P191, P192, P193, P194, P195, P196, P198, P199, P200, P202, P203, P204, P205, P206, P207, P208, P209, P210, P211, P213, P214, P215, P216, P217, P218, P219, P220, P222, P223, P224

A.3.3 Place dependent transitions

Marked place	Dependent transition	Adjusted firing time (days)
P80, P90, P103, P126, P154, P164, P178, P186, P206, P216	T0	28
P38, P53, P54	T23, T32, T33, T34, T41, T45, T48, T55, T68, T69, T70, T71, T85, T86, T87, T93, T94, T98, T116, T117, T120, T121, T139, T140, T141, T142, T150, T156, T162, T173, T174, T181, T186, T193, T199, T204, T213, T229, T230, T231, T232, T233, T249, T250, T251, T262, T269, T270, T275, T288, T289, T290, T302, T303, T310	2
P37, P63, P73, P74, P78, P79, P92, P94, P105, P106, P107, P112, P116, P120, P128, P137, P141, P149, P167, P169, P180, P181, P208, P209, P218, P223	T23, T32, T33, T41, T45, T48, T55, T68, T69, T70, T85, T98, T116, T120, T139, T173, T181, T199, T213, T229, T230, T232, T249, T262, T269, T270, T275, T288, T302	7
P29	T32, T33, T41, T45, T48, T55, T68, T69, T70, T85, T98, T116, T120, T139, T173, T181, T199, T213, T229, T230, T232, T249, T262, T269, T270, T275, T288, T302	14
P36, P62, P81, P91, P93, P104, P166, P168, P195	T32, T41, T45, T48, T55, T68, T69, T85, T173, T181, T199, T213, T229, T249, T262, T269, T275, T288, T302	28
P35, P41, P44, P47, P52, P61, P72, P127, P133, P145, P155, P165, P179, P189, P194, P199, P207, P217	T68	91

A.4 Ballast

Transitions

Sub-model type	Transition number	Input places (multiplicity)	Output places (multiplicity)	Firing time in days	Inhibited by	Resets
Inspection	T0	P0	P1, P35 (3)	40	-	-
	T1	P1	P0	1	-	-
Degradation (Isolated vertical alignment)	T2	P2	P3	~W(0.42, 26597)	-	-
	T3	P3	P4	~W(0.65, 7208020)	-	-
	T4	P4	P5	~W(1, 1581720)	-	-
	T5	P5	P6	~W(1, 1)	-	-
	T6	P6	P7	~W(1, 1)	-	-
	T7	P7	P8	~W(1, 1)	-	-
	T8	P1, P3	P1, P3, P9	0	P9, P10, P11	-
	T9	P1, P4	P1, P4, P12	0	P12	-
	T10	P1, P5	P1, P5, P13	0	P13	-
	T11	P1, P6	P1, P6, P14	0	P14	-
	T12	P1, P7	P1, P7, P15	0	P15	-
	T13	P1, P8	P1, P16	0	-	-
	T14	P9	P11	~U(0,1)	-	-
	T15	P9	P10	0.56	-	-
	T16	P11	P25	60	P12-P16	-
	T17	P12	P25	30	P13-P16	-
	T18	P13	P25	20	P14-P16	-
	T19	P13	P13, P26	1	P14-P16	-
	T20	P14	P25	10	P15, P16	-
	T21	P14	P14, P26	1	P15, P16	-
	T22	P15	P25	2	P16	-
	T23	P15	P15, P26	1	P16	-
	T24	P16	P27	2	-	-
Degradation (SD vertical alignment)	T25	P17	P18	~W(1.089, 4759)	-	-
	T26	P18	P19	~W(2.98, 4295)	-	-
	T27	P1, P18	P1, P18, P20	0	P20, P21, P22	-
	T28	P1, P19	P1, P23	0	-	-
	T29	P20	P22	~U(0,1)	-	-
	T30	P20	P21	0.82	-	-
	T31	P21	P22	100000	-	-
	T32	P22	P24	90	P23	-
	T33	P23	P24	30	-	-
Maintenance	T34	P24	P28, P29, P32, P35 (20), P36 (2), P44	0	-	P2-P24, T2-T33
	T35	P25	P28, P29, P32, P35 (10), P36 (1),	0	-	P2-P16, T2-T24, P25
	T36	P26	P35	0	-	P26
	T37	P27	P35	0	-	P27
Renewal	T38	P28 (10)	P42	0	P37	-
	T39	P29 (3), P31	P31, P42	0	P38	-
	T40	P30	P31	2190	-	-
	T41	P31	P30	1	-	P29
	T42	P32 (2)	P32, P33	0	-	T43
	T43	P32	P34	1095	-	P33
	T44	P33 (2)	P42	0	P39	-
	T45	P35 (10)	P42	0	P40	-
	T46	P36 (10)	P42	0	P41	-
	T47	P42	P43	14	-	P0-P42, T0-T47

A.4.1 Initial marking

Initial marking value	Place no.
1	P0, P2, P17, P30
0	P1, P3, P4, P5, P6, P7, P8, P9, P10, P11, P12, P13,P14, P15, P16, P18, P19, P20, P21, P22, P23, P24, P25, P26, P27, P28, P29, P31, P32, P33, P34, P35, P36, P37, P38, P39, P40, P41, P42, P43, P44, P45, P46, P47

A.4.2 Place dependent transitions

Marked place (multiplicity)	Dependent transition	Adjusted firing time (days)
P44	T25	~W(0.87, 3833)
P44	T26	~W(2.56, 4924)
P44 (2)	T25	~W(0.82, 3950)
P44 (2)	T26	~W(2.75, 5352)
P44 (3)	T25	~W(0.77, 3284)
P44 (3)	T26	~W(3.45, 5293)
P44 (4)	T25	~W(0.85, 2283)
P44 (4)	T26	~W(3.24, 3389)
P44 (5)	T25	~W(0.8, 2293)
P44 (5)	T26	~W(3.24, 3389)
P44 (6)	T25	~W(0.92, 2759)
P44 (6)	T26	~W(3.30, 3597)
P44 (7)	T25	~W(0.76, 1589)
P44 (7)	T26	~W(3.02, 2695)
P21	T14	1
P10	T31	0