Smart Asset Management: Data-driven Approaches for Degradation Prediction and Maintenance Policy

Guo Shi

Department of Management Science

University of Strathclyde, Glasgow

May 21, 2025

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

Abstract

With the use of advanced technologies and digital tools, asset management is advancing to a smarter stage with increasing availability of operational data. However, this also presents challenges in effectively utilising data to drive informed decisions. This thesis aims to address two critical tasks within the context of smart asset management through data-driven approaches: predicting degradation and developing maintenance policies. Three parts under different scenarios are involved in this thesis.

To address the challenge of degradation prediction with limited observations, the first part of this thesis introduces a method that generates, selects, and reweights synthetic data to enhance prediction performance. Unlike existing methods that mix synthetic and real data without considering sample selection or weighting, this approach uses multiple data augmentation methods to generate time-series data, then applies an influence function to select effective synthetic samples, followed by reweighting via gradient descent. To further improve the performance of the deep learning algorithm, transfer learning is applied by pre-training the deep learning model and then fine-tuning it with real data. Numerical experiments demonstrate the framework's effectiveness, especially for highly stochastic degradation data.

The second part of the thesis explores a data-driven preventive maintenance problem where the true time-to-failure model is unknown, but past time-to-failure data and working conditions are observable. Traditional estimate-then-optimise methods separate estimation from optimisation, potentially propagating errors into the decision-making process. To overcome this, an end-to-end framework has been proposed to directly determine the optimal preventive replacement time under specific working conditions without assuming a time-to-failure model. The end-to-end approach treats

historical working conditions as features, mapping them to optimal maintenance decisions by minimising the maintenance cost rate. Supervised learning algorithms then train these features against the optimal decisions. The findings suggest that end-to-end learning can reduce error propagation and that a linear model, when aligned with the learning objective, may outperform more complex alternatives.

Lastly, the third part of the thesis presents a condition-based maintenance policy considering component heterogeneity and dynamic working conditions. Bayesian Poisson and linear regressions are applied to analyse the shock occurrence and magnitude, updating parameters with new observations during online monitoring. The maintenance planning problem is framed as a Markov decision process. This approach establishes a tractable degradation model that accounts for heterogeneity and dynamic conditions and explores the structural properties of optimal maintenance policy. A heuristic algorithm based on the most likely distribution has been introduced to reduce computational complexity. Results reveal that maintenance thresholds, which serve as control limits, fluctuate with covariates and demonstrate the advantages of the proposed policy in handling varying working conditions and component heterogeneity.

Collectively, this thesis tackles the challenges of degradation prediction and maintenance optimisation in smart asset management through data-driven approaches. The proposed methods offer improved solutions for degradation prediction with limited data, and maintenance policies with dynamic working conditions and heterogeneity, contributing to more effective and informed asset management strategies.

Contents

List of Figures

List of Tables

List of Abbreviations

AI	Artificial Intelligence
AFT	Accelerated Failure Time
ANN	Artificial Neural Network
ADRT	Enhancement of deep learning algorithm by integrating data Augmentation, Dropout, Reweight and Transfer learning
ARIMA	Auto Regressive Integrated Moving Average Model
CBM	Condition-Based Maintenance
CDF	Cumulative Distribution Function
CM	Corrective Maintenance
CP	Compound Poisson
DA	Data Augmentation
DL	Deep Learning
E2E	End-to-End
ETO	Esitimate-Then-Optimise
LASSO	Least Absolute Shrinkage and Selection Operator
LGBM	Light Gradient Boosting Machine
LSTM	Long Short-Term Memory model
LSTM_ADRT	Long-Short Term Memory model with pre-processed Augmented samples after Dropout, Reweight and Transfer learning
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MDP	Markov Decision Process
ML	Machine Learning

MLE	Maximum Likelihood Estimation
GB	Gradient Boosting
SVI	Sludge Volume Index
TL	Transfer Learning
TBM	Time-Based Maintenance
PH	Proportional Hazard
PI	Proportional Intensity
PM	Preventive Maintenance
PR	Preventive Replacement
PDF	Probability Density Function
POMDP	Partially Observable MDP
RFR	Random Forest Regressor
RMSE	Root Mean Squared Error
XGB	eXtreme Gradient Boosting

Acknowledgements

I would like to express my deepest gratitude to my supervisors, Dr. Bin Liu and Professor Lesley Walls. Their guidance and support have been instrumental in my PhD journey. Dr. Bin Liu has consistently encouraged me, helping to illuminate the path toward learning and researching cutting-edge topics. Professor Lesley Walls has always inspired me to think critically and approach problems from new perspectives. Without their mentorship and support, I would not have been able to complete my PhD. Additionally, I would like to thank Professor John Quigley, who has been my internal reviewer for three years and provided reflective suggestions after each annual review.

I would also like to extend my heartfelt thanks to my boyfriend, Yu Chen, who was the only person aside from my supervisors to discuss my research with me during the COVID-19 pandemic. His unwavering support, both in my academic work and in life, has been invaluable.

Special thanks go to my family—my mom, dad, and sister—for always supporting my decisions and providing mental and emotional encouragement throughout this journey.

Lastly, I want to thank my friends, both my old friends in China and the new ones I've met in Glasgow. A strong social life has been essential to maintaining mental well-being during this challenging PhD journey, and I am deeply grateful for their companionship.

Chapter 1

Introduction

1.1 Background

Unplanned failure of a critical asset poses a considerable risk, potentially leading to severe societal and financial consequences. For instance, China Northwest Airlines Flight 2303, a domestic flight from Xi'an to Guangzhou, experienced a catastrophic failure on June 6, 1994. The aircraft, a Tupolev Tu-154M, broke up mid-flight and crashed due to an autopilot malfunction, which caused violent shaking and overstressed the airframe. As of 2024, it remains the deadliest airplane crash in mainland China (Flight Safety Foundation, 2024). Motivated by the consequences of unplanned breakdowns, asset management is proposed as a strategy to prevent failures. Assets can be classified into two categories: tangible assets (physical objects such as infrastructure, manufacturing plants, buildings, etc.) and intangible assets (such as financial assets) (Van der Lei et al., 2012). This thesis focuses on tangible assets, particularly in engineering systems.

The International Organisation for Standardisation (ISO) has defined Asset management as "an integrated approach to deriving value from asset systems" in ISO 55000 (ISO, 2014). Asset management plays a critical role in ensuring the long-term performance and value of physical assets across various industries. It encompasses a variety of activities, from acquisition and operation to maintenance and eventual disposal, all with the goal of optimising the lifecycle and cost-efficiency of assets. Among these activities, maintenance stands out as a pivotal component that directly influences the

reliability, availability, and performance of assets. Effective maintenance strategies not only prolong the operational life of assets but also reduce downtime and operational costs. In this context, maintenance can be viewed as the backbone of asset management, providing the essential processes needed to sustain asset functionality.

Today, advanced technologies like the Internet of Things (IoT), artificial intelligence, and big data analytics are drawing researchers' attention to smart asset management. These technologies can be applied to optimise the use, performance, and maintenance of physical assets (Teoh et al., 2021). By integrating real-time data collection, predictive analytics, and automation, smart asset management systems enable organisations to make informed decisions, reduce downtime, and extend the lifespan of their assets. These approaches not only enhance operational efficiency but also improve cost-effectiveness and sustainability (Liu et al., 2023). This thesis delves into the maintenance aspect of asset management, with a particular focus on predictive maintenance techniques and their role in enhancing the overall management of assets in complex, data-driven environments.

Four key tasks are included in smart asset management:

1. **Identify a set of health indicators.** Health indicators are constructed to evaluate the health state of the asset and its critical components. These indicators can be identified directly or indirectly. For example, in a bearing system, health indicators include kurtosis and the root mean square of vibration signals (Wang et al., 2017). In a wastewater system, the Sludge Volume Index (SVI) can reflect the tendency of activated sludge solids (Liu et al., 2020).

2. **Establish a monitor system to capture the degradation.** With diverse sensors and IoT systems, real-time monitoring of health indicators can be achieved. For instance, a piezoelectric sensor placed under the insert of a cutting tool can measure the force applied. As the tool wears, the measured force decreases, allowing the wear rate to be determined. The monitored data can be sent to the cloud, which has the computational power for prediction and maintenance scheduling (Mekid, 2021).

3. **Predict future levels of degradation.** With real-time data indicating the health state, asset managers can apply statistical, physical, or machine learning models to predict future degradation levels (Tian and Zuo, 2010). These predictions help identify potential future failures and estimate the remaining useful life of assets (Lei et al., 2018).

4. **Design a maintenance policy to improve asset availability or reduce the maintenance cost.** Once degradation patterns are analysed through historical health indicators, asset managers can adopt either time-based or condition-based maintenance policies, depending on the type of data and the chosen strategy (Ahmad and Kamaruddin, 2012).

This thesis focuses on prediction model development and maintenance policy design.

1.2 Research gaps and objectives

As stated, this thesis mainly focuses on developing algorithms to predict the future levels of health indicators and design maintenance policies. Although the two tasks are broadly discussed, several research gaps in the existing literature are summarised as follows.

Accurate degradation prediction is fundamental to predictive maintenance, as it enables decision-makers to intervene before failures occur and avoid costly unplanned downtime. Recent advances in Machine Learning (ML) and Deep Learning (DL) have introduced powerful tools for modelling complex degradation patterns. However, these models typically require large and diverse datasets to generalise well. In some industrial contexts, such data is scarce due to the high cost of sensor deployment, long asset lifecycles, or the infrequency of failure events in high-reliability systems. As a result, DL-based methods may suffer from overfitting or fail to extract meaningful features, leading to unreliable forecasts. Therefore, developing DL prediction methods that can operate effectively under limited-data regimes is essential.

While accurate prediction is crucial, the ability to translate predictions into optimal maintenance decisions is equally important. Maintenance decision-making tradition-

ally follows a two-stage process: model parameters—such as time-to-failure distributions—are first estimated from historical data and then used as inputs for a separate optimisation procedure that aims to minimise maintenance costs or maximise system availability. Although this separation is analytically convenient, it creates a disconnect between predictive accuracy and decision quality. The objective functions used for model estimation (e.g., minimising mean squared error) are often misaligned with downstream goals like cost efficiency or operational risk. Consequently, prediction errors or distributional biases may propagate into the decision layer, leading to suboptimal or even misleading maintenance policies. To address this challenge, recent studies advocate for end-to-end, decision-aware frameworks that couple predictive modelling with task-specific objectives (Qi et al., 2023; Donti et al., 2017). This thesis contributes to this paradigm by exploring data-driven end-to-end approaches to unify prediction and optimisation.

A further complication arises from the observation that asset degradation does not occur in isolation from environmental and operational context. Assets operate under dynamic conditions, such as fluctuating temperature, humidity, production rate, and operating speed (Kong et al., 2021). These factors can affect the degradation trajectories of components, and failure behaviours observed under one set of conditions may not generalise to others. Nevertheless, many existing models assume stationary or homogeneous environments, thereby neglecting the temporal and spatial variability inherent in practical applications. To overcome these limitations, this thesis proposes maintenance optimisation frameworks that are explicitly conditioned on environmental and operational variables, allowing degradation models to adapt dynamically to varying contexts.

Finally, even when environmental factors are controlled for, latent variability among components remains a modelling challenge. Components produced from the same design and operated under similar conditions may still degrade in different ways due to subtle differences in raw materials, manufacturing processes, or historical usage. This latent heterogeneity undermines the assumption that all components of the same type can be treated as statistically identical. Learning a single deterministic model

across all units risks obscuring individual degradation paths, potentially leading to inaccurate forecasts and misdirected maintenance actions. Moreover, applying uniform maintenance thresholds may result in over-servicing some components while leaving others at higher risk of failure. Addressing this issue calls for modelling strategies that can capture unit-level variability while still leveraging shared structures in the data (Yang et al., 2019; Sun et al., 2021). In response, this thesis explores hierarchical and latent-variable-based models to design maintenance strategies that are robust to hidden heterogeneity and capable of delivering maintenance policies.

Together, these challenges highlight gaps in the current predictive maintenance literature. They motivate the need for methods that can operate under limited data, align predictions with decision objectives, adapt to dynamic environments, and account for hidden heterogeneity across components. This thesis addresses these gaps by developing data-driven frameworks for degradation modelling and maintenance policy optimisation.

This thesis aims to address the aforementioned research gaps, with the objectives outlined as follows:

- Improve the performance of degradation prediction when applying DL with limited data size.
- Develop data-driven maintenance policies that can integrate the data within decision-making processes.
- Formulate effective maintenance policies that can account for varied working conditions.
- Design maintenance policies that consider the latent heterogeneity among different components.

1.3 Thesis outline

The rest of this thesis is organised as follows:

- Chapter 2 presents a literature review on smart asset management, focusing on methods for degradation prediction and categories of maintenance policies, including time-based and condition-based maintenance strategies.

- To address the challenges of degradation prediction with limited data, Chapter 3 extends the applicability of DL algorithms for health indicator prediction with limited observations by incorporating data augmentation, dropout, and transfer learning. Numerical experiments and a case study are provided to validate the proposed method.

- To consider the dynamic working conditions and derive an end-to-end maintenance policy, Chapter 4 proposes time-based preventive maintenance that directly determines the optimal maintenance decision using data, bypassing the parameter estimation stage and thereby reducing error propagation from estimation to decision-making.

- To incorporate both dynamic working conditions and heterogeneity, Chapter 5 proposes a Condition-Based Maintenance (CBM) policy that utilises real-time data to update the parameters of the degradation model, ensuring an accurate representation of degradation behaviour and enabling informed maintenance decisions. This approach provides an analytically tractable model, offering a framework for understanding the dynamics of CBM policies. Additionally, a heuristic solution algorithm is proposed to improve the model's computational efficiency.

- Chapter 6 provides the conclusion, which summarises the results by revisiting the thesis and discusses potential future directions in smart asset management.

Chapter 2

Literature Review

Asset management is proposed as a strategic and comprehensive approach for systematic planning, operation, and maintenance to maximise lifetime performance. As the industry integrates more technology and equipment to meet consumer demands, asset management has garnered notable interest from researchers across various disciplines. Industrial assets need to ensure their reliability for both safety and economic purposes. This necessitates a thorough understanding of their failure dynamics to accurately forecast future potential failure, and meet the maintenance needs. Consequently, developing quantitative models to predict the level of degradation and schedule maintenance plans is essential.

This Chapter will examine the existing research on two crucial tasks in industrial asset management: the prediction of degradation and the optimisation of maintenance schedules. We will explore the relevant methods for degradation prediction, including physics-based and data-driven methods. Regarding maintenance scheduling, we will discuss time-based and condition-based maintenance strategies, considering working conditions and the heterogeneity of components. In addition, this thesis will be positioned within the context of the existing literature. A summary will be provided to conclude the literature review, highlighting key findings and identifying gaps that this thesis aims to address.

2.1 Degradation prediction

Health indicators are designed to evaluate the degradation of engineering systems using time series data (Zhou et al., 2016; Ghofrani et al., 2022). Prediction of the health indicators plays an important role in asset monitoring systems by providing early warnings of potential failures and informing maintenance actions (Zakikhani et al., 2020). In general, the mainstream of degradation prediction approaches can be divided into two categories: physics-based and data-driven methods (Ren et al., 2022).

2.1.1 Physics-based methods

Physics-based methods utilise the underlying degradation mechanisms of components to model their deterioration over time by explicitly capturing the relationship between condition variables and system lifetime (Shahraki et al., 2017). These approaches quantitatively characterise failure behaviour using physical laws under specific loading conditions, such as thermal, mechanical, chemical, or electrical stress (Escobar and Meeker, 2006; Fan et al., 2011).

A variety of physics-based methods have been applied to describe electronic device degradation processes, including the Arrhenius model, Peck model, and Coffin–Manson model (Escobar and Meeker, 2006; Cui, 2005; Liu et al., 2019b). Among these, the Arrhenius model is used when temperature is the dominant acceleration factor. For instance, the degradation of EVA (ethyl vinyl acetate) and Tedlar layers in monocrystalline silicon photovoltaic modules has been investigated using Arrhenius-based approaches to simulate indoor and outdoor ageing under composite climate conditions (Rajput et al., 2017). The Peck model and the Coffin–Manson model have also been adopted to characterise degradation under environmental and mechanical stresses in electronic devices (Deng et al., 2023).

In addition, crack growth models have been applied to structural elements to establish the relationship between crack growth rate and stress. For example, finite element analysis is used to evaluate stress and strain, and is then integrated with Paris' law to estimate the crack growth rate based on a set of bearing run-to-failure data (Liao,

2013). A series of bending fatigue tests has also been conducted to analyse fatigue crack propagation behaviour by measuring the residual stress and identifying critical tooth regions (Yan et al., 2022). To account for the diverse nature of degradation, four physics-based models have been incorporated for fatigue crack growth prediction: the Paris-Erdogan model, polynomial model, global function-based model, and curve-fitting model (Nguyen et al., 2019). Paris' law has been combined with the extended finite element method and an active learning Kriging surrogate model to predict fatigue life and assess the reliability of cracked CFRP laminates under material and load uncertainties (Liu et al., 2025).

Physics-based models play an important role in modelling degradation processes, as they provide physically interpretable insights across a wide range of operating conditions. However, their applicability is constrained by the increasing complexity of modern systems and the lack of precise knowledge about failure mechanisms. In contrast, the growing availability of industrial data has made data-driven approaches an attractive alternative, supporting the development of flexible and adaptive degradation models.

2.1.2 Data-driven methods

Statistical methods

Statistical models for degradation modelling are typically classified into three subcategories: regression-based models, general path models, and stochastic process models (Escobar and Meeker, 2006). For example, logistic regression has been employed to estimate the degradation-related failure probabilities of bearings using run-to-failure datasets. These probabilities are then used as target vectors, which are subsequently combined with AutoRegressive-Moving Average (ARMA) model to predict future degradation trends (Caesarendra et al., 2010). In addition, filtering-based approaches have also been utilised. For instance, the drift rate—a key health indicator for gyroscopes that reflects long-term stability—has been predicted using random filtering theory to estimate future system drift (Wang et al., 2014b). Kalman filter has similarly been applied to infer the degradation states of systems affected by sensor deterioration (Liu

et al., 2019a).

Many failure mechanisms can be traced to underlying degradation processes (Gorjian et al., 2010). General path models and stochastic process models are used to represent the underlying degradation process and make inferences about future degradation. General path models describe the continuous degradation process as a linear/nonlinear function of time. A general non-linear regression model has been presented to describe the degradation paths of a population of units by introducing fixed-effect parameters for the common characteristics of the population and a random error effect for a specific unit (Lu et al., 2021). A linear degradation path model has been applied to estimate the lifetime distribution of train wheels, with random parameters assumed to follow lognormal, Weibull and normal distributions (Freitas et al., 2009). Although the linear random-effect model has advantages, nonlinear models are better suited to capturing complex degradation paths and provide a better fit to nonlinear data. To account for non-monotonic degradation patterns while preserving random effects, additive and multiplicative models have been developed (Bae et al., 2007). General path models are applied because of their simplicity, capability to model continuous degradation paths, and their compatibility with different variance-covariance structures of the response vector (Shahraki et al., 2017). Stochastic dynamics are present in degradation paths due to the uncertainties in the system working environments, measurement errors, and variability among units in a population (Zhang et al., 2018). For continuous degradation processes, models such as the Gamma process, Gaussian mixture and the Wiener process have been considered as independent increment processes and have been applied in the field of reliability analysis (Nicolai et al., 2007). In terms of parameter estimation, both general path models and stochastic process models identify model parameters based on frequentist methods, such as maximum likelihood estimation, or Bayesian alternatives.

Machine learning methods

Both tree-based and neural network-based models have been applied in degradation prediction. For instance, a study on the Melbourne tram network demonstrated that

vehicle acceleration, used to assess ride comfort, can effectively serve as a health indicator for predicting track degradation using a random forest model (Falamarzi et al., 2019). To improve short-term water quality prediction, two hybrid models integrating decision tree algorithms with a data denoising technique were developed and demonstrated high accuracy and stability across six water quality indicators using real-world data from the Tualatin River Basin (Lu and Ma, 2020).

Regarding neural network-based approaches, Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), their variants, and hybrid architectures are among the most commonly used in this field (Guo et al., 2021). For example, the root mean square value of vibration data has been shown to reflect gearbox health conditions, and an extended RNN has been proposed for one-step data prediction, outperforming the ARMA method (Tian and Zuo, 2010). In another study, five machine learning models—including gradient boosting, LightGBM, random forests, extra trees, and XGBoost—alongside two deep learning models, Long Short-Term Memory (LSTM) and LSTM-CNN, have been used to predict battery health states, with the gradient boosting model achieving the best performance (Huotari et al., 2021).

Unlike statistical methods, machine learning methods do not require explicit mathematical formulations and operate under fewer clearly defined assumptions. It is worth noting, however, that these studies generally rely on the availability of sufficient training samples, and there is limited research on improving DL performance in small-sample scenarios.

2.2 Maintenance policy

When a failed component is inexpensive and easy to replace, and its failure does not greatly disrupt operations, run-to-failure maintenance—also called failure-based maintenance—can be a strategic asset management approach, allowing equipment or components to be used until they break down. However, this approach is inappropriate for critical assets because the consequences of failure are too severe to be tolerated. With the collected historical failure data, preventive maintenance policy has been extensively studied since the 1960s, with early works proposing general mathematical frameworks

to decide preventive maintenance actions before the system failure occurs (Barlow and Hunter, 1960). This early work on maintenance optimisation focuses on failure events by modelling failure rates. A comprehensive review work provides an overview of the preventive maintenance policy and summarises the related literature(Ahmad and Kamaruddin, 2012).

Preventive maintenance policies can be broadly categorised by their trigger schemes into Time-Based Maintenance (TBM) and CBM (Ahmad and Kamaruddin, 2012). Both policies aim to minimise the objective of interest, such as minimisation of maintenance cost or maximisation of availability, to decide the appropriate maintenance plan. The main differences between TBM and CBM are listed as follows:

- TBM: TBM depends on the relevant time-to-failure records. It involves deciding the optimal maintenance interval to either minimise maintenance costs or maximise asset availability.
- CBM: CBM is guided by the degradation behaviour and the monitoring data. It aims to determine the optimal degradation threshold to achieve the same goals of minimising maintenance costs or maximising availability.

2.2.1 Time-based maintenance policy

TBM, also known as age-based maintenance policy in some contexts, is an approach used to determine the optimal time or age for replacing a system (Zhao et al., 2017). The first step of TBM is to collect the historical time-to-failure data. Then these data can be used for time-to-failure modelling based on some typical reliability-related characteristics, such as mean time to failure, reliability distribution, and failure rate. The parameters of the time-to-failure model can be estimated from the time-to-failure data and then put into the maintenance scheduling model to output the preventive maintenance intervals. Under this policy, the system is replaced at the scheduled age or at the time of failure, whichever occurs first. The process for a TBM policy is illustrated in Figure 2.1.

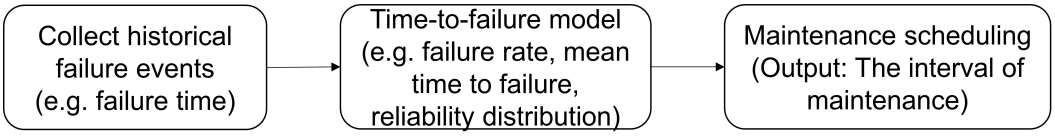

Figure 2.1: TBM process

Failure rate

It is worth noting that a body of research focuses on modelling failure rates (also referred to as hazard rates in some contexts) to develop a reliability function for maintenance scheduling in TBM (Ahmad and Kamaruddin, 2012; Xia et al., 2015; de Jonge et al., 2017).

For non-repairable systems, the failure rate represents the probability of being failed at a given age. It is denoted by $\lambda(t)$ and calculated as:

$$\lambda(t) = \frac{f(t)}{R(t)}, t \geq 0$$

where $f(t)$ is the Probability distribution function (PDF) of failure and $R(t)$ is the reliability function.

One of the conceptual failure rate models is the bathtub curve, which divides the lifetime of an asset into three phases: infant mortality stage (unstable and with a decreasing failure rate), random failures stage (stable and with a constant failure rate), and wear-out stage (unstable and with an increasing failure rate) (Wang et al., 2002; Zeng et al., 2016). This conceptual model contributes to understanding the different stages of asset life and optimising maintenance schedules accordingly.

The Weibull distribution is applied in reliability engineering due to its flexibility in modelling increasing, constant, and decreasing failure rates (Ahmad and Kamaruddin, 2012; Chan and Shaw, 1993; Xia et al., 2015). This flexibility can be achieved by adjusting its shape parameter:

- When the shape parameter is greater than 1, it represents an increasing failure rate.
- When the shape parameter is equal to 1, it represents a constant failure rate.

- When the shape parameter is less than 1, it represents a decreasing failure rate.

In addition, exponential, gamma, and lognormal distributions are also studied (Yin et al., 2013; Proschan, 1963; Jiang, 2010).

For repairable systems, designed to be restored to functionality after a failure, repeated failures are a frequent occurrence. In this context, the term "failure rate" refers to the rate at which failures occur in such systems and can be modelled using a non-homogeneous Poisson process. The failure rate in repairable systems $m(t)$ that varies over time t (Krivtsov, 2007). Given that a failure has just occurred at time T, the probability that the next failure occurs after time t is given by $F(t) = 1 - \exp\left(-\int_0^t m(T+x)\,dx\right), t \geq 0$. The power law function and log-linear function are used for modelling failure rate in repairable systems (Cox and Lewis, 1966; Krivtsov, 2007).

With the parametric form of the time-to-failure distribution and the available time-to-failure data sets, statistical inference techniques like Maximum Likelihood Estimation (MLE) are utilised to estimate the relevant parameters (Fouladirad et al., 2018). Furthermore, when the form of time-to-failure distribution is unknown, few studies have applied the non-parametric kernel density estimator to estimate reliability function (Sidibé et al., 2016).

Maintenance scheduling

TBM aims to determine the optimal maintenance interval to achieve various objectives, such as minimising maintenance costs, maximising asset availability, or addressing multiple objectives that combine both cost reduction and increased availability. Before setting up the objective function, it is important to figure out whether the asset is repairable or not. For non-reparable assets, the replacement policy is applicable. Under this policy, the asset is replaced when its age reaches the optimal preventive replacement age or failure occurs (Ahmad and Kamaruddin, 2012).

For non-repairable assets, a specific repair policy is applied: the asset is replaced at an optimal replacement age, and any failure occurring before that age is addressed with minimal repair. Replacement makes the asset "as-good-as-new", while minimal repair

restores functionality without altering the asset's health state, meaning the failure rate remains unchanged. Additionally, some research has discussed imperfect repair policies, where after an imperfect repair, the asset's state is assumed to be better than before but not as good as new. Instead, it is considered to be at a younger age with a reduced failure rate (Xia et al., 2015).

Renewal reward theory is applied in maintenance scheduling for TBM policy (Arts et al., 2024). Renewal reward theory applies to a process that accrues reward continuously over a renewal cycle, which is defined as the duration between two consecutive replacements (Zhang and Jardine, 1998). The general form of the long run maintenance is defined as:

$$C(t) = \lim_{t\to\infty} \frac{C(t)}{t} = \frac{\mathbb{E}[C(Z)]}{\mathbb{E}[Z]}, t \geq 0$$

where Z is the duration of the renewal cycle, $C(Z)$ is the total maintenance cost of a renewal cycle. Assuming only corrective maintenance and preventive maintenance action are taken, the expected of maintenance cost is given as:

$$\mathbb{E}[C(Z)] = F(Z)C_c + R(Z)C_p,$$

where $F(Z)$ is the probability of failure, $R(Z)$ is the reliability that $R(Z) = 1 - F(Z)$, C_c is the cost of corrective maintenance and C_p is the cost of preventive maintenance. Without loss of generality, $C_c > C_p > 0$.

TBM considering working condition

The ageing process of assets is influenced by various fixed or dynamic working conditions, such as working load, speed, and temperature. These working conditions are modelled as covariates and are incorporated into the failure rate model. The modelling approaches can be divided into three categories: parametric methods, semi-parametric methods, and non-parametric methods.

The Accelerated Failure Time (AFT) model is one of the parametric methods to determine the failure distribution when considering covariates (Hu et al., 2017). In addition, it also simplified the working condition as a piecewise function of time, which

is not only practical but also easy to extend from existing models. This work is based on an accelerated-time intensity Poisson process model where the failure rate function is given by:

$$\lambda(t|\mathbf{x}) = \lambda_0(e^{\mathbf{x}^\intercal\boldsymbol{\beta}}t)e^{\mathbf{x}^\intercal\boldsymbol{\beta}}, t \geq 0$$

where $\lambda(t|\mathbf{x})$ is the failure rate function, λ_0 is the baseline failure rate function that considers the age of the asset, $\mathbf{x}$ is the vector of working conditions, $\boldsymbol{\beta}$ is the coefficient vector that measures the impact of the covariates, and $e^{\mathbf{x}^\intercal\boldsymbol{\beta}}$ is the acceleration factor. if $e^{\mathbf{x}^\intercal\boldsymbol{\beta}} > 1$, the impact of the covariate is decelerated, and vice versa if $e^{\mathbf{x}^\intercal\boldsymbol{\beta}} < 1$. It is a fully parametric model as it specifies a complete probability distribution for the baseline failure rate.

As a semi-parametric model, Cox Propotional Hazard (PH) model is applied to make regression on features and failure time (Debón et al., 2010; Zhang et al., 2014). The Cox PH model is represented by the failure rate:

$$\lambda(t) = \lambda_0(t)e^{\mathbf{x}^\intercal\boldsymbol{\beta}}, t \geq 0$$

where λ_0 is the baseline failure rate. Unlike the AFT model, the Cox PH model does not specify the form of the baseline failure rate, allowing it to be estimated non-parametrically from the data. $e^{\mathbf{x}^\intercal\boldsymbol{\beta}}$ is the parametric part that captures the effect of covariates on failure rate. The combination of parametric and non-parametric components contributes to the robustness of the Cox PH model.

Non-parametric methods model the failure rate function without assuming a specific parametric form, allowing the failure rate to be estimated directly from the data. Techniques such as Bayesian non-parametric methods, kernel smoothing, and additive models are used for this purpose. However, according to the literature, non-parametric models are not as popular as the aforementioned parametric or semi-parametric methods. This is likely due to the increased complexity and computational demands associated with non-parametric approaches, as well as the challenges in interpreting their results compared to more straightforward parametric or semi-parametric models.

TBM considering heterogeneity

The collected time-to-failure data might show heterogeneity patterns caused by latent heterogeneity, such as the differences in manufacturing process, raw materials. To account for the latent heterogeneity, different methods are proposed and developed. These methods are based on the modelling of the failure rate.

Change-point method is applied to break down the failure rate function into segments. The general form of the change-point method is:

$$\lambda(t) = \lambda_1(t)I(1 \leq t \leq t_1) + \lambda_2(t)I(t > t_1),$$

where $I(x)$ is an indicator function, defined as $I(x) = 1$ if x is true, and $I(x) = 0$ otherwise. Though the change-point method can detect shifts in the failure rate, it still assumes a single failure rate function for the entire lifetime of the asset.

To capture the heterogeneity among the population, the frailty model is introduced. This model introduces an additional variability into the failure rate function, known as the frailty term, allowing for a more flexible and accurate construction of the failure rate function (Vaupel et al., 1979; Huber-Carol and Vonta, 2004; Li and Liu, 2016). Taking the Cox PH model as an example to show how to incorporate the frailty term, for an individual i. the failure rate function is given as:

$$\lambda_i(t|\mathbf{x}_i, z_i) = \lambda_0(t)e^{\mathbf{x}_i^\intercal \boldsymbol{\beta}} z_i,$$

where z_i is the frailty term for asset i and it is assumed to follow a specific parametric distribution. A proportional failure rate function with random effects is proposed to generalise the frailty model by allowing certain coefficients to be random variables that follow specific prior distributions (Vaida and Xu, 2000; Hu and Chen, 2020).

Instead of modelling the failure rate, another approach is to model the time-to-failure distribution as a mixture distribution. This method models the time-to-failure distribution as a weighted sum of component distributions of the same type, but with different parameter distributions. Distributions such as Weibull and Lognormal are used as components to model the sub-populations of assets (Bučar et al., 2004; Cheng

and Yuan, 2013).

2.2.2 Condition-based maintenance policy

For TBM policy, the collected data includes the lifetime of each asset, which is used to inform maintenance scheduling. With advancements in modern sensors and wireless communication technologies, the health condition of assets can now be monitored and collected, providing a richer and more comprehensive dataset for maintenance planning (De Jonge and Scarf, 2020; Kim and Makis, 2013; Drent et al., 2023).

Generally, The CBM process consists of three steps, which are illustrated in Figure 2.2 and summarised as follows:

- Condition monitoring. During the monitoring, the condition monitoring signals, such as the vibration of rotating equipment, electoral and temperature, are monitored using certain types of sensors. The monitoring process can be carried out either online or offline. Online monitoring is carried out during the operating state, while offline monitoring is performed when the asset is not running. In addition, the monitoring can also be carried out periodically or continuously (Ahmad and Kamaruddin, 2012).

- Degradation process modelling. The collected condition monitoring signals, also referred to as degradation data, generally contain the health information of the asset and can be further modelled to reflect the degradation process. These degradation processes are stochastic and can be either continuous or discrete. By analysing the degradation process, the failure time distribution and its parameters can be determined (Zhang et al., 2018).

- Maintenance scheduling. Once the degradation reaches a certain threshold, the asset is failed. For CBM policy, it is important to dynamically decide the optimal maintenance threshold to take action in advance of reaching the failure threshold (Elwany et al., 2011).

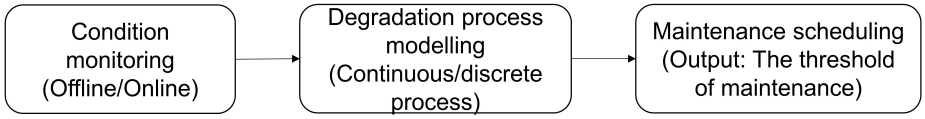

Figure 2.2: The process of CBM

Stochastic degradation process

Stochastic processes serve as fundamental model-based methods, providing benefits in the optimisation of CBM policies. The degradation process is categorised into continuous degradation processes and discrete degradation processes.

For continuous degradation processes, popular models include Wiener processes (Wang et al., 2014a; Zhang et al., 2018), Gamma processes (Wang et al., 2021), and inverse Gaussian processes (Ye and Chen, 2014), selected based on the nature of degradation increment. When the degradation model is a Wiener process with a positive drift, the first-passage time (i.e., the time when the degradation reaches the failure threshold) follows an inverse Gaussian distribution. Additionally, the Wiener degradation process is not strictly increasing; the degradation increment in each unit of time can be either positive or negative. However, monotonicity is inherent in Gamma processes and inverse Gaussian processes, which have only positive increments.

Many systems are subject to random shocks caused by sudden environmental changes (Rafiee et al., 2015). For example, bridges probably experience random shocks such as rainstorms, floods, earthquakes and overloads (Wang et al., 2020). This type of degradation can be modelled using a discrete degradation process, where the arrival of shocks follows a Poisson distribution and the increments caused by these shocks are independent and identically distributed. This shock-based discrete degradation model is also referred to as the Compound Poisson degradation model (Junca and Sanchez-Silva, 2013; Drent et al., 2023).

Given the degradation process $X(t)$, the lifetime of system is defined as:

$$T = \inf\{t : X(t) \geq L | X(0) < L\},$$

where L is the failure threshold. According to first-passage-time theory, when modelling

the degradation process as a Wiener process, the first passage time (also referred to as the remaining useful time) for a fixed level follows an inverse Gaussian distribution (Bian and Gebraeel, 2012; Ye et al., 2013). This distribution of remaining useful time can further enhance maintenance scheduling.

In addition, the selection among these stochastic processes depends not only on the underlying mechanism but also on criteria such as AIC (Akaike Information Criterion) and BIC (Bayesian Information Criterion) (Nguyen et al., 2018). It can be summarised that, regardless of the type of stochastic process selected, the degradation is assumed to be stationary with independent increments, which ensures favourable mathematical properties (Wang et al., 2014c).

Maintenance scheduling

Given the degradation model and the estimations of the parameters, the maintenance scheduling models in CBM can be categorised into two classes: renewal reward theory and Markov Decision Process (MDP) (Arts et al., 2024). In addition, renewal cycle theory is applied in both TBM and CBM. The main difference between TBM and CBM when applying for renewal reward theory is the derivation of reliability depends on the type of data, i.e., the failure event or condition monitoring signals.

MDPs are applied to make optimal sequence decisions under uncertainty (Elwany et al., 2011; Drent et al., 2023). A standard MDP consists of 4 elements that can be represented as a tuple (S, A, P_a, R_a), where S is the state space, A is the set of actions, P_a is the transition probability matrix when taking action a, and R_a is the corresponding reward. When incorporated with CBM policy, the state space generally includes the age of the asset and the degradation level. The action space includes the possible action that the decision-maker can make, such as preventive maintenance, or corrective maintenance. The reward is the cost for corresponding maintenance action. The transition probability matrix can be derived from the stochastic degradation model when taking the degradation levels as states.

When modelling maintenance policies via MDP in CBM, both finite horizon and infinite horizon approaches are broadly discussed. For the infinite horizon problem, it

is assumed that the asset operates indefinitely, and this problem can be solved using policy iteration or value iteration (Elwany and Gebraeel, 2008; Chen et al., 2015). One proposed advantage of modelling an infinite horizon is that the asymptotic long-run cost rate converges to the cost ratio with renewal reward theory, which is relatively easy to obtain. However, most assets are designed with a limited operating time due to factors such as demand changes or product upgrades (Zhang et al., 2024a). Thus, it is necessary to model a finite horizon for such scenarios. Backward dynamic programming can be applied to solve finite horizon MDPs (Liu et al., 2021). Additionally, the optimal maintenance policy can be represented as a control limit policy, regardless of whether the problem is modelled as a finite or infinite horizon.

Some other extensions of MDP are also incorporated into CBM policy. Partially Observable MDP (POMDP) is applied when the underlying state cannot be directly observed (Kim, 2016; Deep et al., 2023). In addition to a standard MDP, POMDP includes a set of observations, a belief vector, and an emission matrix. The belief vector represents the decision-maker's belief about the underlying state, while the emission matrix contains probabilities that link the underlying states to the observations. However, computing an optimal policy for a POMDP is an intractable problem due to the "curse of dimensionality" (Lee et al., 2007). It has been noted that solving a finite POMDP is PSPACE-complete and finding an optimal policy over an infinite horizon is undecidable. In addition to POMDP, Semi-MDP is applied in CBM when the sojourn time in each state is a general continuous random variable. Furthermore, POMDP and Semi-MDP can be combined to address scenarios involving both unobservable states and random transition sojourn times (Khaleghei and Kim, 2021).

CBM considering working condition

In order to incorporate working conditions that influence degradation processes, various studies have extended and generalised variants from conventional stochastic degradation models. Static and dynamic working conditions have been studied. Static working condition means the working conditions remain stable when operating (Zhao et al., 2019). However, dynamic means that the working condition is time-varying, which

can be represented as a function of time, or the evolution of the working condition is modelled by a continuous-time Markov process (Hu et al., 2021; Bian et al., 2015).

One of the frequently used approaches utilises a link function to describe the effects of working conditions on the parameters in the stochastic degradation process. Linear, power law, and exponential functions are selected as link functions based on the goodness-of-fit (Ye and Xie, 2015). Existing works link the working condition to the drift parameter in both the Wiener process and inverse Gaussian process (Liao and Tian, 2013), the shape parameter in the Gamma process (Bagdonavicius and Nikulin, 2001; Zhao et al., 2019), and the occurrence of shock in the Compound Poisson (CP) process (Zhu et al., 2015).

CBM considering component heterogeneity

The issue of component heterogeneity has been addressed in existing literature, regardless of whether working conditions are taken into consideration or not. Observed components within the same population may exhibit different degradation paths due to various reasons, such as variability in raw material. For unobservable heterogeneity, the random effects model treats certain parameters in the stochastic degradation model as random variables that follow specific parametric distributions (Ye and Xie, 2015; Lawless and Crowder, 2004; Peng and Tseng, 2009; Ye and Chen, 2014). It is worth noting that the maintenance policy exhibits different structural properties when heterogeneity is considered (Chen et al., 2015; Zhang et al., 2016). Considering component heterogeneity in maintenance policy, the earliest work models the degradation paths as Brownian motion and sets the drift parameter as a random variable that follows normal distribution (Elwany et al., 2011). The maintenance decision is modelled as an infinite horizon MDP to minimise the maintenance cost and find the optimal Preventive Maintenance (PM) threshold. The structural property of the maintenance shows that the PM threshold increases with the component age. This work has been further extended to other stochastic degradation processes, such as Gamma (Liu et al., 2021), inverse Gaussian process (Chen et al., 2015).

One more related and recent work on discrete degradation paths modelled by the CP

process has addressed the heterogeneity of components (Drent et al., 2023). This work considers heterogeneity in the number of shocks and the shock magnitude separately. Specifically, the number of shocks is assumed to follow a Poisson distribution, with the rate parameter regarded as a Gamma-distributed random variable. The shock magnitude is assumed to follow a one-parameter exponential family distribution, with the parameter also exhibiting unobserved heterogeneity and regarded as a Gamma-distributed random variable. The prior parameters are initialised using historical data, and further updated when new observations are available. This process follows the Bayesian learning paradigm, which has been frequently used in the literature on price and inventory management (Wang, 2021; Harrison et al., 2012; Chuang and Kim, 2023). Similar to the aforementioned CBM policies, the MDP has been applied to derive the optimal PM threshold, which has been proven to increase with the component age.

2.2.3 Summary and research trends

TBM and CBM policies have been studied over the past decades, forming the foundation of modern maintenance optimisation. TBM primarily relies on historical failure data to determine optimal preventive replacement intervals, while CBM leverages real-time condition monitoring to dynamically schedule maintenance activities. TBM typically uses reliability models based on failure rates or time-to-failure distributions, including parametric, semi-parametric, and non-parametric approaches. CBM, in contrast, utilises stochastic degradation models such as Wiener, Gamma, or Compound Poisson processes, in conjunction with decision-making frameworks like MDPs and their extensions. Both approaches aim to improve reliability and reduce maintenance costs, but follow distinct modelling frameworks and data requirements.

Despite substantial progress, several limitations remain. TBM models often assume stationary operating environments and fail to accommodate time-varying conditions or latent heterogeneity. CBM models, while more flexible, rely heavily on high-quality sensor data and face challenges related to model complexity, computational demands, and interpretability. In multi-component systems, dependencies among components further complicate scheduling, particularly when integrated with production planning

or resource constraints.

Recent research has increasingly focused on extending maintenance policies to account for system-level interactions. TBM and CBM approaches have evolved from applications in single-component systems (Das and Acharya, 2004; Elwany et al., 2011) to more complex multi-component systems (Scarf et al., 2009; Nakagawa and Yasui, 2005; Liu et al., 2021; Chen and Hao, 2025). In such systems, various forms of interdependence—including stochastic, structural, and economic dependencies, pose challenges to the design and implementation of effective maintenance strategies (Shi and Zeng, 2016). Structural dependence refers to how system functionality is influenced by the configuration of components, which can follow series (Dao and Zuo, 2015), parallel (Berrichi et al., 2009), mixed series-parallel (Xia et al., 2013), or k-out-of-n arrangements (Liu et al., 2024). Stochastic dependence is typically captured through failure rate or stochastic process models (Fan et al., 2021; Liu et al., 2021; Rasmekomen and Parlikad, 2016). Economic dependence implies cost-saving potential through group or opportunistic maintenance, where components are maintained jointly—even if some are serviced earlier or later than ideal—to reduce overall maintenance costs (Zhou et al., 2009; Xia et al., 2015). These interdependencies affect not only the reliability modelling of the system but also the optimisation of maintenance decisions.

There is increasing emphasis on integrating maintenance with broader decision-making domains such as production planning, inventory control, and resource management, supported by Bayesian learning, reinforcement learning, and hybrid models for adaptive, data-driven maintenance. The joint optimisation of maintenance and production planning has been explored in studies where asset degradation is influenced by production load, making it necessary to coordinate scheduling decisions across both domains (Cassady and Kutanoglu, 2005; Paprocka, 2019; uit het Broek et al., 2020). Similarly, maintenance policies have been jointly optimised with spare parts inventory management (Elwany and Gebraeel, 2008; Wang, 2012; Vaughan, 2005), where challenges arise due to the intermittent and lumpy nature of spare parts demand, which is directly driven by the maintenance schedule. Integration with resource management has also gained attention, particularly when maintenance-related resources such

as personnel, tools, and time are limited (Glazebrook et al., 2005; Martorell et al., 2010; Do Van et al., 2013). Selective maintenance, for example, has been formulated as a MDP that incorporates stochastic action durations and resource constraints, and solved using a rollout-based approximate dynamic programming algorithm enhanced with heuristics and machine learning to enable scalability (Zhang et al., 2024b). In the offshore wind sector, where maintenance costs are particularly high, a holistic opportunistic maintenance strategy has been proposed that integrates degradation trends, failure events, and weather forecasts to optimise turbine grouping, vessel routing, and scheduling, yielding substantial cost reductions in real-world applications (Si et al., 2025).

In summary, maintenance policy research is shifting towards system-wide, data-driven strategies that integrate maintenance with broader operational decisions. Emphasis is being placed on handling uncertainty, exploiting interdependencies, and leveraging learning-based methods to enhance adaptability and cost-effectiveness in complex, resource-constrained environments.

2.3 Positioning in the literature

In this section, we position each chapter in the literature, highlight its intersections with other research fields, and identify gaps in the current literature for each chapter.

2.3.1 Positioning of Chapter 3

Chapter 3 focuses on degradation prediction using DL methods, as reviewed in Section 2.1. However, within the field of degradation prediction using time series data, although various data-driven methods have been proposed, there has been no study addressing scenarios with insufficient data.

Some works in the research field of Computer Vision tackled the overfitting problem caused by insufficient training data (Park et al., 2019; Ke et al., 2019; Cai et al., 2017). Various image-based Data Augmentation (DA) methods—including rotation (randomly rotating the image), solarising (inverting pixel values above a threshold),

histogram equalisation (enhancing contrast), and posterising (reducing the number of pixel intensity levels)—are applied to increase the training data size. These transformations introduce visual diversity while preserving the essential shape and structural patterns relevant. Moreover, the optimal DA strategy with the best combination of DA methods and suitable parameters is explored to enhance the performance of the DL algorithm in the classification of images (Cubuk et al., 2019). The development and enhancement of DL algorithms for time series data are relatively new and there is still no related work for degradation prediction.

The most relevant work to our study is the enhancement of DL algorithms with limited time series data by applying DA method (Bandara et al., 2021; Demir et al., 2021). The basic DA techniques of time series data focus on transforming in the time domain, frequency domain and time-frequency domain respectively (Wen et al., 2020). Decomposition methods, statistical generative methods and some learning methods like the deep generative model and auto augmentation are the recent works for advanced DA techniques. However, it can be observed from these studies that most existing works focus on classification tasks—such as predicting discrete health states or failure types—while only a few address regression tasks, where the goal is to estimate continuous values such as degradation levels or remaining useful life (Fons et al., 2021; Iwana and Uchida, 2021). So far, only one work explored the combination of a DA method and Transfer Learning (TL) strategy and gained improvement on LSTM in the presence of limited time series data sets (Bandara et al., 2021). Our work is distinguished from the previous studies in three aspects. Firstly, we adopt multiple DA techniques and construct a novel framework based on these augmented data. Secondly, we integrate both data selection and reweight to improve the quality of training data. Thirdly, DA methods, data dropout, data reweight and TL are stacked in the whole framework to get better performance on degradation prediction.

2.3.2 Positioning of Chapter 4

Chapter 4 aligns with the literature on TBM policy considering the working conditions. However, the requirement for prior information about the failure is difficult

for decision-makers, thereby limiting the application (Ahmad and Kamaruddin, 2012). When applying the parametric methods, the optimisation results will be amplified if the underlying parametric form is misspecified. Although non-parametric methods alleviate assumptions on the reliability distribution, separating prediction and optimisation can result in sub-optimality. Therefore, there is a need to develop a distribution-free framework that integrates prediction and optimisation in preventive maintenance problems.

Recently, one work has been devoted to the condition-based maintenance problem based on a fully data-driven approach with an unknown deterioration process and unknown failure behaviour (Cai et al., 2023). Under the setting of the CBM policy, the work aims to decide the optimal degradation threshold with a minimal maintenance cost rate. In order to obtain the decision directly, the cost rate function can be adjusted into the data-driven formula by replacing the failure probability with approximated expressions from the collected runs-to-failure data. Notably, the approximated failure probability becomes more accurate when more runs-to-failure data becomes available. Similar to our work, the maintenance decision is made by finding a close form of the cost rate function using the data-driven method instead of putting any prior assumption to underlying failure distribution. To distinguish it from our work, our work not only provides the theoretical bound of the approximated maintenance cost rate function for the TBM problem but also incorporates the working conditions into the preventive maintenance decision policy.

Beyond the maintenance decision area, several works have explored the integration of prediction and optimisation. To the best of our knowledge, two works have been dedicated to proposing general integration frameworks. One of the works constructs the objectives as a weighted linear combination of each sample, where the weights can be estimated using collected data sets through machine learning algorithms like k-nearest neighbours algorithm, trees, and kernel methods (Bertsimas and Kallus, 2020). Another paradigm known as "smart predict, then optimise" has introduced a general framework to minimise decision errors during the training of the prediction model (Elmachtoub and Grigas, 2022). The loss function during the training process is the absolute gap of

cost between the predicted decision and the optimal decision, rather than the prediction accuracy typically employed in traditional two-step optimisation methods.

In addition, two studies have explored solving feature-based newsvendor problems using big data sets. One work takes the optimal solutions as a linear combination of features while minimising the empirical cost function. Linear programming and kernel optimisation methods are leveraged to find solutions and show better performance compared with sample average approximation and other two-step methods (Ban and Rudin, 2019). Another study considers the optimal ordering number as the label for the machine-learning algorithm and adopts the empirical cost minimisation as the loss function (Oroojlooyjadid et al., 2020). The work leverages deep learning frameworks as the learning algorithm. The effectiveness of the proposed method is demonstrated through experiments conducted on small data sets, real data sets, and randomly generated data sets.

Chapter 4 is inspired by the end-to-end machine learning algorithm, which refers to the end goal of the learning algorithm directly from the raw input of data without any intermediate step (Donti et al., 2017). Recently, there has been a growing body of research that applies this paradigm to various domains, including finance, text recognition, and also feature-based inventory problems that consider uncertain demand and lead time (Qi et al., 2023; Bengio, 1997; Wang et al., 2011). In this line of work, the main challenge is to find the labels for the training data set first and subsequently apply well-established machine learning or deep learning algorithms. However, to the best of our knowledge, no prior work has explored the integration of such a framework in the preventive maintenance policy considering varied working conditions.

2.3.3 Positioning of Chapter 5

Chapter 5 aligns with the literature on CBM policy considering heterogeneity and time-varying working conditions. Similar to the aforementioned CBM policies, the MDP has been applied to derive the optimal PM threshold, which has been proven to increase with the component age (Drent et al., 2023). To distinguish from this work, our contribution lies in being the first to incorporate working conditions into the

maintenance policy with heterogeneous components leveraging real-time data. Under such a scenario, developing an analytically tractable result becomes challenging due to the increased complexity of the state space.

Chapter 5 is also related to Contextual MDP (CMDP). In the literature, CMDP was first proposed as an extension of MDP to integrate environmental information as a form of contexts (Hallak et al., 2015; Benjamins et al., 2022). To the best of our knowledge, the earliest paper describes a contextual MDP as a set of MDP models that share the same state and action space, with each MDP model corresponding to an environmental context (Hallak et al., 2015). To enhance its generality, a unified contextual MDP framework has been proposed to accommodate both MDP and partially observable MDP. An application area of CMDP is personalised service decisions, such as web advertising and precision medicine (Modi et al., 2018). For instance, in precision medicine, optimal medication dosages can be determined using patient-specific information, such as gender and age, extracted from electronic health records.

However, CMDP faces challenges such as limited observations under the same context, the trade-off between exploration and exploitation, and arbitrarily large context space (Bastani and Bayati, 2020). These challenges have been addressed in the existing research. To efficiently utilise the observations, the Q function, also known as the action-value function, is approximated via a set of functions once a low Bellman rank is achieved (Jiang et al., 2017). During the approximation, the valid functions are then selected when the Bellman error is low, which further contributes to optimistic exploration (Jiang et al., 2017). When facing large context space, the context representation can be factorised into sub-blocks and referred to as block CMDPs (Sodhani et al., 2021).

Although this thesis integrates working conditions into the MDP model for optimal maintenance control, our work differs from CMDP in the following aspects. First, we incorporate dynamic working conditions that vary over time periods, wherein CMDP typically deals with static contexts or changes occurring at a much slower pace than the decision periods. Second, the CMDP framework lacks the considerations of the underlying heterogeneity, which we address comprehensively in this thesis. Additionally,

unlike CMDP which optimises the tradeoff between exploration and exploitation, our work focuses on exploitation. This is due to the fact that exploration in a relatively short episode, such as the maintenance decision epochs, may lead to a sub-optimal decision (Sutton et al., 1998; Siraskar et al., 2023).

2.4 Summary

This chapter provides a comprehensive review of the literature on two critical tasks in asset management: health indicator prediction and maintenance policy. The discussion begins with an in-depth exploration of health indicator prediction, where three primary methodologies are highlighted: model-based, physics-based, and data-driven approaches. Model-based methods rely on mathematical models to forecast asset health, while physics-based methods use the physical principles governing the asset's operation. Data-driven methods, on the other hand, utilise statistical and DL techniques to predict asset health based on historical data.

The chapter then shifts focus to maintenance policies, specifically reviewing TBM and CBM, along with their various extensions. TBM policies are scrutinised for their scheduled approach to maintenance, highlighting the importance of regular intervals regardless of the asset's condition. In contrast, CBM strategies are analysed for their reliance on real-time data and working conditions, enabling more responsive and potentially cost-effective maintenance actions. Furthermore, the chapter contextualises the positions of Chapters 3, 4, and 5 within the broader scope of the literature, and addresses the existing research gaps in the literature. It clarifies how each chapter contributes to the overall understanding of asset management.

In summary, this chapter systematically reviews the literature on degradation prediction and maintenance policies, offering a detailed examination of various methodologies and their applications. It not only integrates the discussions of Chapters 3, 4, and 5 but also points out the critical gaps in existing research, paving the way for future contributions to the field.

Chapter 3

Extending the Applicability of Deep Learning Algorithms to Degradation Prediction with Limited Time Series Data

DL algorithms, such as deep neural networks have shown advantages over classical time series prediction methods. However, the performance of DL algorithms depends on the size of the data set with a risk of overfitting for smaller data sets. To extend the applicability of DL algorithms in degradation prediction with such relatively small datasets, this Chapter proposes DA methods to increase the data volume by effectively generating synthetic data. Different from the existing studies that simply mix synthetic and real data without considering the selection of synthetic samples or the weight of each sample, a novel method is proposed to generate, select and reweight synthetic samples to improve the prediction accuracy. After generating a collection of time-series data with multiple DA methods, this Chapter develops an influence function to select the effective synthetic samples and then reweight the selected samples using the gradient descent method. To improve the accuracy of DL algorithms with the mixture of synthetic and real samples, we pre-train the DL models drawing upon the idea of transfer learning and

use real samples to adjust the model parameters with a small learning rate. Simulation experiments are conducted under various settings to illustrate the effectiveness of the proposed method and indicate the scenarios where our approach exhibits its advances. In addition, we describe an illustrative example that involves predicting the values of health indicator in a real wastewater treatment plant, where the data of system health indicator is collected daily. Compared with classical forecasting models, the developed approach indicates superior performance in degradation prediction. The results from the illustrative example show that the proposed method can improve the prediction accuracy of DL algorithms in a less data-abundant scenario.

3.1 Introduction

Degradation prediction is fundamental for designing and providing cost-effective maintenance strategies. Health indicators refer to condition-monitoring signals that reflect the degradation state of a system over time. These indicators are typically collected as time series during system operation and are used to assess the system's health and support predictive maintenance decisions. For instance, the failures of the traction motor can be reflected in the changes in the temperature signal, which is a sequence of data points collected over an interval of time (Dong et al., 2022). Also, the ratio of current capacity to the initial capacity of a battery is defined as the health indicator for lithium-ion batteries to assess the state of health (Huotari et al., 2021). Prediction of these health indicators is useful to capture the dynamic degradation trend based on historic time-stamped data (Han et al., 2019). In the research field of degradation prediction with time-series data, various methods have been developed, including statistical forecasting models, ML models and DL models, to improve the accuracy of prediction (Tian and Zuo, 2010).

The study of statistical forecasting models for degradation prediction with time series data has been widely developed. Among these classical prediction algorithms, the Autoregressive Integrated Moving model (ARIMA) has continuously received scholarly attention to deal with time series data (Liu et al., 2020). Nevertheless, conventional statistical forecasting models encounter challenges when processing and predicting datasets

characterised by non-linear relationships.(Qin et al., 2017; Kabir et al., 2024). Nowadays, ML algorithms such as Decision Tree, Random Forest, Support Vector Machine emerge as the forefront in time series prediction (Zhang et al., 2021; He et al., 2017). As one of the subfields of ML, DL shows good performance with multi-layers to represent the latent features of data at a higher level (Goodfellow et al., 2016; Ajayi et al., 2020). Among DL algorithms, RNN has been designed to deal with time-dependence data. In order to overcome the gradient exploding problem of RNN, variants of RNN have been proposed, i.e., LSTM, Gate Recurrent Unit (Yu et al., 2019). Nevertheless, the performance of DL algorithms highly relies on the size of input data since insufficient data usually leads to overfitting with poor prediction performance in the test data set (Stødle et al., 2023). Engineering systems that are subject to discrete inspections, such as wastewater treatment systems generally measure parameters related to degradation two or three times a week (Choi and Park, 2001). There can be relatively few inspection observations that limit the applicability of DL algorithms (Fu et al., 2020).

In less data-abundant settings, DA and TL are two effective ways to avoid overfitting in DL algorithms (Bandara et al., 2021). DA techniques aim to generate artificial data to increase the sample size of available training data and the effectiveness has been verified in various applications like Computer Vision (Shorten and Khoshgoftaar, 2019), Recognition Technology (Cai et al., 2022). Also, there is work that generates samples with high similarity and diversity in gearbox fault diagnosis (Chen et al., 2022). Nonetheless, DA methods are mainly applied in classification tasks as aforementioned applications. In the research field of degradation prediction with time series data, DA approaches like data jittering and deep generative models have been proposed and successfully applied in regression tasks (Fawaz et al., 2018; Nourani et al., 2021). On the other hand, the challenge induced by the shortage of data can also be alleviated by transferring the pre-trained model from a source domain with abundant data to a target domain that has limited observational data in prediction with time series. In the TL process, the common features can be captured in the pre-trained model and then transferred to the target domain with a relatively small learning rate (Weiss et al., 2016; Tan et al., 2018). Moreover, in the existence of DA, the effectiveness of TL

prediction with time series has been verified through comparison with simply pooling the real and artificial data (Bandara et al., 2021). However, in the research field of degradation prediction with time series data, while various data-driven methods have been proposed, no study has investigated the scenario with a lack of sufficient data.

To the best of our knowledge, this is the first work that applies DA in degradation prediction. Different from existing work in prediction with limited time series data, when applying the DA methods, the training data is appropriately pre-processed and selected to improve its performance instead of leveraging the augmented samples directly to train the DL models. Firstly, a subset of training data should be selected from the augmented samples, since not all of the augmented samples are favorable. Secondly, the weight of augmented data should be assigned considering different contributions to the prediction task. Thirdly, the combination of augmented data with real data needs further exploration in the remit of TL. However, there is no work integrating the selection, reweight and TL to the augmented samples based on multi-DA methods.

In this Chapter, we propose a novel framework to integrate data selection and data reweight to optimize the synthetic samples generated by DA techniques, then apply TL between the synthetic samples and real data. The proposed method is named as ADRT, which integrates data Augmentation, synthetic samples Dropout, Reweight and Transfer learning. Moreover, we utilise LSTM, one of the state-of-the-art methods for dealing with sequence data (Karevan and Suykens, 2020), as the baseline model. The proposed method is called LSTM_ADRT in the following. Finally, a numerical experiment is conducted using simulated data and an illustrative example is performed with a real data set from a wastewater plant to show the effectiveness of the proposed method. The contribution of the study is summarised as follows:

1. We increase the size of training data by integrating multi-DA methods in degradation prediction with limited observational data.

2. We develop an approach to select the important augmented data with dropout and reweight, to improve the prediction performance before channelling it into the DL algorithms.

3. Instead of simply mixing the real data and synthetic data, we train the DL models based on the TL scheme in advance and then use the real samples to adjust the parameters with a small learning rate so as to improve the effectiveness of the approach.

The remainder of this Chapter is organised as follows. The proposed ADRT framework is discussed in Section 3.2. The general description and elements of the approach are discussed in detail, including data augmentation, pre-processing, sample dropout, reweight and TL scheme. Section 3.3 presents the experiment results using simulated data. Section 3.4 presents an illustrative example of a wastewater treatment plant to illustrate the developed approach. Extensive experiments are conducted to show the effectiveness of the approaches and investigate various factors influencing the algorithm's performance. Finally, concluding remarks and future research directions are provided in Section 3.5.

3.2 Method description

3.2.1 Overview of the proposed method

The overall structure of the proposed method ADRT is given in Figure 3.1. Six steps are involved in this framework including data collection, augmentation, pre-processing, dropout, reweight, and transfer learning. With the collected time series data of health indicator, we use multi DA techniques to generate a collection of synthetic data based on real data. Subsequently, we normalise these data to the range of [0,1] and then use the moving time window to generate synthetic samples during the pre-processing stage. In the step of sample dropout, the top negative influential samples generated in the first stage are dropout. In addition, the top positive influential samples are kept for the data reweight. The influence of each sample is determined by assessing its effect on the prediction accuracy of the real dataset when the corresponding sample is removed. A positive influence value indicates that the synthetic sample can enhance prediction accuracy. In the last step, the designed TL scheme is applied to retain the model with a small learning rate to adjust the model parameters in the transfer learning stage.

By and large, DA, sample dropout, reweight and a designed TL scheme are stacked to enhance the DL algorithm step by step. The steps mentioned above will be further discussed in the following sections.

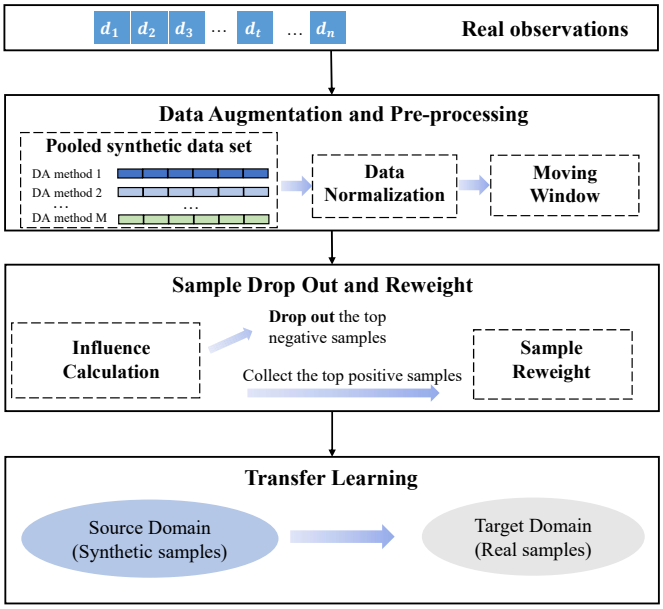

Figure 3.1: Overview of the ADRT framework

3.2.2 Data augmentation and pre-processing

Data augmentation methods

To overcome the overfitting of DL algorithms with a small data set, synthetic data is generated by leveraging DA methods. In this Chapter, we utilise DA techniques in time series data including jittering, scaling permutation, magnitude warping, window slicing and window warping (Fawaz et al., 2018; Iwana and Uchida, 2021).

One of the simplest ways to generate synthetic data is **jittering**, which generates artificial data by adding noise to the real data. The process of jittering can be defined as:

$$D_{\mathrm{A}}^{\mathrm{J}} = [d_1 + \xi_1, ..., d_i + \xi_i, ..., d_n + \xi_n] \tag{3.1}$$

where $D_{\mathrm{A}}^{\mathrm{J}}$ is the synthetic data sequence augmented by jittering. $d_1, \ldots, d_n$ represent

the input real data sequence, where d_i denotes the i_{th} data element. ξ_i is the noise, which follows the Gaussian distribution $\xi_i \sim N(0, \sigma_\xi^2)$ and can be added into the time series at each step.

Scaling aims to change the magnitude of time series with a random scalar value α_{s}:

$$D_{\mathrm{A}}^{\mathrm{S}} = [\alpha_{\mathrm{s}} d_1, ..., \alpha_{\mathrm{s}} d_i, ..., \alpha_{\mathrm{s}} d_n], \tag{3.2}$$

where the random scaling parameters α_s follows the Gaussian distribution $\alpha_s \sim N(1, \sigma_{\alpha_{\mathrm{s}}}^2)$, $D_{\mathrm{A}}^{\mathrm{S}}$ is the synthetic data set augmented by scaling.

Permutation as one of the DA methods has been proposed to disrupt the order of time series (Um et al., 2017). We first slice the data into N_{S} same-length segments and randomly permute the segments to create a new time series. Therefore, the number of segments N_{S} is the critical variable ranging from 1 to 5 (Um et al., 2017). The data set generated by permutation is noted as $D_{\mathrm{A}}^{\mathrm{P}}$.

Magnitude Warping is a specific augmentation method to warp the time series data by a smoothed curve and is defined as:

$$D_{\mathrm{A}}^{\mathrm{M}} = [M_1 d_1, ..., M_i d_i, ..., M_n d_n] \tag{3.3}$$

where $D_{\mathrm{A}}^{\mathrm{M}}$ is the synthetic data set augmented by magnitude warping and $M_1, ..., M_t, ..., M_n$ is sequence generated by interpolating a cubic spline $S(u)$ with the knots parameters $u = u_1, ..., u_m, ..., u_M$. These knots parameters are random values and follow the Gaussian distribution $N(1, \sigma_M^2)$. σ_M and the total number of knots M are both the parameters of magnitude warping.

Window Slicing is similar to cropping in the image data augmentation by slicing time steps as:

$$D_{\mathrm{A}}^{\mathrm{WS}} = [d_\varphi, ..., d_i, ..., d_{W+\varphi}] \tag{3.4}$$

where $D_{\mathrm{A}}^{\mathrm{WS}}$ is the synthetic data set augmented by window slicing and the parameter W is the size of the window which can be decided by the warp ratio R_{WS} multiply the length of the sequence, and φ is a random integer in the range of 1 to $n - W$. Thus, R_{WS} is the parameter of window slicing. After window slicing, the time series data

are interpolated back to the original length to maintain a consistent input size for the model.

Window Warping is a data augmentation technique adapted from time warping which is more time-series specific. It involves selecting a random segment (window) of the time series and altering its temporal structure by compressing (speeding up) or stretching (slowing down) that segment along the time axis. The rest of the sequence, outside the selected window, remains unchanged. This operation preserves the overall pattern of the sequence but introduces local time distortions. The degree of warping is controlled by a warp ratio R_{WW}, which defines the scaling factor applied to the selected window. The rest of the sequence remains unchanged. The resulting augmented dataset is denoted as $D_{\mathrm{A}}^{\mathrm{WW}}$, and the value of R_{WW} can be adjusted experimentally.

The effectiveness of different DA methods varies with the data set. Thus this Chapter mixes all the generated data and then selects the appropriate data adaptively. In application, the parameters in these DA methods can be found in the existing works (Fawaz et al., 2018). The examples of different data augmentation methods are illustrated in Figure 3.2. A broad range of data augmentation techniques, including permutation, scaling, and jittering, are adopted to generate diverse augmented samples. Given the initial uncertainty regarding the suitability of different augmentation methods, a data dropout and reweighting process are implemented to automatically evaluate, select, and adjust the augmented samples. This strategy ensures that the final training set maintains good quality information while gaining improved diversity and robustness through augmentation.

Data pre-processing

After generating augmented data by multi DA methods, we mix all of the synthetic data set to generate a pooled data set as $D_{\mathrm{A}} = \left\{ d_{ij}^{\mathrm{a}}, 1 \leq i \leq P, 1 \leq j \leq N \right\}$, where d_{ij}^{a} is the j^{th} synthetic data point generated by the i^{th} DA method, P is the total number of DA methods and N is the length of data generated by each DA method. Subsequently, in order to reduce the calculation time, normalisation is a necessary step to scale down the data. Min-Max normalisation is a simple and widely used technique that can keep

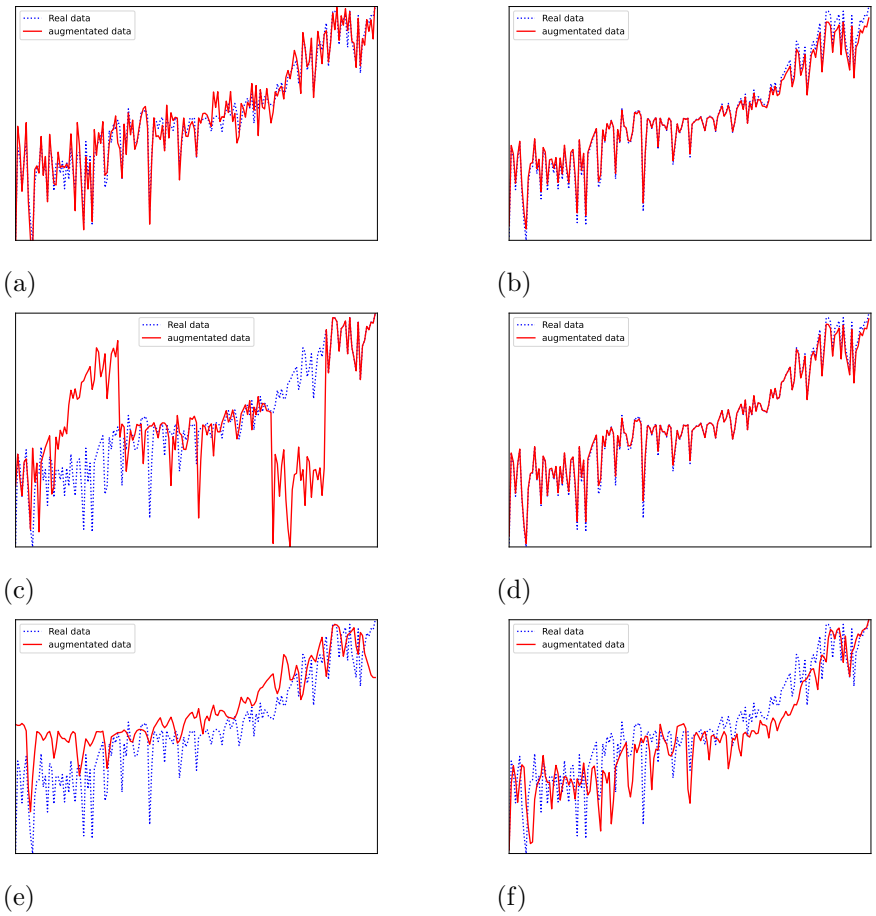

Figure 3.2: Illustration of data augmentation methods. (a) Jittering (σ_ξ=0.05). (b) Scaling ($\sigma_{\alpha_s} = 0.5$). (c) Permutation ($N_S = 12$). (d) Magnitude Warping ($\sigma_M = 0.1, M = 4$). (e) Window Slicing ($R_{WS} = 0.6$). (f) Window Warpping ($R_{WW} = 0.5$)

.

the relationships among the data (Patro and Sahu, 2015). Here, we utilise the Min-Max normalisation technique for the pooled synthetic data set D_A according to the following equation:

$$d_{ij}^{n} = \frac{d_{ij}^{a} - d_{min}^{a}}{d_{max}^{a} - d_{min}^{a}}, d_{ij}^{a} \in D_A \tag{3.5}$$

where d_{ij}^{n} is the synthetic data after normalisation, d_{min}^{a} and d_{max}^{a} are the minimal and maximal value among the pooled synthetic data before normalisation respectively.

The moving window method is used as a data preparation strategy, rather than a prediction model itself. It constructs training samples by sliding a fixed-size window

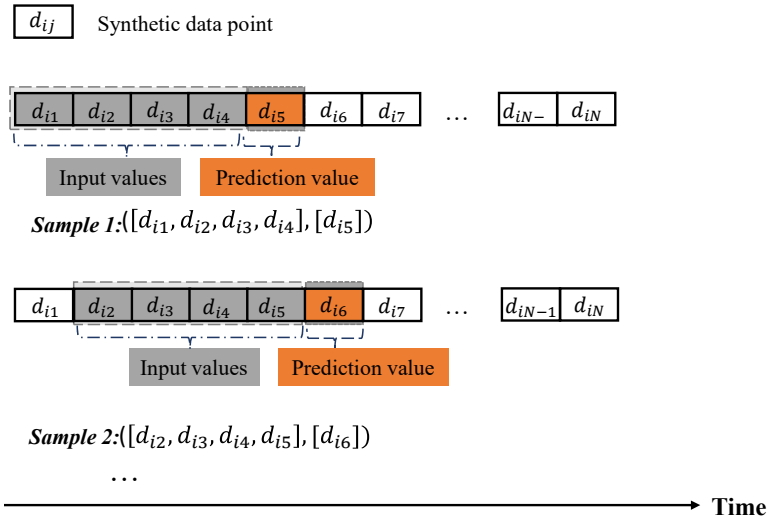

Figure 3.3: Illustration of moving window method

over the historical time series to create multiple input–output pairs. This allows the model to incorporate the most recent information and simulate dynamic updates in a rolling prediction setting. A sample generated by the moving window method can be described as the previous l data points used to predict the next data points at time t. The moving window method is used for each synthetic data set separately. The whole process of the moving window method is described in Figure 3.3, where the moving window length $l = 4$, indicating the previous four data points are set as the input data for the DL model and the next data point is the corresponding output value. The pooled synthetic samples generated moving widow are denoted in the form of pairs $\{(x_i^{\mathrm{s}}, y_i^{\mathrm{s}}), 1 \leq i \leq J\}$, where x_i^{s} is the input value for prediction from the i^{th} synthetic sample after data pre-processing, y_i^{s} is the associated output value, J is the size of the synthetic samples after pre-processing. It is worth noting that all of the data used in the moving window method come from the pooled synthetic data set after normalisation. We also define the real observations after pre-processing as $\{(x_i^{\mathrm{r}}, y_i^{\mathrm{r}}), 1 \leq i \leq K\}$, which is the pre-processed real samples, where $(x_i^{\mathrm{r}}, y_i^{\mathrm{r}})$ is the i^{th} real sample after pre-processing.

After normalisation and the moving window method, the pooled synthetic samples still need to be properly selected and assigned weight before being treated as the training

data set. The details of data dropout and reweight will be discussed in the following.

3.2.3 Sample dropout and reweight

Sample dropout

The pooled synthetic samples are generated by multi DA methods that may have different contributions to the prediction task. Thus, the quality of these pooled synthetic samples is unknown for the prediction task. Before treating these synthetic samples as the training data set, unfavorable samples should be removed considering the negative influence on the prediction accuracy of real data. In this Chapter, we leverage the influence function to pick up the subset of pooled synthetic samples and then reassign the weight for each synthetic sample. We use the influence function (Cook and Weisberg, 1982) to evaluate the influence of deleting a synthetic sample on the prediction accuracy of the real data set. In the previous study, the influence function has been utilised to evaluate the effect of upweighting and perturbing a training input on the model parameters and the loss of test data (Koh and Liang, 2017). Moreover, conjugated gradients and stochastic estimation have been developed to reduce the computational burden for calculating the Hessian matrix in the influence function. In this Chapter, we compute the influence of dropping each synthetic sample on the prediction accuracy of the real sample set and then delete the top negative samples after ranking their impact.

Let $f_\theta(x)$ be our baseline DL prediction model, θ the parameter of the model, $L(f_\theta(x))$ the loss function between the true value and the predicted value. The influence of dropping the j^{th} synthetic sample $(x_j^{\text{s}}, y_j^{\text{s}})$ on the k^{th} real sample $(x_k^{\text{r}}, y_k^{\text{r}})$ is given as

$$I_{j,k} = L(f_\theta(x_k^{\text{r}})) - L(f_{\theta'}(x_k^{\text{r}})) \tag{3.6}$$

where θ is the optimal parameter that minimises the loss function $\theta = \arg\min_\theta \frac{1}{J}\sum_{i=1}^{J} L(f_\theta(x_i^{\text{s}}))$, and J is the total number of the synthetic samples after pre-processing. Denote θ' as the optimal DL model parameter by removing the synthetic sample $(x_j^{\text{s}}, y_j^{\text{s}})$, $\theta' = \arg\min_\theta \frac{1}{J}\sum_{i=1, i\neq j}^{J} L(f_\theta(x_i^{\text{s}}))$. However, one challenge lies in the computational efficiency to remove the data point and then retrain the model. Following the devel-

oped influence function (Koh and Liang, 2017), the effect of dropping the j^{th} synthetic sample $(x_j^{\text{s}}, y_j^{\text{s}})$ on the loss of predicting the k^{th} real sample $(x_k^{\text{r}}, y_k^{\text{r}})$ can be written as

$$I_{j,k} = - \bigtriangledown_\theta L(f_\theta(x_k^{\text{r}}))^{\text{T}} H_\theta^{-1} \bigtriangledown_\theta L(f_\theta(x_j^{\text{s}})) \tag{3.7}$$

where $H_\theta = \frac{1}{J} \sum_{i=1}^{J} \bigtriangledown_\theta^2 L(f_\theta(x_j^{\text{s}}))$ is Hessian matrix which is positive define. $\bigtriangledown_\theta L(f_\theta(x_k^{\text{s}}))$ represents how the synthetic sample influences the model parameters during training, and $\bigtriangledown_\theta L(f_\theta(x_k^{\text{r}}))$ captures how sensitive the real sample's loss is to changes in those parameters. The inverse Hessian H_θ^{-1} adjusts for the curvature of the loss surface, ensuring the direction and magnitude of influence are properly scaled.

In order to enhance the efficiency to calculate the influence function, Hessian-vector products is applied to avoid directly calculating H_θ^{-1}. Let $s_{\text{test}} = H_\theta^{-1} \bigtriangledown_\theta L(f_\theta(x_k^r))$ and the influence can be rewritten as $I_{j,k} = -s_{\text{test}} \cdot \bigtriangledown_\theta L(f_\theta(x_j^{\text{s}}))$. s_{test} can be approximating by

$$s_{\text{test}} \equiv \arg\min_p (p^{\text{T}} H_\theta p - \bigtriangledown_\theta L(f_\theta(x_k^{\text{r}}))^{\text{T}} p). \tag{3.8}$$

This is a convex quadratic problem where p is the solution vector that approximates $H_\theta^{-1} \bigtriangledown_\theta L(f_\theta(x_k^r))$. This problem can be efficiently solved using conjugate gradient methods, leveraging automatic differentiation frameworks that support Hessian-vector products without constructing the full Hessian (Martens et al., 2010).

The influence of removing one synthetic sample $(x_j^{\text{s}}, y_j^{\text{s}})$ on the whole real data set $\{(x_i^{\text{r}}, y_i^{\text{r}}), 1 \leq i \leq K\}$ can be summed as $I_{j,K} = \sum_{i=1}^{K} I_{j,k}$. We sort the influence values in ascending order, then delete the top negative synthetic samples that might reduce the prediction accuracy on real data. Besides, in order to choose a validation set to reweight the synthetic samples, we collect the top positive ones that are helpful in the prediction of real samples. After sorting the pooled synthetic samples according to the influence function, the synthetic samples $\{(x_i^{\text{s}}, y_i^{\text{s}}), 1 \leq i \leq J\}$ after dropout the top J' negative synthetic samples are defined as $\{(x_i^{\text{d}}, y_i^{\text{d}}), 1 \leq i \leq J - J'\}$.

Sample reweight

Traditionally, training a DL model aims to minimise the expected loss in the training samples, where all the samples are of the same importance. However, the synthetic data generated from different DA approaches have different influences on the performance of predicting performance. Thus, we need to reweight the synthetic samples to highlight the important samples before being used in the pretraining and transfer learning stage. The objective for a DL model with weighted loss in the synthetic samples $\{(x_i^{\mathrm{d}}, y_i^{\mathrm{d}}), 1 \leq i \leq J - J'\}$ after dropout can be written as

$$\theta^*(w) = \arg\min_{\theta} \sum_{i=1}^{J-J'} w_i L(f_\theta(x_i^{\mathrm{d}})) \tag{3.9}$$

where w_i is the weight of the i^{th} synthetic sample and $\sum_{i=1}^{J-J'} w_i = 1$. Besides, we select the top V positive influential samples by using the influence function as the validation set. Since these positive samples are most responsible for the prediction and can be used to guide the reweight process. The result of the optimal w_i can be obtained by minimising the loss on the validation data set $\{(x_i^{\mathrm{v}}, y_i^{\mathrm{v}}), 1 \leq i \leq V\}$. The optimal w^* is based on the performance of the validation set:

$$w^* = \arg\min_{w \leq 0} \frac{1}{V} \sum_{i=1}^{V} L(f_{\theta^*(w)}(x_i^{\mathrm{v}})) \tag{3.10}$$

As can be observed, the above is a complex and time-consuming two-loop optimisation process. A single loop optimisation was proposed in the mini-batch data $\{(x_i^{\mathrm{d}}, y_i^{\mathrm{d}}), 1 \leq i \leq n\}$ from the pooled synthetic samples after dropout and $\{(x_i^{\mathrm{v}}, y_i^{\mathrm{v}}), 1 \leq i \leq v\}$ from the validation data set based on single gradient descent step at step p, which can be utilised to update the parameter θ as (Alessandri and Gaggero, 2017):

$$\theta_{p+1} = \theta_p - \beta \bigtriangledown (\frac{1}{n} \sum_{i=1}^{n} (f_\theta(x_i^{\mathrm{d}}))) \tag{3.11}$$

where β is the step size, n is the mini-batch size of synthetic data and $n \ll J - J'$, v is the mini-batch size of validation data set and $v \ll V$. By perturbing the weight by for

each synthetic sample $L_{\delta_i}(f_\theta(x_i^{\mathrm{v}})) = \delta_i L(f_\theta(x_i^{\mathrm{v}}))$, then the sample weight $W_{i,p}^{'}$ at step p without normalisation can be estimated by single descent step in the mini-batch as (Ren et al., 2018):

$$W_{i,p}^{'} = -\epsilon \frac{\partial}{\delta_{i,t}} \frac{1}{v} \sum_{i=1}^{v} L(f_{\theta_{p+1}}(x_i^{\mathrm{v}}))\Big|_{\delta_{i,p=0}} \tag{3.12}$$

where ϵ is the descent step size. In the case where $W_{n,p}$ is negative, we can use $W_{i,p} = 0$ instead to diminish the negative effect on the weight of other samples. In order to normalise the weight set, we define $W_{i,p} = \frac{W_{i,p}^{'}}{\sum_i^m W_{i,p}^{'}}$. After optimising the weight for P steps, the final weight $W_{i,P}$ can be obtained. The detailed computation graph of reweight is given in Figure 3.4, which can be concluded as running the left graph firstly to train the LSTM model and then using backward automatic differentiation to make a second-order gradient. The pseudo-code of the sample reweight algorithm can be found in Algorithm 1.

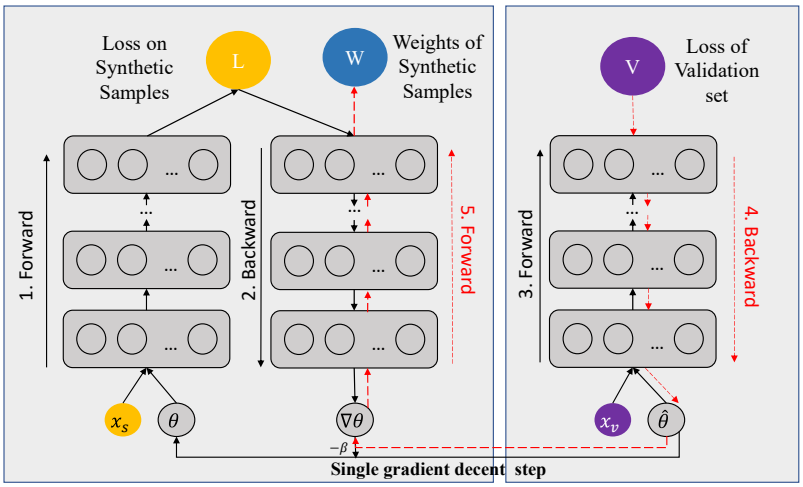

Figure 3.4: The single loop computation graph for the reweight process

3.2.4 Transfer learning scheme

By adapting the knowledge learned in a similar model, TL has been recognised as an effective approach to learn a DL model in the presence of insufficient data. Instance transfer, feature-based transfer, parameter transfer, and knowledge transfer are the

Algorithm 1: Data reweight algorithm

Initialise: Initialised DL model parameter θ_0,
mini batch size of pre-processed synthetic samples after dropout n,
mini batch size of pre-processed synthetic samples with top m positive influence,
pre-processed synthetic samples after dropout $\{(x_i^{\mathrm{d}}, y_i^{\mathrm{d}}), 1 \leq i \leq J - J'\}$,
pre-processed synthetic samples with high positive influence
$\{(x_i^{\mathrm{v}}, y_i^{\mathrm{v}}), 1 \leq i \leq V\}$

for $t = 0...T-1$ **do**

$\{(x_i^{\mathrm{d}}, y_i^{\mathrm{d}}), 1 \leq i \leq n\} \leftarrow \mathrm{GetMiniBatch}(\{(x_i^{\mathrm{d}}, y_i^{\mathrm{d}}), 1 \leq i \leq J - J'\})$

$\{(x_i^{\mathrm{v}}, y_i^{\mathrm{v}}), 1 \leq i \leq v\} \leftarrow \mathrm{GetMiniBatch}(\{(x_i^{\mathrm{v}}, y_i^{\mathrm{v}}), 1 \leq i \leq V\})$

end

for $i = 1...n$ **do**

$\hat{y}_i^{\mathrm{d}} \longleftarrow \mathrm{Forward}\ (x_i^{\mathrm{d}}, y_i^{\mathrm{d}}, \theta_t)$

end

$\epsilon \leftarrow 0$; $L^{\mathrm{d}} \leftarrow \sum_{i=1}^{n} \epsilon_i L(f_{\theta_t}(x_i^{\mathrm{d}}))$

$\nabla\theta_t \leftarrow \mathrm{BackwardAD}(L_{\mathrm{d}}, \theta_t)$

$\hat{\theta}_t \leftarrow \theta_t - \beta\nabla\theta_t$

for $i = 0...m$ **do**

$\hat{y}_i^{\mathrm{v}} \longleftarrow \mathrm{Forward}\ (x_i^{\mathrm{v}}, y_i^{\mathrm{v}}, \hat{\theta}_t)$

end

$L^{\mathrm{v}} \leftarrow \sum_{i=1}^{m} \epsilon_i L(f_{\hat{\theta}_t}(x_i^{\mathrm{v}}))$

$\nabla\epsilon \leftarrow \mathrm{BackwardAD}(L_{\mathrm{v}}, \epsilon)$

$W' = \max(-\nabla\epsilon, 0)$; $W = \frac{W'}{\sum_i^m W'}$

$\hat{L}^{\mathrm{d}} \leftarrow \sum_{i=1}^{i=n} W_i L(f_{\hat{\theta}_t}(x_i^{\mathrm{d}})$

$\nabla\theta_t \leftarrow \mathrm{BackwardAD}(\hat{L}^{\mathrm{d}}, \theta_t)$

$\theta_{t+1} \leftarrow \mathrm{OptimiserStep}(\theta_t, \nabla\theta_t)$

Output: W

most commonly used methods in TL (Tan et al., 2018). In our method, one of the key issues is to transfer the knowledge from the pre-trained neural network using the synthetic model to the real data. Thus, instead of pooling the real data and synthetic data directly, applying TL scheme is effective to borrow the knowledge from real data to the DL process trained with augmented data.

In the field of TL, a domain D generally consists of two elements: the feature space X and the corresponding probability distribution $P(x)$, where the input $x \in X$. The domain can be denoted as $D = \{(X, P(x))\}$. In a traditional TL process, the data sets are usually segmented into two domains: source domain and target domain. The two domains are generally with different probability distributions $D_{\text{source}} \neq D_{\text{target}}$. Each task T consists of two components: the label space y and the predictive model $f(\cdot)$. Thus the task can be represented as $T = \{(y, f(\cdot))\}$. Considering the two different domains, the objective of TL is to improve the model performance of target predictive function $f_{\text{target}}(\cdot)$ in D_{target} by utilising the knowledge in D_{source} and T_{source}.

Furthermore, recent evidence indicates that the transferability extracted from different layers can capture the latent features from general to specific when the layer goes deeper (Goodfellow et al., 2016). Four schemes have been proposed for RNNs in small data amounts to achieve fast transfer. In the proposed method, we develop a novel architecture for LSTM with 3 LSTM residual layers and 2 dense layers. Two stages are involved in the proposed TL scheme. Firstly, in the pre-train stage, we use synthetic samples after dropout to train the LSTM model with weight obtained in the sample reweight stage. Then, in the transfer stage, we freeze the trainable parameters in the top 3 LSTM residual layers, and then the dense layers are re-trained with a small learning rate. Besides, the synthetic samples after dropout are regarded as the source domain $D_{\text{source}} : \{(x_i^{\text{d}}, y_i^{\text{d}}), 1 \leq i \leq J - J'\}$ to pre-train the LSTM model, and the real samples are regarded as the target domain $T_{\text{source}} : \{(x_i^{\text{r}}, y_i^{\text{r}}), 1 \leq i \leq K\}$ that faces the practical prediction problem. The detailed transfer scheme is presented in Figure 3.5.

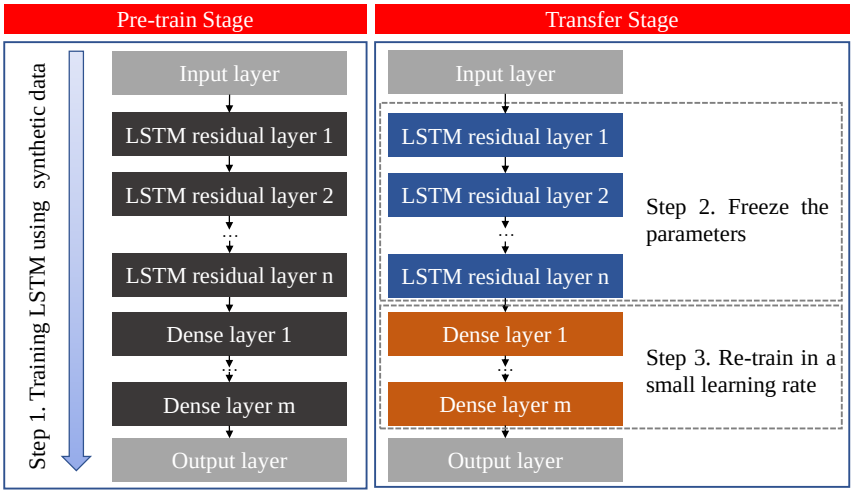

Figure 3.5: The detailed transfer scheme for LSTM structure

3.3 Numerical experiments

3.3.1 Simulation settings

The Wiener process has been widely applied in degradation modelling due to its mathematical property of independent increments. In this section, we leverage the Wiener process to simulate the degradation process, which is used to indicate the variation of health indicator $H(t)$. Assume that during the degradation process, the health indicator $H(t); t \geq 0$ is modelled as a Wiener process and formulated as:

$$H(t) = h_0 + \mu t + \sigma_B B(t)$$

where h_0 is the initial value of the health indicator, μ is the drift coefficient describing the rate of degradation, σ_B is the diffusion coefficient and $B(t)$ is the standard Brownian motion that represents the stochasticity of the degradation process. In order to analyse the performance of different prediction methods, we simulate the values of health indicator under different μ and σ_B, where $\mu \in \{0.1, 0.2, 0.3, 0.4\}$ and $\sigma \in \{1, 2, 3, 4\}$. In total, we generate 16 degradation curves in the time series form that follow the Wiener

process with different μ and σ_B. For each degradation curve, 100 time series data are randomly generated, as presented in Figure 3.6.

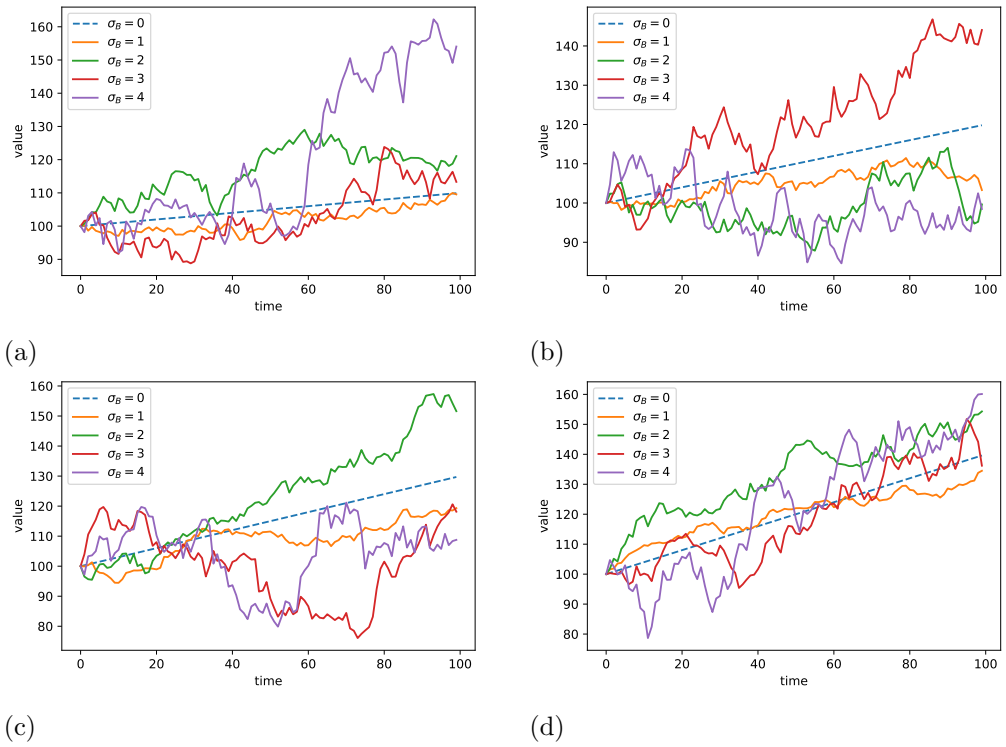

Figure 3.6: Illustration of simulated health indicator under different combinations of μ and σ_B. (a) Simulated data follow the Wiener process with $\mu = 0.1$, $\sigma_B \in \{1, 2, 3, 4\}$. (b) Simulated data follow the Wiener process with $\mu = 0.2$, $\sigma_B \in \{1, 2, 3, 4\}$. (c) Simulated data follow the Wiener process with $\mu = 0.3$, $\sigma_B \in \{1, 2, 3, 4\}$. (d) Simulated data follow the Wiener process with $\mu = 0.4$, $\sigma_B \in \{1, 2, 3, 4\}$.

In the data preprocessing stage, the moving window length l is set as 4, which means the previous 4-day observations are used to predict the value of the fifth day. After data preprocessing, 80% of the simulated samples are used as the training samples for data augmentation, reweight and transfer learning, and 20% of the simulated samples are regarded as the test samples to show the performance of the prediction methods.

Jittering, scaling, permutation, magnitude warping, window slicing, and window warping are used as data augmentation methods to extend the simulated data. Parameters of each augmentation method can refer to the existing research (Fawaz et al., 2018). The parameters of each method are given in Table 3.1.

During the data dropout and reweight process, after sorting the synthetic samples

Table 3.1: Parameters of the augmentation methods

Methods	Tunable Parameters
Jittering	$\sigma_{\xi} \in \{0.03, 0.06, 0.09\}$
Scaling	$\sigma_{\alpha_{\mathrm{s}}} \in \{0.05, 0.1, 0.15, 0.2\}$
Permutation	$N_{\mathrm{S}} \in \{5, 6, 7, 8, 9\}$
Magnitude warpping	$\sigma_M \in \{0.1, 0.2, 0.3\}$ $M \in \{4, 5\}$
Window slicing	$R_{\mathrm{WS}} \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$
Window Warping	$R_{\mathrm{WW}} \in \{0.05, 0.1.0.15, 0.2\}$

by the influence values, we remove the last 10% samples and select the top 10% samples to guide the data reweight. Given the LSTM's demonstrated ability to handle long-term dependencies in time series data, it is employed as the baseline model (Xia et al., 2020). Adaptive Moment Estimation is the optimisation method to minimise the square error loss during the training process in this work. Table 3.2 shows the structure of the baseline model, which consists of LSTM layers followed by Feedforward Neural Network (FNN) layers. We use the Rectified Linear Unit (ReLU) and Hyperbolic Tangent (Tanh) functions as activation functions in different layers.

Table 3.2: Structure of the baseline LSTM

Model	Lay No.	Layer type	Neurons	Activation
LSTM	1	LSTM	128	ReL)
	2	LSTM	128	ReLU
	3	LSTM	128	ReLU
	4	FNN	16	Tanh
	5	FNN	1	Tanh

3.3.2 Evaluation criteria

Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE) are two widely employed to measure the differences between the predicted value and observed value, and therefore are used to evaluate the prediction performance of the proposed approach. The RMSE is defined as follows:

$$RMSE = \sqrt{\frac{1}{N_{\mathrm{t}}} \sum_{i=1}^{N_{\mathrm{t}}} (y_i^{\mathrm{t}} - f_\theta(x_i^{\mathrm{t}}))^2} \tag{3.13}$$

where y_i^{t} represents the real observations in the test data set after pre-processing, $f_\theta(x_i^{\mathrm{t}})$ represents the output of the prediction model and N_{t} is the total number of the test samples. The definition of MAPE is given as

$$MAPE = \frac{100\%}{N_{\mathrm{t}}} \sum_{i=1}^{N_{\mathrm{t}}} \frac{|y_i^{\mathrm{t}} - f_\theta(x_i^{\mathrm{t}})|}{y_i^{\mathrm{t}}} \tag{3.14}$$

Clearly, the smaller the RMSE and MAPE, the better the performance of a given prediction method. In addition, each experiment is repeatedly verified 20 times to eliminate the effects of stochasticity and local minimal problems in DL and ML algorithms, and we calculate the mean values and standard deviation (Std) of RMSE and MAPE.

3.3.3 Comparisons

Six prediction algorithms including ARIMA, Random Forest Regressor (RFR), Gradient Boosting Machine (GBM), Light Gradient Boosting Machine (LGBM), eXtreme Gradient Boosting (XGB) and Least Absolute Shrinkage and Selection Operator (LASSO) are conducted to compare the performance of the DA methods and the transfer learning. ARIMA is a recognised statistical method to predict time series data and it will be used for comparison to illustrate the effectiveness of the DA-powered ML algorithms. RFR, GBM, LGBM, XGB and LASSO are typical ML algorithms utilised in both regression and classification tasks, and LSTM is the baseline DL model adopted in this Chapter. The parameters in these algorithms can be selected by the cross-validation method. The RMSE and MAPE of different prediction methods for the simulated health indicator data are shown in Table 3.3. Since each run of the ARIMA model returns the same performance, the Std values of RMSE and MAPE in ARIMA are not shown. Besides, the predicted value of LGBM and XGB remains the same in the result because of overfitting, the Std values of RMSE and MAPE are also not given.

The highlighted numbers are the smallest mean values of RMSE and MAPE. ARIMA shows the best performance on the prediction of the simulated data with $\mu = 0.1, \sigma_B = 1$ and LSTM_ADRT is very close to ARIMA regarding the mean of RMSE and MAPE. For other simulated data, LSTM_ADRT shows the smallest mean value of RMSE and

Table 3.3: RMSE and MAPE of different prediction methods for simulated data the under combinations of σ_B and μ

		$\sigma_B = 1$				$\sigma_B = 2$				$\sigma_B = 3$				$\sigma_B = 4$			
		RMSE		MAPE(%)		RMSE		MAPE(%)		RMSE		MAPE(%)		RMSE		MAPE(%)	
		Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
$\mu = 0.1$	ARIMA	**1.122**	/	**0.890**	/	1.562	/	**0.998**	/	3.554	/	2.567	/	5.418	/	2.747	/
	RFR	2.656	0.049	2.075	0.047	2.027	0.052	1.460	0.049	5.592	0.238	4.030	0.237	9.873	0.161	5.887	0.097
	GB	2.506	0.031	1.936	0.027	2.860	0.052	2.153	0.039	6.201	0.007	4.521	0.007	10.803	0.015	6.298	0.012
	LGBM	3.111	/	2.469	/	1.524	/	0.923	/	8.801	/	6.737	/	15.021	/	9.148	/
	XGB	2.519	/	1.947	/	2.489	/	1.862	/	6.6410	/	4.998	/	10.460	/	6.164	/
	LASSO	1.659	0.005	1.276	0.005	1.263	0.002	0.733	0.002	4.033	0.019	2.829	0.016	5.590	0.020	2.889	0.013
	LSTM	2.519	0.095	1.948	0.085	1.511	0.064	1.029	0.080	3.415	0.192	2.372	0.098	7.202	0.486	4.328	0.306
	LSTM_ADRT	1.135	0.035	0.893	0.038	**1.243**	0.030	1.004	0.072	**3.013**	0.021	**2.134**	0.020	**4.937**	0.222	**2.625**	0.195
$\mu = 0.2$	ARIMA	1.204	/	0.950	/	3.095	/	2.436	/	6.184	/	3.748	/	3.538	/	3.182	/
	RFR	1.165	0.029	0.953	0.029	3.613	0.052	2.784	0.041	12.507	0.360	8.493	0.260	2.951	0.088	2.361	0.065
	GB	1.245	0.016	0.973	0.013	4.384	0.025	3.522	0.025	15.907	0.107	10.956	0.077	3.294	0.050	2.360	0.044
	LGBM	1.713	/	1.397	/	5.256	/	4.047	/	14.973	/	10.282	/	3.393	/	3.149	/
	XGB	1.202	/	0.932	/	4.283	/	3.356	/	14.419	/	9.930	/	3.143	/	2.499	/
	LASSO	0.822	0.000	0.643	0.000	3.196	0.005	2.505	0.003	5.931	0.031	4.005	0.025	2.978	0	2.597	0
	LSTM	0.928	0.033	0.732	0.026	2.877	0.063	2.392	0.025	7.436	0.370	4.945	0.250	2.964	0.043	2.608	0.026
	LSTM_ADRT	**0.819**	0.010	**0.621**	0.020	**2.794**	0.040	**2.372**	0.022	**5.227**	0.204	**3.400**	0.146	**2.860**	0.046	**2.593**	0.020
$\mu = 0.3$	ARIMA	1.297	/	0.873	/	3.427	/	2.047	/	3.461	/	2.415	/	3.996	/	2.964	/
	RFR	4.985	0.041	3.874	0.037	15.139	0.132	8.377	0.085	4.454	0.112	3.194	0.079	3.364	0.130	2.604	0.091
	GB	4.962	0.033	3.844	0.029	14.602	0.042	8.012	0.030	5.138	0.074	3.695	0.056	4.278	0.062	3.111	0.039
	LGBM	5.769	/	4.588	/	17.420	/	10.131	/	4.068	/	3.266	/	4.303	/	2.945	/
	XGB	4.958	/	3.842	/	15.898	/	8.838	/	4.607	/	3.364	/	3.379	/	2.540	/
	LASSO	1.686	0.011	1.260	0.007	5.392	0.006	3.226	0.004	3.485	0.003	2.899	0.002	3.940	0.014	3.024	0.011
	LSTM	3.513	0.052	2.744	0.048	8.485	0.396	4.817	0.247	3.271	0.040	2.576	0.101	4.006	0.138	3.076	0.112
	LSTM_ADRT	**1.197**	0.079	**0.873**	0.055	**3.407**	0.227	**1.901**	0.133	**3.134**	0.023	**2.356**	0.033	**3.908**	0.054	**2.928**	0.097
$\mu = 0.4$	ARIMA	1.305	/	0.784	/	2.687	/	**1.580**	/	4.356	/	2.741	/	3.400	/	2.044	/
	RFR	1.690	0.026	1.041	0.022	7.044	0.093	4.312	0.064	6.432	0.053	3.282	0.038	7.026	0.154	3.293	0.090
	GB	2.322	0.012	1.556	0.008	5.510	0.032	3.251	0.027	5.821	0.007	2.934	0.014	8.184	0.082	4.258	0.055
	LGBM	3.773	/	2.424	/	7.479	/	4.548	/	10.761	/	6.370	/	8.013	/	4.039	/
	XGB	1.773	/	1.084	/	8.533	/	5.287	/	6.388	/	3.156	/	9.000	/	4.763	/
	LASSO	1.945	0.008	1.234	0.006	4.446	0.005	2.644	0.002	4.319	0.032	2.317	0.021	4.505	0.020	2.534	0.011
	LSTM	1.452	0.067	0.825	0.044	4.039	0.164	2.329	0.130	4.081	0.089	2.348	0.036	4.502	0.237	2.442	0.083
	LSTM_ADRT	**1.251**	0.034	**0.751**	0.008	**2.676**	0.088	1.590	0.052	**3.633**	0.076	**2.138**	0.026	**3.296**	0.061	**1.945**	0.045

MAPE. The comparative results indicate that the proposed approach is superior in most cases, especially for the degradation processes with high stochasticity. Traditional data-driven methods such as ARIMA, work well in the scenarios of degradation with less volatility.

3.3.4 Analysis of moving window length

In order to analyse the influence of different moving window length l on the prediction performance, we choose the simulated data generated by the Wiener process with $\mu = 0.4, \sigma_B = 4$ for the moving window length l equals to 4, 8 and 12. Table 3.4 presents the mean and Std of RMSE and MAPE of the 8 prediction methods on the test samples for different moving window lengths, wherein the highlighted values show the best performance for the corresponding settings. It can be observed that the mean of RMSE and MAPE shows an increasing trend with the increase of l in ARIMA and LSTM_ADRT. LASSO has the smallest Std among these prediction methods.

Table 3.4: RMSE and MAPE of different prediction methods for simulated data ($\mu = 0.4, \sigma = 4$) with moving window length $l = 4, 8, 12$

	$l = 4$				$l = 8$				$l = 12$			
	RMSE		MAPE(%)		RMSE		MAPE(%)		RMSE		MAPE(%)	
	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
ARIMA	3.400	\	2.044	\	3.589	\	2.068	\	3.822	\	2.103	\
RFR	7.026	0.154	3.293	0.090	6.815	0.235	3.123	0.119	6.621	0.234	2.978	0.087
GB	8.184	0.082	4.258	0.055	6.431	0.105	3.097	0.037	6.671	0.221	3.251	0.089
LGBM	8.013	\	4.039	\	7.795	\	3.897	\	6.545	\	3.079	\
XGB	9.000	\	4.763	\	9.233	\	4.512	\	8.900	\	4.167	\
LASSO	4.505	0.020	2.534	0.011	4.450	0.025	2.502	0.014	4.644	0.039	2.579	0.021
LSTM	4.502	0.237	2.442	0.083	4.490	0.196	2.425	0.041	4.803	0.073	2.460	0.041
LSTM_ADRT	**3.296**	0.061	**1.945**	0.045	**3.479**	0.035	**2.051**	0.020	**3.708**	0.063	**2.069**	0.030

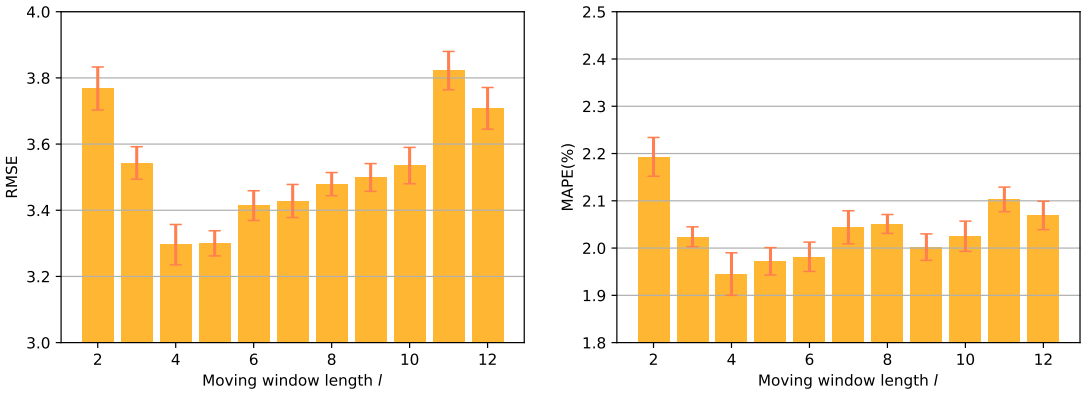

Figure 3.7: RMSE and MAPE of LSTM_ADRT for simulated data ($\mu = 0.4, \sigma = 4$) with moving window length l from 2 to 12

LSTM_ADRT has smaller Std values compared with the LSTM baseline model. Note that LSTM_ADRT still shows the best performance in the mean values of RMSE and MAPE compared with other prediction methods with different moving window lengths.

The detailed RMSE and MAPE of LSTM_ADRT prediction with simulated data generated by the Wiener process with $\mu = 0.4, \sigma_B = 4$ under different moving window lengths l from 2 to 12 are given in Figure 3.7. Both the RMSE and MAPE show an increasing first and then decreasing trend with the increasing of moving window length. When choosing moving window length $l = 4$, the LSTM_ADRT shows the best performance.

3.3.5 Discussion

The comparative results demonstrate the effectiveness of the proposed LSTM_ADRT method across a range of degradation processes with varying levels of stochasticity. When $\mu = 0.1$ and $\sigma_B = 1$, ARIMA shows the best performance in terms of RMSE and MAPE; however, the LSTM_ADRT model achieves very close results, indicating its capability to handle degradation scenarios with low volatility. For other combinations of μ and σ_B, particularly under higher stochasticity levels, LSTM_ADRT consistently outperforms other ML and DL baselines, achieving the lowest RMSE and MAPE values. These results highlight the robustness of the proposed method in capturing complex degradation patterns where traditional methods such as ARIMA struggle due to the increased uncertainty.

Further analysis on the influence of the moving window length, reveals that the window size has a significant impact on prediction performance. Both RMSE and MAPE exhibit a trend of initially increasing and then decreasing as the moving window length increases from 2 to 12. This behaviour suggests that an optimal window length exists, balancing between capturing sufficient temporal information and avoiding excessive noise. Specifically, a window length of $l = 4$ yields the best performance for LSTM_ADRT in the considered setting, achieving the lowest RMSE and MAPE compared to other methods. It is also observed that LSTM_ADRT maintains relatively low standard deviations, indicating better stability across different window lengths.

Overall, the results validate the effectiveness of the data augmentation, selection, and reweighting strategies in improving predictive performance. By managing the quality of augmented samples, the LSTM_ADRT model demonstrates generalisation and robustness across varying degradation dynamics and data segmentation settings.

3.4 Illustrative example

3.4.1 Description of the example and data

In this section, an illustrative example from a wastewater treatment plant with a modified activated sludge process is used to show the prediction performance of the proposed

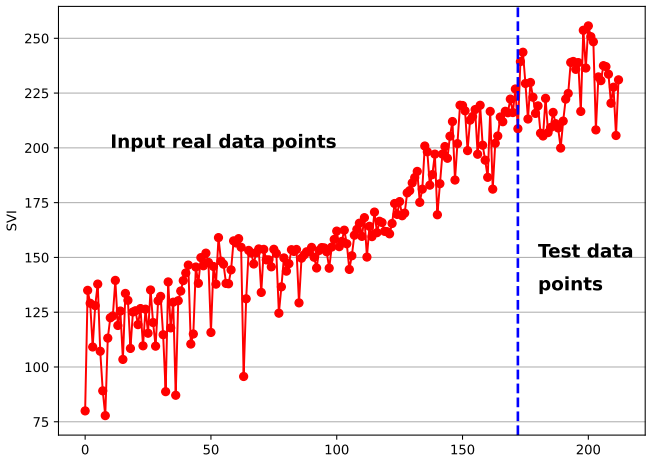

Figure 3.8: Real observations of SVI values

method. One of the most common failures in a wastewater treatment process is sludge bulking, which is caused by the imbalance of bacteria in the secondary clarifiers with a rapid degradation trend (Liu et al., 2020). SVI is the health indicator to imply the degradation degree and alarm the failure in advance. The values of SVI in this illustrative example are collected daily with 213 observations in total. The data collected from the operational process is given in Figure 3.8, which varies from 75mg/L to 253mg/L.

To implement the developed approach, the length of the moving window l is set as 4 days to make an illustration, which means that the SVI values of every four days are utilised to predict the value of the fifth day. In order to show the effectiveness of the proposed method, we split the real data set into two parts: 80% of the real data is used for the DA and training process, and 20% is used for testing the performance of the proposed method.

3.4.2 Experimental results and comparisons

Similar to the previous numerical experiments, six prediction algorithms including Arima, RFR, GBM, LGBM, XGB and LASSO are conducted to compare the performance of the DA methods and the transfer learning. Each algorithm runs repeatedly 20 times to show the performance. The mean and Std of RMSE and MAPE are calculated. The results are shown in Table 3.5. It can be observed that the proposed enhancement

framework LSTM_ADRT shows the best performance among these methods, in terms of the mean and Std values.

Table 3.5: Comparison of RMSE and MAPE among different prediction methods

Methods	RMSE		MAPE(%)	
	Mean	Std	Mean	Std
ARIMA	13.1248	/	4.6731	/
RFR	30.4910	0.5160	11.5360	0.2629
GB	19.7372	0.2340	6.64805	0.0753
LGBM	20.7906	/	6.9912	/
XGB	25.0841	/	8.7912	/
LASSO	13.1752	0.1938	4.6319	0.0934
LSTM	19.4894	0.1872	6.6509	0.0887
LSTM_ADRT	**12.4175**	0.0297	**4.4815**	0.0097

The estimations of SVI based on the ARIMA, RFR, GBM, LGBM, XGB, LASSO, LSTM and LSTM_ADRT are also illustrated in Figure 3.9. Considering sludge bulking as a soft failure in wastewater plants, defined by the SVI value exceeding 230 mg/L, we note that the first instance of this condition occurs in the 173rd set of real data. In this context, a soft failure implies that the system can still operate in the presence of the failure but with an elevated risk of disrupting the operating process. By anticipating the prediction value surpassing the threshold for soft failure, engineers can proactively take corrective actions. Then we take the 173rd data point as the real soft failure data point. Among these prediction methods, the proposed LSTM_ADRT predicts the 176th data point is the first time exceeding the soft failure threshold and the closest to the real soft failure data point.

3.4.3 Effectiveness of multi DA methods

We also draw line graphics to compare the performance for predicting SVI values with and without DA, as presented in Figure 3.10. It is obvious that the performance of these ML methods is improved with the augmented data. One extreme example is the LGBM, which fails to capture the dynamic behaviour of SVI. The flat line of LGBM shown in Figure 3.10c indicates the poor performance of LGBM when training only with the real data set. However, the performance of LGBM with augmented data

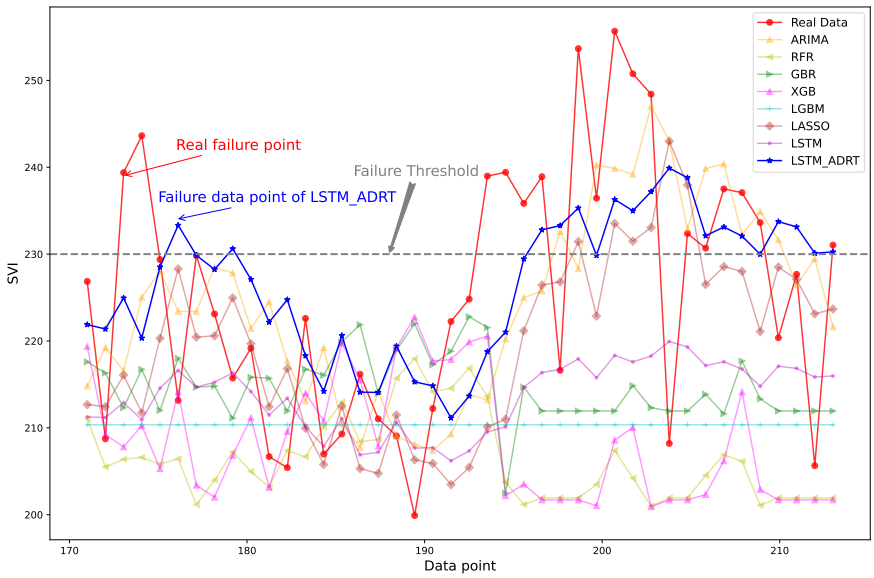

Figure 3.9: Estimated SVI value of test samples by different prediction methods

(LGBM_A) improves compared to without augmented data, and the errors between the real observations and predicted values are narrowed as shown in Figure 3.10c.

The details of RMSE and MAPE for each method with augmented data are shown in Table 3.6. Compared with Table 3.5, the six methods trained with augmented data outperform those with real data set. For example, the mean values of RMSE and MAPE based on RFR that trained with augmented data (RFR_A) are 13.6273 and 5.0433%, which is improved greatly from 30.4910 and 11.5360% when the RFR trained with limited real data set.

Table 3.6: RMSE and MAPE of different prediction methods with augmented data

Methods	RMSE		MAPE(%)	
	Mean	Std	Mean	Std
RFR_A	13.6273	0.0829	5.0433	**0.0257**
GBR_A	18.2938	**0.0521**	7.6743	0.0307
LGBM_A	18.4046	/	7.3856	/
XGB_A	19.4645	/	8.2831	/
LASSO_A	**12.9232**	0.1735	**4.5617**	0.0824
LSTM_A	13.0561	0.4023	4.7465	0.2038

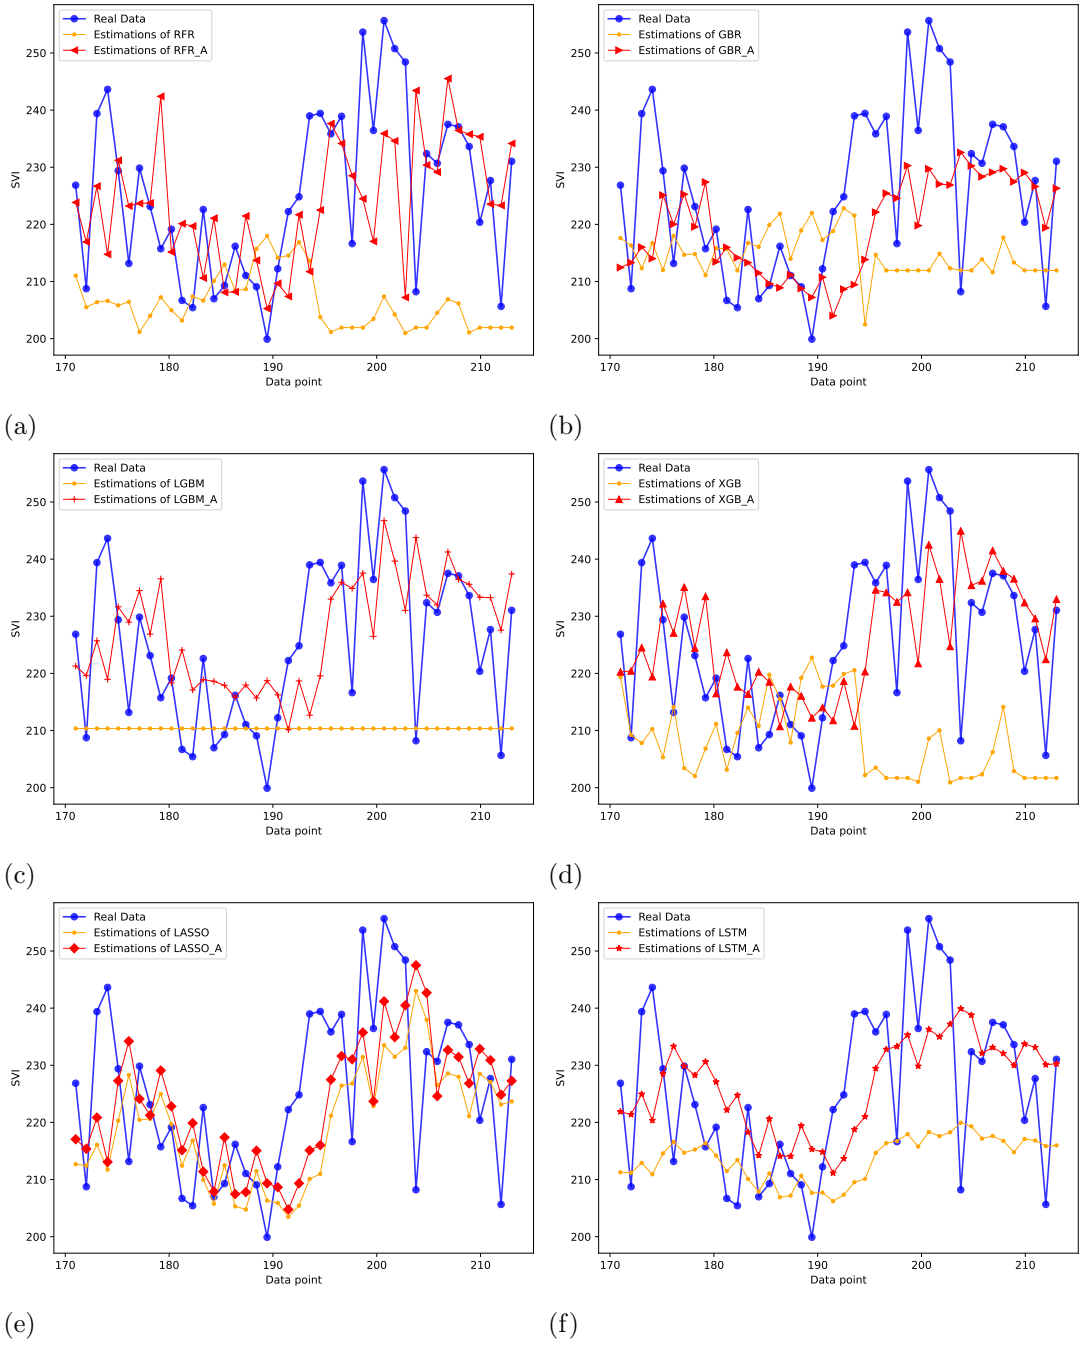

Figure 3.10: The effect of augmented data for the SVI prediction under different prediction methods

3.4.4 Performance of TL scheme

As discussed, the developed approach consists of four techniques: DA methods, sample dropping out, sample reweight and TL, to improve the performance of DL algorithms with limited observations. In order to illustrate the effectiveness of the TL scheme, we conduct TL separately on LSTM, LSTM with augmented data (LSTM_A), LSTM with augmented data and data dropout (LSTM_AD), LSTM with augmented data and reweight (LSTM_AR), and LSTM with augmented data, data dropout and reweight (LSTM_ADR). The mean values and Std of RMSE and MAPE for running 20 times are given in Table 3.7. To give a straightforward view of the effectiveness of TL, the RMSE and MAPE before and after TL are sketched in Figure 3.11. It can be observed that the TL scheme can effectively decrease the fluctuation and improve the stability of algorithms. Besides, LSTM_ADRT achieves the highest performance among these methods with respect to the mean and Std of MAPE. In terms of RMSE, LSTM_ADT has the smallest mean value and LSTM_ARRT has the smallest Std value, which indicates that both TL and DA methods can be used for DL algorithms with limited observations.

Table 3.7: The transfer effects on RMSE and MAPE

Methods	RMSE		MAPE(%)	
	Mean	Std	Mean	Std
LSTM	19.4894	0.1872	6.6509	0.0887
LSTM_A	13.0561	0.4023	4.7465	0.2148
LSTM_AD	12.5047	0.1669	4.5071	0.0312
LSTM_AR	12.4548	0.0919	4.5071	0.0690
LSTM_ADR	12.6871	0.0990	4.4870	0.0341
LSTM_AT	12.5736	0.1009	4.6072	0.0314
LSTM_ADT	**12.3983**	0.0397	4.5411	0.0313
LSTM_ART	12.4601	0.0346	4.4921	0.0161
LSTM_ADRT	12.4175	**0.0297**	**4.4815**	**0.0102**

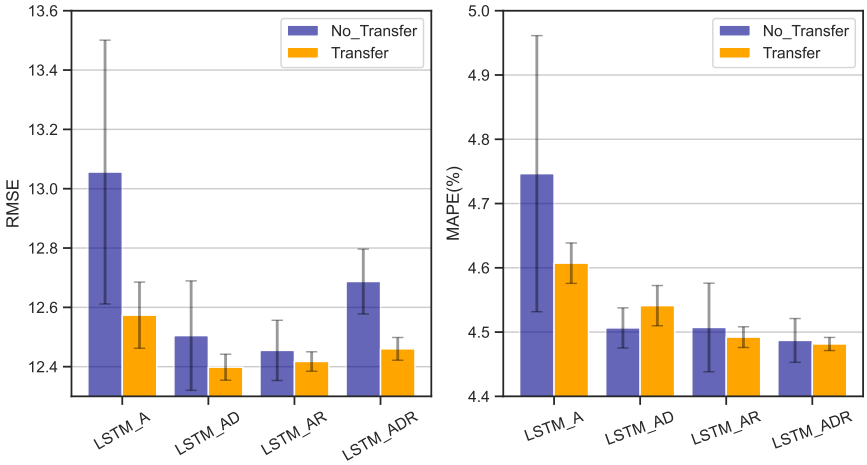

Figure 3.11: Comparing MAPE and RMSE of each method with and without transfer learning

3.4.5 Sensitivity analysis

3.4.6 Effects of the size of real samples

The proposed method aims to extend the ability of DL algorithms in the limited data scenario. In order to illustrate the effect of sample size, we adjust the size of the original data to the first 50 days and first 100 days. Similarly, 80% of the data is used to train the model and 20% is used to test the performance of the proposed method. Each method run 20 times to test the mean and Std values considering the randomness of the models.

Table 3.8 shows the mean value and Std of RMSE and MAPE under different methods with 50 real data and 100 real data that are collected in the first 50 days and 100 days. As for 50 real data, it can be observed that LSTM has poor performance with RMSE over 30 and MAPE over 20%. The DA techniques are effective in reducing the prediction error by nearly 40% in RMSE and MAPE. Besides, the proposed framework LSTM_ADRT achieves the best performance on RMSE value in the scenario of 50 samples.

Similarly, DA techniques can improve the performance of the LSTM model when the sample size is 100. The prediction ability of LSTM with the real data set of 100 days is better than the 50-day data set, which is 10.7751 on RMSE and 6.339% on

Table 3.8: RMSE and MAPE of different methods with 50 samples and 100 samples

Methods	Sample size = 50				Sample size = 100			
	RMSE		MAPE(%)		RMSE		MAPE(%)	
	Mean	Std	Mean	Std	Mean	Std	Mean	Std
LSTM	31.7497	7.7497	20.5531	4.4661	10.7751	5.1204	6.3399	2.2772
LSTM_A	17.3895	0.9755	11.5546	0.7544	8.1537	0.6045	4.8663	0.4454
LSTM_AD	16.0429	0.7075	11.3628	0.5066	7.5018	0.4930	4.3941	0.3461
LSTM_AR	14.4203	0.4104	9.7334	0.2593	6.3905	0.1410	**2.6178**	.0218
LSTM_ADR	14.3929	0.1274	8.6970	0.2396	6.2061	0.0695	2.6712	**0.0211**
LSTM_AT	15.3628	0.7830	**8.1850**	0.4287	7.0031	0.7131	4.0373	0.5087
LSTM_ADT	14.9845	0.8878	8.8311	0.4564	6.7965	0.3854	3.8616	0.2249
LSTM_ART	14.4390	0.1238	8.7174	0.2148	6.1570	0.1807	2.7219	0.1096
LSTM_ADRT	**14.3880**	**0.0731**	8.6386	**0.1761**	**5.9628**	**0.0248**	2.8707	0.0762

MAPE. The performance of LSTM is improved by about 20% with the enhancement of DA techniques. It indicates that the size of the data set can influence the performance. Besides, we can see that LSTM is improved gradually in the stacking process of data dropping, reweighting and TL scheme. In conclusion, LSTM_ADRT achieves the best performance on the mean values and their Std values of RMSE.

To illustrate the variations of the outputs for the 20 runs using 50 real samples, we plot in Figure 3.12a the RMSE and MAPE for each run with respect to various techniques: sample dropout, reweight, TL and their combinations. It can be observed that the dispersity, the mean value and Std of the outputs are reduced. This is due to the fact that the data reweight is continuously optimising the training data generated by multi-DA methods and reassigning the weight according to the rank of the influence for real data prediction accuracy.

The RMSE and MAPE for the repeated 20 experiments using 100 real samples are given in Figure 3.12b. It can be seen that reweight is the most effective method among these three methods when applying them separately. Besides, when stacking these methods, the method LSTM_ADRT shows the lowest RMSE value while LSTM_AR has the lowest MAPE value. One of the possible reasons is that the objective function for the LSTM when training is minimising the square error between the predicted values and observations.

Overall, the capability of LSTM is greatly influenced by the size of the training data

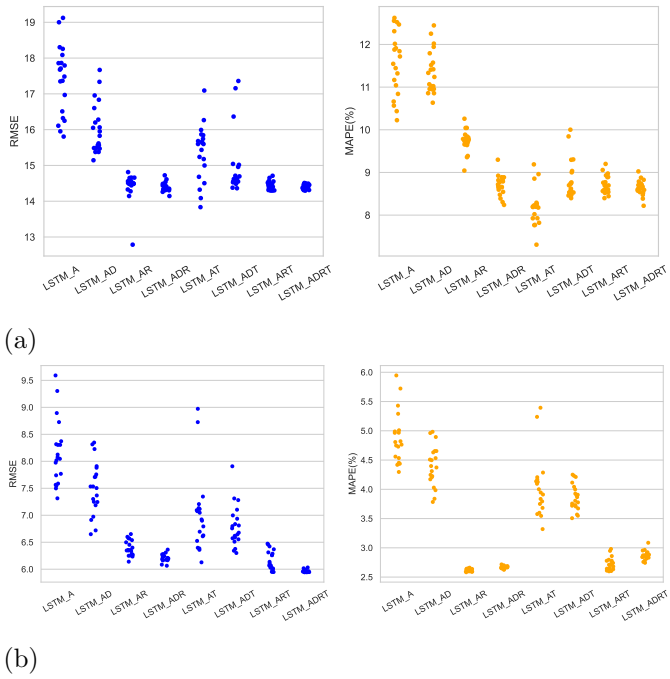

Figure 3.12: Outputs of RMSE and MAPE for each run (a) with 50 samples. (b) with 100 samples.

set, since the 100-day data set shows better performance on LSTM when training with 50-day data set and DA techniques can improve the original LSTM with large data size. Moreover, the stacking process shows improvements when applying each method and its combinations. Among these comparisons, LSTM_ADRT always shows the best performance in decreasing RMSE values.

Effects of Data Dropout

As previously discussed, the top negative influential samples are removed since they may lead to a decrease in the prediction accuracy of the health indicator. In this section, we aim to analyse the effects of the delete ratio on the performance of LSTM_ADRT. Moreover, in order to illustrate the effectiveness of dropout according to the rank of influence function, we also make a comparison with the method of deleting data randomly with a similar delete ratio in the proposed method.

Figure 3.13 gives the results of different delete ratios on the proposed method LSTM_ADRT from 0 to 50% with 5% as the step size in the bar plot. The result

of deleting data randomly is also shown in Figure 3.13 in the line graph. It can be concluded that random delete shows a worse performance compared to removing data according to the influence value. Besides, the RMSE and MAPE with different delete ratios show a trend of first falling and then rising, indicating that there is an optimal delete ratio. Unfortunately, we cannot provide a universal conclusion on the optimal delete ratio. In this illustrative example, the delete ratio of 20% shows the best performance.

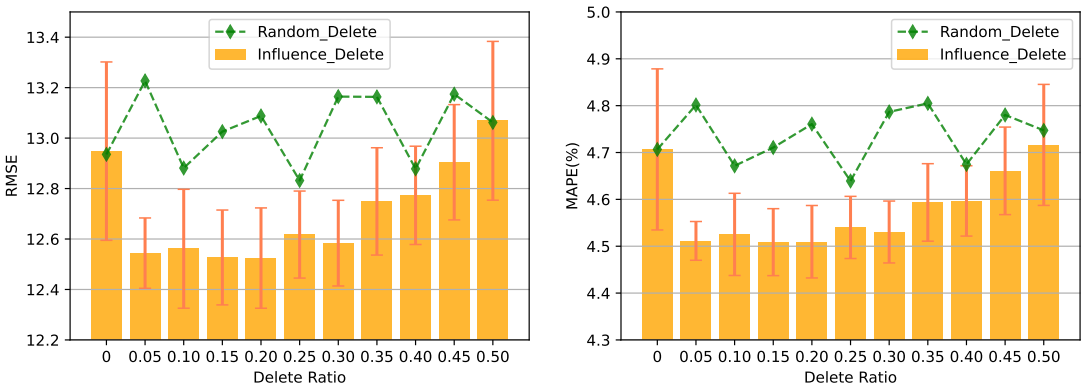

Figure 3.13: The effects on RMSE and MAPE of different sizes of delete samples and random delete samples

Effects of Data Reweight

As stated previously, the top influential samples generated by DA techniques are selected as the validation set to reweight the rest augmented samples. Thus, the size of the samples used as the validation set is one of the essential parameters to explore the effects on the results. In addition, we also explore the comparison between the top positive influential samples and randomly selected samples to illustrate the effectiveness of the proposed approach.

Figure 3.14 shows the variation of RMSE and MAPE with the number of reweight samples. The bar chart represents the results of reweight based on the influence function, while the dotted line stands for random reweight. It can be concluded that reweighting using the top influential samples instead of random selection leads to superiority in all the given sizes.

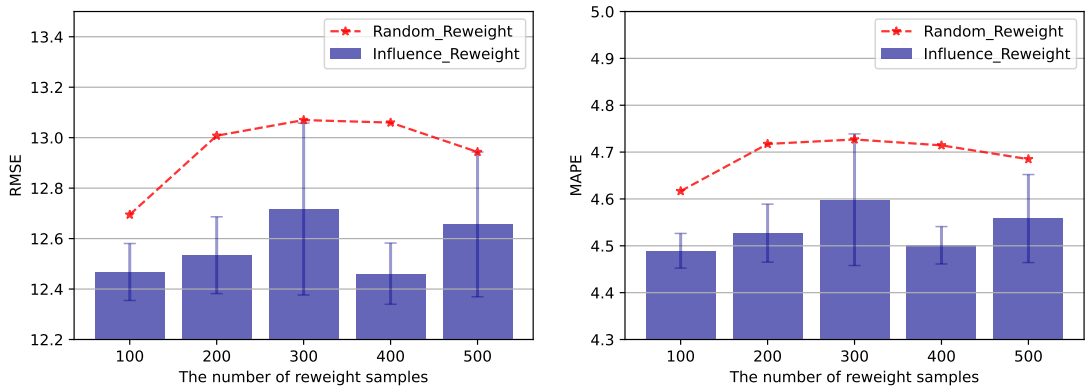

Figure 3.14: The effects on RMSE and MAPE of different size of influential samples and randomly selected samples

3.4.7 Discussion

An illustrative example from a real wastewater treatment plant was conducted in this section to illustrate the effectiveness of the proposed approach, which aimed to solve the overfitting of DL algorithms due to insufficient observations in the degradation prediction process. It shows that the proposed LSTM_ADRT outperforms other classic time series prediction methods in terms of RMSE, MAPE. Experiments also indicate the effectiveness of multi DA methods and the designed TL scheme. The classical machine learning methods applied in this Chapter are improved with augmented data. In addition, the designed TL scheme can effectively enhance the stability of the DL algorithms.

In addition, extensive experiments were conducted to test the effects of multiple elements, including the size of original samples, the delete ratio, and the validation sample size. When modifying the actual sample size, we maintain the parameters for augmentation methods unchanged, ensuring that the ratio of the augmented sample size to the real sample size remains constant. In terms of the size of real samples, more improvements by the proposed approach were observed with a smaller sample size. When adjusting the delete ratio from 5% to 50%, it shows the best performance when the delete ratio is 20%. Also, data dropout according to the rank of influence function shows superior performance compared with random dropout. Similar to data dropout, data reweight based on the rank of influence function illustrates its effective-

ness compared with that based on randomly selected data.

Furthermore, in comparison with simpler time series forecasting methods such as ARIMA, the proposed LSTM_ADRT model demonstrates slightly improved predictive performance, as evidenced by reductions in RMSE and MAPE metrics. When real-world data are limited, the benefits of such models remain uncertain and context-dependent. Although the improvements in this illustrative example are limited, the proposed approach may offer advantages in scenarios where the degradation process exhibits high variability.

3.5 Summary

This Chapter extends the applicability of DL algorithms in system degradation prediction with limited observations. DA is used to increase the training sample size and improve the performance of DL algorithms in degradation prediction. Multiple DA methods, including jittering, scaling, window warpping, etc., are used to generate synthetic data. Influence function is applied to drop out the synthetic samples that decrease the prediction accuracy. Considering the importance of different synthetic samples, gradient descent approach is used to reweight the selected synthetic samples before channelling into the DL algorithms. In addition, in order to further improve data efficiency, a TL scheme is designed to retrain the model with real samples. The numerical experiment examines the performance of the proposed method under various settings. The illustrative example conducted with a real data set from a wastewater treatment process shows the improved performance of DL algorithms and illustrates the effectiveness of the developed approach. The findings suggest that the proposed model excels in accurately predicting degradation, even with limited real observations, when compared to some state-of-the-art prediction methods. This is particularly evident in cases where the degradation process exhibits high stochasticity.

While this study has successfully improved the performance of deep learning algorithms in degradation forecasting with limited observations, more endeavours are demanded to further explore the potential of the developed approach. In this study, six augmentation methods were utilised; nevertheless, the impact of each augmentation

method on real-data prediction accuracy remains unknown. In future research, there is a need to explore the development of an instructive indicator to guide the selection of a specific augmentation method. In the context of application, decision-makers encounter challenges in efficiently utilising degradation prediction for informed decision-making, such as scheduling predictive maintenance and allocating maintenance resources. Furthermore, in the extension, the integration of health indicator prediction through the developed approach could be explored as a means to guide the formulation of an optimal maintenance policy.

Chapter 4

A Data-Driven End-to-End Preventive Maintenance Policy with Varied Working Conditions

We investigate a data-driven preventive maintenance problem in which the underlying failure distribution is unknown, but the past run-to-failure data and working conditions can be observed. The traditional Estimate-Then-Optimise (ETO) framework decouples the estimation process from optimisation to identify the optimal maintenance decisions. However, this approach may lead to the propagation of estimation errors into the optimisation phase, affecting the overall efficacy of the decision-making process. To address this issue, we propose an End-To-End (E2E) framework to directly suggest the optimal preventive replacement time under specific working conditions, without relying on any preconceived assumptions of the failure distribution. In the E2E framework, the historical working condition contexts are treated as featured data and mapped with the approximated optimal maintenance decisions obtained by minimising the data-driven maintenance cost rate functions. These functions have been analytically validated to closely mirror the actual maintenance cost rate function, ensuring reliability and accuracy. Subsequently, supervised learning algorithms are leveraged to train the structured data that consists of the features and corresponding approximated optimal maintenance decisions. This approach enables us to streamline the decision-making process, facil-

itating rapid and automated maintenance directives for systems under both new and previously observed conditions. Numerical experiments illustrate the superiority of the proposed E2E framework over the traditional ETO approach, in the context of both repairable and non-repairable systems.

4.1 Introduction

Effective maintenance decision-making is crucial in mitigating unplanned failures and reducing economic losses across diverse sectors, including manufacturing, national infrastructure, and the automotive industry (Elwany and Gebraeel, 2008). A system's ageing process is influenced by its working conditions, including both the environmental and operating conditions (Martorell et al., 1999). The advent of the Internet of Things has enabled the accumulation of extensive data on failure events, under varying working conditions, highlighting the need for adaptive preventive maintenance policies (Compare et al., 2019).

Consider the example of vehicle tyres, which are typically recommended for replacement after about six years, irrespective of their visible conditions (Kalsher et al., 2005). The research shows that tyres are more prone to failure under conditions like high-speed operation or extreme temperatures, underscoring the importance of developing maintenance strategies that are responsive to specific operational scenarios. Modern "smart" tyres, equipped with sensors, offer valuable insights into conditions like temperature, pressure, and speed (Garcia-Pozuelo et al., 2017). These data, combined with historical time-to-failure records, are stored in cloud databases, presenting a complex challenge: efficiently deriving maintenance decisions from vast datasets and intricate patterns.

Maintenance policies are broadly classified into two groups: TBM and CBM (Ahmad and Kamaruddin, 2012). Both policies aim to minimise the objective of interest, such as minimisation of maintenance cost or maximisation of availability, to decide the appropriate maintenance plan. The difference lies in their foundational data: TBM depends on the relevant time-to-failure records, while CBM is guided by the degradation behaviour and the condition monitoring data (de Jonge et al., 2017). Despite the potential advantages of CBM, TBM continues to be used due to its straightforward

implementation using existing time-to-failure data (de Jonge et al., 2015). The primary goal of TBM is to identify the optimal time for Preventive Replacement (PR), guided by objectives such as expected maintenance costs per unit time, and system availability.

This study focuses on the time-based preventive maintenance policy. The traditional ETO procedure of preventive maintenance policy consists of two steps (Elmachtoub et al., 2023): parameter estimation upon the relevant time-to-failure model, and the optimisation procedure to determine the optimal PR time. During the estimation step, typically it is assumed that the time to failure follows a distribution, such as the Weibull or Gamma distribution (Ahmad and Kamaruddin, 2012). The contexts of working conditions are modelled as covariates to describe the influence upon the ageing process. For non-repairable systems, Cox PH regression models the relationship between the covariates and the hazard rate (Thijssens and Verhagen, 2020). For repairable systems, the number of recurrent failure events is modelled with a non-homogeneous Poisson process. Cox Proportional Intensity (PI) function models the intensity of the recurrent failures (Jiang et al., 2006). The estimated parameters of these models are then factored into the optimisation model to determine the optimal maintenance policy.

However, this two-step method suffers from several shortcomings. First, it decouples the estimation and optimisation stages, meaning that predictive models typically compress high-dimensional input features—such as contextual variables and uncertainty indicators—into a single point estimate or simplified distribution. In doing so, high-dimensional information is discarded during the optimisation stage (Qi et al., 2023). Second, the objective function used in estimation (e.g., minimising prediction error) is misaligned with the objective function used in optimisation (e.g., minimising cost, risk, or regret). This misalignment means that even highly accurate forecasts may not translate into good decisions (Donti et al., 2017). Third, in some cases, the underlying data-generating process of time-to-failure is unknown or only partially observable. If the estimation is inaccurate, especially due to model misspecification or sparse data, then decisions derived from the estimated model may systematically deviate from the true optimum. If the model error is small and the decision is not highly sensitive, the performance loss may be negligible. However, in high-stakes or highly non-linear

decision contexts, small model misspecifications might propagate into flawed decisions (Elmachtoub and Grigas, 2022).

In this Chapter, we develop a data-driven E2E preventive maintenance policy that handles features and does not require the knowledge of the underlying time-to-failure model. Here, the term "features" is synonymous with covariates used to represent the working conditions. Working condition features may include environmental variables (e.g., temperature, humidity), operational load levels (e.g., pressure, speed, vibration), or usage patterns (e.g., operating hours, duty cycles). These covariates are treated as the input to the supervised learning model. We derive approximated maintenance cost functions for both repairable and non-repairable systems to obtain data-driven maintenance decisions for varied working conditions. For each working condition, we utilise supervised learning algorithms to map the features of working conditions with maintenance decisions. The E2E framework is a one-step method that aims to find the near-optimal maintenance time for preventive maintenance policy based on the input features. While this work represents the first exploration of one-step method in the maintenance decision domain, the integration of estimation and optimisation has been examined in feature-based inventory management and the newsvendor problem (Qi et al., 2023; Ban and Rudin, 2019; Oroojlooyjadid et al., 2020).

Labelling the optimal PR time poses a challenge when applying the supervised learning algorithm, particularly in a data-driving setting where the true time-to-failure model is unknown. We consider the scenario where historical time-to-failure data and features can be collected. This work applies two supervised learning algorithms: LASSO and Artificial Neural Network (ANN). LASSO is well-suited for tackling high-dimensional linear problems, and ANN is suitable for handling non-linear and high-dimensional data. The contribution of our work in relation to the existing literature can be summarised in three aspects:

- We are the first to propose approximated maintenance cost rate functions for repairable systems and non-repairable systems, enabling us to decide near-optimal PR time.

- We theoretically establish the bound between the approximated maintenance cost

rate and the true maintenance cost rate function, providing a theoretical guarantee of the accuracy of our approach.

- We develop a data-driven E2E framework for preventive maintenance policy. The optimal maintenance decisions are provided directly from the data. This framework neither specifies the form nor estimates the parameters of the time-to-failure model, thereby avoiding potential error propagation.

The remainder of this Chapter is organised as follows. Section 4.2 describes the preventive maintenance problem and the conventional two-step ETO procedures. Section 4.3 presents an overview of the E2E framework for repairable and non-repairable systems, followed by a discussion of preventive maintenance policy under time-varying working conditions. Section 4.4 presents numerical experiments to illustrate the effectiveness of the proposed E2E framework. Finally, the conclusion and future research directions are provided in Section 4.5.

4.2 Problem description

In this Chapter, we focus on developing a preventive maintenance policy that takes into account the varied working conditions for single-unit systems. A single-unit system is defined as a system containing only one primary component or unit, such as a light bulb or a pump in a water supply system. The reliability of the entire system depends solely on this single unit. The examination of single-unit systems serves as a foundation for understanding more complex configurations, including series, parallel, and redundant systems (Wang, 2002). We will use the term "systems" to refer to the single-unit systems in the following.

The main objective is to determine the optimal PR time while minimising the maintenance cost rate, given a specific working condition. We use the term "features" to refer to the covariates that describe the working conditions. We assume the relevant features have been identified and appropriate historical time-to-failure data has been collected. Furthermore, we assume that the types of features are stable and the values are constant.

The preventive maintenance policy differs based on whether the system is repairable or non-repairable. Repairable systems are subject to repair intended to restore functionality and so can experience recurrent failures. On the other hand, non-repairable systems experience a single failure event and will have to be replaced upon failure. In the following sections, we introduce the preventive maintenance policy for repairable systems and non-repairable systems, respectively.

4.2.1 Preventive maintenance policy for repairable systems

For repairable systems, failures are recurrent events, and as a result, multiple failure times are observed for each system. Denote the arrival time of failures during a decision interval under working condition i as $\mathbf{z}_i = [z_{ijk}, ..., z_{ijk}]$, where $1 \leq i \leq M, 1 \leq j \leq J_i, 1 \leq k \leq K_{ij}$, M is the number of working conditions, J_i is the number of systems operated under working condition i, and K_{ij} is the total number of failures of the jth system under working condition i. The accumulative time-to-failure data of repairable systems with working conditions is denoted as $X_M = [(\mathbf{x}_1, \mathbf{z}_1), ..., (\mathbf{x}_M, \mathbf{z}_M)]$, where $\mathbf{x}_i$ is the feature vector of working condition i, $\mathbf{x}_i \in \mathbb{R}^p$, and p indicates the number of features.

We consider two types of maintenance actions for repairable systems: PR and minimal repair. If the system fails during the decision-making interval, a minimal repair can restore the system to an "as-bad-as-old" state. This indicates that the system will be restored to the operational state immediately before failure, and the failure intensity will remain identical. By contrast, PR can make the system "as-good-as-new", and the failure intensity reverts to the initial level when the system was installed. The downtime duration of both minimal repairs and PR is assumed to be negligible in relation to the operational time. The optimisation objective is to minimise the cost rate per unit of time within a finite period. However, the maintenance time is affected by other operating constraints such as resource availability, and therefore we assume that the maintenance decision is within the interval $\left[T_{min}^{RS}, T_{max}^{RS}\right]$, $T_{min}^{RS} > 0$.

The maintenance cost rate function is given by:

$$C_{RS}(t|\mathbf{x}) = \frac{c_m \mathbb{E}\left[N(t|\mathbf{x})\right] + c_p}{t}, \quad t \in \left[T_{min}^{RS}, T_{max}^{RS}\right], \tag{4.1}$$

where $C_{RS}(t|\mathbf{x})$ is the maintenance cost rate at time t, $\mathbb{E}\left[N(t|\mathbf{x})\right]$ denotes the expected number of minimal repairs based on the vector of features $\mathbf{x}$, c_m and c_p represent the costs associated with each minimal repair and PR, $c_m < c_p < \infty$. When the failure intensity $m(t|\mathbf{x})$ is known, the expected number of minimal repairs is given by:

$$\mathbb{E}[N(t|\mathbf{x})] = \int_0^t m(t|\mathbf{x})dt.$$

The optimal time for PR under working condition $\mathbf{x}$ is:

$$TM^*(\mathbf{x}) = \arg \min_{T_{min}^{RS} \leq t \leq T_{max}^{RS}} C_{RS}(t|\mathbf{x}). \tag{4.2}$$

4.2.2 Preventive maintenance policy for non-repairable systems

In terms of non-repairable systems, we consider a scenario where there are Q_i systems observed under the same working condition i. The time to failure of systems observed under working condition i can be represented as a vector $\mathbf{t}_i = [t_{i1}, ..., t_{iQ_i}]$. Additionally, for each set of time-to-failure observations $\mathbf{t}_i$, the corresponding q features are known. The time-to-failure data set with working conditions can be represented as $S_N = [(\mathbf{s}_1, \mathbf{t}_1), ..., (\mathbf{s}_N, \mathbf{t}_N)]$, where $\mathbf{s}_i$ is the feature vector and $\mathbf{s}_i \in \mathbb{R}^q$ for $i \in \{1, ..., N\}$, N is the number of working conditions, and q denotes the number of features.

Since failure is a single event for a non-repairable system, the system can either replace the failed system upon failure or at the optimal PR time. We assume both PR and failure replacement recover the system to an "as-good-as-new" state with different costs. Within the decision interval $\left[T_{min}^{NRS}, T_{max}^{NRS}\right]$, the preventive maintenance model proposed to decide the optimal PR time while minimising the maintenance cost rate function $C_{NRS}(t|\mathbf{s})$ is expressed as follows:

$$C_{NRS}(t|\mathbf{s}) = \frac{c_f F(t|\mathbf{s}) + c_r R(t|\mathbf{s})}{\int_0^t R(t|\mathbf{s})dt}, \quad t \in \left[T_{min}^{NRS}, T_{max}^{NRS}\right], \tag{4.3}$$

where $C_{NRS}(t|\mathbf{s})$ is the maintenance cost rate at time t with the given feature vector $\mathbf{s}$ for non-repairable systems. $F(t|\mathbf{s})$ is the probability of failure at time t, $R(t, \mathbf{s})$ is the reliability at time t, and c_r is the cost of preventive replacement and c_f is the cost of failure replacement, $c_r < c_f < \infty$. When the hazard rate $\lambda(t|\mathbf{s})$ is known, the reliability $R(t|\mathbf{s}) = \exp\left[-\int_{x=t}^{t} \lambda(x|\mathbf{s})dx\right]$, thereby the cumulative probability of failure $F(t|\mathbf{s}) = 1 - \exp\left[-\int_{x=t}^{t} \lambda(x|\mathbf{s})dx\right]$. The optimal PR time within the decision interval $[T_{min}^{NRS}, T_{max}^{NRS}]$ under working condition $\mathbf{s}$ can be obtained as:

$$TR^*(\mathbf{s}) = \arg \min_{T_{min}^{NRS} \leq t \leq T_{max}^{NRS}} C_{NRS}(t|\mathbf{s}).$$

4.2.3 ETO method

The traditional ETO framework for obtaining a preventive maintenance policy requires statistical analysis of time-to-failure data. In the optimisation problem defined in Equation (4.1), accurately estimating the expected number of failures is the key task before taking into optimising maintenance decisions for repairable systems. In Equation (4.3), the parameters of the failure Cumulative Distribution Function (CDF) need to be estimated in order to determine the optimal PR time.

For repairable systems, the number of failure occurrences is modelled as a stochastic point process with a non-stationary failure intensity function $m(t|\mathbf{x})$, which captures the deterioration over discrete periods (Usher et al., 1998). We assume that the failure intensity can be modelled with a parametric function as $m_{\mathbf{w}}(t|\mathbf{x})$ given the feature vector $\mathbf{x}$ and parametric vector $\mathbf{w}$. The power-law and log-linear model are two functions to model failure intensity. The MLE technique can be used to estimate the parameter $\mathbf{w}$ of the failure intensity function for repairable systems. Given the failure data set X_M, the likelihood function for the parametric model is given by:

$$L_{RS}(\mathbf{w}) = \prod_{i=1}^{M} \left\{ \prod_{j=1}^{J_i} \left\{ \prod_{k=1}^{K_{ij}} m(z_{ijk}|\mathbf{x}_i, \mathbf{w}) \right\} \exp\left(- \int_0^{T_{max}^{RS}} m(t|\mathbf{x}_i, \mathbf{w})dt \right) \right\},$$

which is the product of individual likelihoods for each repairable system. The estimated parameter $\mathbf{w}$ can be obtained by setting the first derivative $\frac{\partial \log L_{RS}(\mathbf{w})}{\partial \mathbf{w}} = 0$.

For non-repairable systems, the estimation of the failure probability can be achieved by estimating the hazard rate $\lambda(t|\mathbf{s}, \boldsymbol{\Theta})$, where $\boldsymbol{\Theta}$ denotes the vector of parameters in the hazard rate function. In situations where only time-to-failure data is available, a Weibull function can be used to model the hazard rate due to its ability to represent various types of ageing systems, including those with increasing, decreasing, or constant hazard rates. Similar to repairable systems, MLE can be utilised to estimate the true distributional parameters. Given the time-to-failure data set S_N, the likelihood function is:

$$L_{NRS}(\boldsymbol{\Theta}) = \prod_{i=1}^{N} \left\{ \prod_{j=1}^{Q_i} \lambda(t_{ij}|\mathbf{s}_i, \boldsymbol{\Theta}) \exp\left\{-\lambda(t_{ij}|\mathbf{s}_i, \boldsymbol{\Theta})\right\} \right\}.$$

The estimation of the parameters $\boldsymbol{\Theta}$ can be obtained by equating the partial derivatives of the logarithm of the likelihood function with respect to the parameters to zero.

4.3 End to end preventive maintenance policy

4.3.1 Overview of the E2E framework

The E2E framework is designed to directly derive maintenance decisions based on a historical data set. This approach bypasses the need to estimate the parameters of the hazard rate or failure intensity model in the preventive maintenance policy. The first step involves determining whether the systems are repairable or non-repairable, as the maintenance policy influences the formulation of objective functions for maintenance decisions. After that, during the second stage, we aim to obtain an approximated maintenance cost function based on historical time-to-failure data collected under the same working conditions. Next, a supervised algorithm is trained using the data set generated in the second stage composed of paired features and maintenance decisions. This training process enables the algorithm to learn patterns and relationships between the input features and the corresponding maintenance decisions. Finally, the maintenance decisions of new working conditions can be obtained by the trained model. The overall framework is illustrated in Figure 4.1.

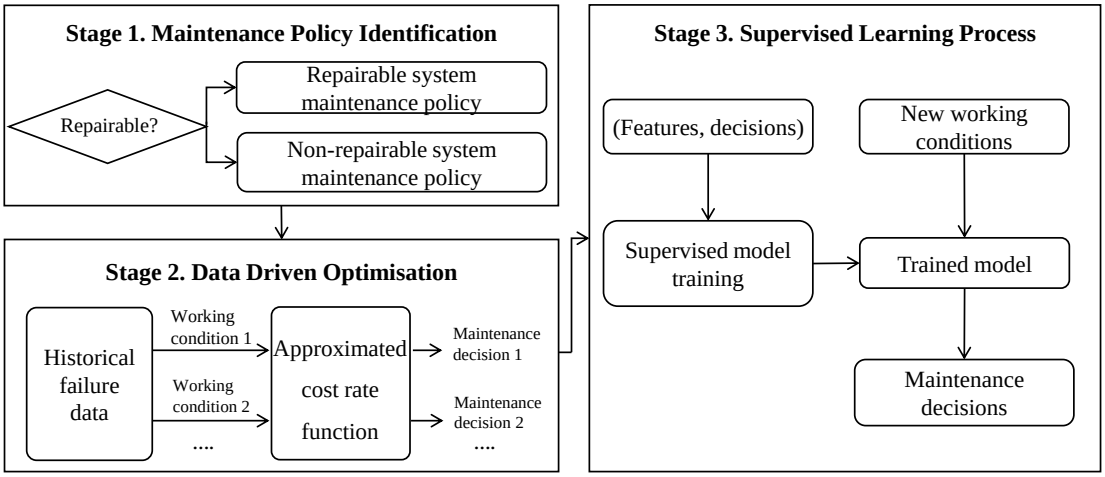

Figure 4.1: Overview of the E2E maintenance decision making framework

4.3.2 Data-driven optimisation of maintenance decisions

Maintenance decisions for repairable system

To obtain approximate PR maintenance decisions with the collected time-to-failure data, we construct a data-driven maintenance cost rate function for repairable systems. Without estimating the parameters for the failure intensity model under varied working conditions, we rely on the sample average to estimate the expected number of failures by a given time for a repairable system under a specific working condition.

For repairable systems with the accumulated time-to-failure data set X_M, since we have J_i systems operated under working condition i, we can empirically estimate the average number of failures at time t under working condition i. Taking working condition i as an example, the number of failures for system j from the beginning to the current time t is $\hat{N}_{ij}(t;\mathbf{x}_i) := \sum_{k=1}^{K_{ij}} \mathbf{1}(z_{ijk} \leq t)$, in which $\mathbf{1}(A)$ is the indicator function that equals to 1 when A is true. Subsequently, the approximated maintenance cost rate function $\hat{C}_{RS}(t;\mathbf{x}_i)$ under working condition i with feature vector $\mathbf{x}_i$ can be obtained as

$$\hat{C}_{RS}(t;\mathbf{x}_i) := \frac{c_m \frac{1}{J_i} \sum_{j=1}^{J_i} \hat{N}_{ij}(t;\mathbf{x}_i) + c_p}{t}, t \in \left[T_{min}^{RS}, T_{max}^{RS}\right]. \tag{4.4}$$

Here we use average $\frac{1}{J_i}\sum_{j=1}^{J_i}\hat{N}_{ij}(t;\mathbf{x}_i)$ to approximate the expected number of failures at time t. The PR time derived from Equation (4.4) is defined as data-driven PR time, which is given by

$$T\hat{M}_i^* := \arg\min_{t\in[T_{min}^{RS},T_{max}^{RS}]} \hat{C}_{RS}(t;\mathbf{x}_i).$$

Next, we obtain a structural data set that consists of features and data-driven PR time for repairable systems. The structural data set is denoted as $D_{RS} := \Big[(\mathbf{x}_i, T\hat{M}_i^*), 1 \le i \le M\Big]$.

Define $N_{ij}(t)$ is the expected number of failures of the repairable system j under working condition i. Assume that $N_{i1}(t), N_{i2}(t), ..., N_{iJ_i}(t)$ are independent and identically distributed, and $N_{min} \le N_{ij}(t) \le N_{max}$. In the following, we present the absolute bound between the approximated maintenance cost rate function and the true maintenance cost rate function by leveraging Hoeffding's inequality.

Proposition 4.1 *Given the decision interval that $t \in [T_{min}^{RS}, T_{max}^{RS}]$, for any accuracy parameter $\omega > 0$, with a probability of at least $1 - 2\exp\{\frac{-2J_i\omega^2}{(N_{max}-N_{min})^2}\}$, the absolute error between approximated maintenance cost rate and true maintenance cost rate is*

$$\Big|C_{RS}(t;\mathbf{x}_i) - \hat{C}_{RS}(t;\mathbf{x}_i)\Big| \le \frac{c_m\omega}{t}.$$

The proof of Proposition 4.1 is presented in the Appendix A.1. Proposition 4.1 suggests that the accuracy of the cost rate approximation improves as more systems are operated under the same working condition. Specifically, increasing the number of systems operated under the same working condition decreases the probability of encountering large errors. This relationship is reflected in the exponential term of the probability expression. This proposition has implications for the design and operation of maintenance strategies. In particular, it suggests that obtaining a large sample of systems operating under the same working condition can enhance the reliability of cost rate approximations. Furthermore, when only a limited number of systems are available, the approximated cost rates should be used with caution due to potentially higher estimation errors.

To illustrate the data-driven maintenance decision-making method, we present a simple numerical example considering two different working conditions, each with two

identical repairable systems. Specifically, we have system 1 and system 2 operating under working condition 1, and system 3 and system 4 operated under working condition 2. The cost of minimal repair is set as $c_m = 200$ and the cost of PR is set as $c_p = 800$. The failures of these systems occur at specific times, as depicted in Figure 4.2a. After employing $\frac{1}{J_i}\sum_{j=1}^{J_i}\hat{N}_{ij}(t;\mathbf{x}_i)$, we can obtain a piecewise function, as shown in Figure 4.2b, representing the empirical expected number of failures for each working condition. Incorporating this empirical expected failure count into the maintenance cost rate function yields the approximated maintenance cost rate, depicted in Figure 4.2c. It is worth noting that the approximated maintenance cost rate remains in a piece-wise form, with the change points corresponding to the failure times. The piecewise function is non-increasing between two consecutive failure times. By obtaining the minimum value from the approximated maintenance cost rate function, the optimal PR time for a specific working condition can be determined, as illustrated in Figure 4.2d. It is worth mentioning that this example involves only four systems operating under two working conditions to illustrate the data-driven method. However, it should be emphasised that the effectiveness of the proposed data-driven approach will improve with the availability of the number of systems operated under the same working condition, as its performance has been guaranteed by the absolute error between the approximated maintenance cost rate function and the true one.

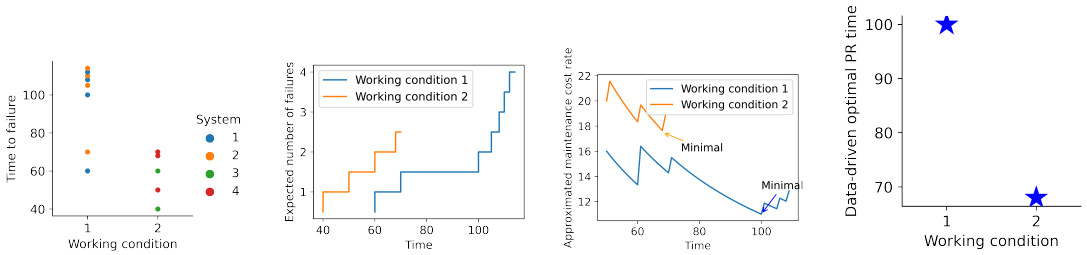

(a) Time to failure of 4 systems

(b) Expected number of failures

(c) Approximated maintenance cost

(d) Data-driven maintenance decison

Figure 4.2: Illustration of data-driven maintenance decision for repairable systems

Maintenance decisions for non-repairable systems

When considering non-repairable systems, the objective function differs from the repairable systems and the key task lies in estimating the reliability function or failure CDF. Leveraging the historical time-to-failure data set S_N with Q_i run-to-failure identical non-repairable systems operated under working condition i, we can empirically estimate the failure CDF and the reliability function as

$$\hat{F}(t;\mathbf{s}_i) := \frac{1}{Q_i}\sum_{j=1}^{Q_i}\mathbf{1}(t_{ij} \leq t), \tag{4.5}$$

$$\hat{R}(t;\mathbf{s}_i) := 1 - \hat{F}(t;\mathbf{s}_i), \tag{4.6}$$

where $\hat{F}(t;\mathbf{s}_i)$ and $\hat{R}(t;\mathbf{s}_i)$ are the empirical failure CDF and empirical reliability function with feature $\mathbf{s}_i$ respectively. Then, the approximated maintenance cost rate under feature $\mathbf{s}_i$ can be obtained:

$$\hat{C}_{NRS}(t;\mathbf{s}_i) := \frac{c_f\hat{F}(t;\mathbf{s}_i) + c_r\hat{R}(t;\mathbf{s}_i)}{\int_{t=0}^{t}\hat{R}(t;\mathbf{s}_i)}, t \in \left[T_{min}^{NRS}, T_{max}^{NRS}\right]. \tag{4.7}$$

By minimising the approximated maintenance cost rate function, the data-driven PR time for working condition i with feature $\mathbf{s}_i$ can be obtained as

$$\hat{TR}_i^* := \arg\min_{t\in\left[T_{min}^{NRS}, T_{max}^{NRS}\right]} \hat{C}_{NRS}(t;\mathbf{s}_i).$$

By constructing the data-driven maintenance cost rate functions in this manner, we can approximate the optimal PR time for the corresponding working condition and obtain the structural data set consisting of features and data-driven PR time for non-repairable systems, denoted as $D_{NRS} := [(\mathbf{s}_1, \hat{TR}_i), ..., \mathbf{s}_N, \hat{TR}_N)]$. Based on Massart's inequality, we can derive the bound for the approximated maintenance cost rate function in the following.

Proposition 4.2 *For any* $t \in [T_{min}^{NRS}, T_{max}^{NRS}]$, *with the probability of at least* $1 - 2\exp^{-2Q_i\varepsilon^2}$, *where* $0 \leq \varepsilon < \hat{R}^2(T_{min}^{NRS};\mathbf{s}_i)$ *is the accuracy parameter, the absolute error between the approximated maintenance cost rate function and true maintenance cost*

function for non-repairable systems is:

$$\left|\hat{C}_{NRS}(t;\mathbf{s}_i) - C_{NRS}(t;\mathbf{s}_i)\right| \leq \frac{(2c_f - c_r)\varepsilon t}{\left[\hat{R}^2(T^{NRS}_{min};\mathbf{s}_i) - \varepsilon\right](T^{NRS}_{min})^2}. \tag{4.8}$$

The proof of Proposition 4.2 is given in Appendix A.2. Based on Proposition 4.2, it is evident that reducing the accuracy parameter ε leads to a tighter bound. However, this tighter bound is associated with a reduced probability of guarantee. Furthermore, a higher probability of guaranteeing the bound is observed when a greater number of systems operate under identical working conditions while maintaining a consistent bound between the approximated function and the true function. This proposition highlights considerations for maintenance decision-making in non-repairable systems. It emphasises the trade-off between approximation accuracy and confidence levels: achieving tighter cost rate approximations requires more operational data. In scenarios where only limited data is available, decision-makers must carefully balance the desire for precision against the risk of relying on less confident estimations.

In a similar vein to the repairable systems, we can provide a simple example to further illustrate the data-driven maintenance decision method. Consider that 5 systems are operated under working condition 1 and 4 systems operating under working condition 2. The time to failure of these systems is presented in Figure 4.3a. By incorporating the empirical probability of failure, as depicted in Figure 4.3b, we can derive the approximated maintenance cost rate function, which is shown in Figure 4.3c. Finally, by minimising the approximated cost rate function, we obtain the maintenance decision for non-repairable systems under the two varied working conditions, as shown in Figure 4.3d.

Supervised learning for mapping features and decisions

With the collected feature data and failure events, our goal is to determine the optimal maintenance time through a supervised learning method, that is $f_{RS} : \mathbf{x}_i \rightarrow TM^*_i$ for repairable systems and $f_{NRS} : \mathbf{s}_i \rightarrow TR^*_i$. Taking the linear supervised learning algorithm as an example, the mapping between the working conditions and data-driven

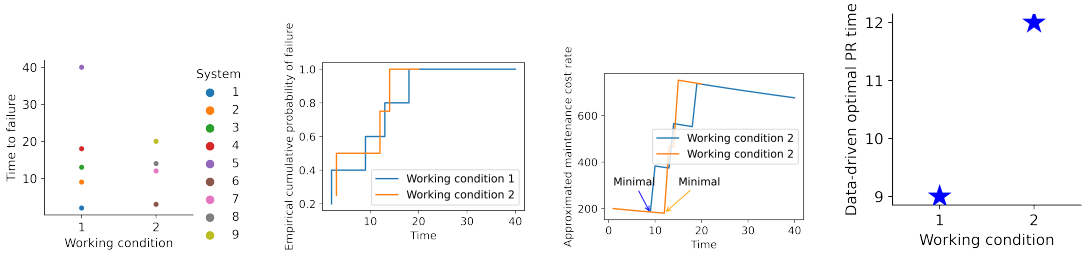

(a) Time to failure of 9 systems

(b) Empirical failure probability

(c) Approximated maintenance cost

(d) Data-driven maintenance decision

Figure 4.3: Illustration of data-drive maintenance decision for non-repairable systems

PR time is illustrated in Figure 4.4. The training objective is to make the predicted PR time close to the true optimal value. We define the respective training objective functions for repairable systems and non-repairable systems as

$$\min \sum_{i=1}^{M} \mathcal{L}(f_{RS}(\mathbf{x}_i), T\hat{M}_i^*)$$
$$\min \sum_{i=1}^{N} \mathcal{L}(f_{NRS}(\mathbf{s}_i), T\hat{R}_i^*),$$

where $\mathcal{L}(\cdot)$ is the loss function to calculate the distance between predictions and labels. Various supervised learning models can be employed to establish the mapping between features and maintenance decisions. In this work, we illustrate the approach using two representative methods: LASSO, which captures linear and sparse relationships, and ANN, which model complex and non-linear mappings (Hoff et al., 2009; Mishra and Srivastava, 2014). These examples are not intended to suggest a strict preference, but rather to demonstrate how different types of models can be integrated into the policy learning framework. Other supervised learning models, such as decision trees, random forests, or support vector machines, may also be appropriate depending on the data characteristics and application requirements (Muhammad and Yan, 2015).

4.3.3 Discussion: extension to time-varying working conditions

When dealing with time-varying working conditions, our data-driven preventive maintenance decision method remains applicable, provided that multiple identical systems

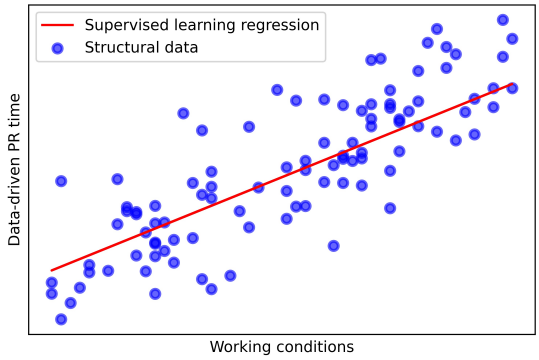

Figure 4.4: Illustration of mapping working conditions with data-driven PR time

are observed concurrently. These systems can be measured either from the same starting time or from each system's base time, as long as they operate under the same working condition after beginning operation. However, the form of features will differ as the values are not fixed over time. To account for time-varying features, we can introduce an additional dimension to represent the temporal aspect by utilising a piecewise method to divide time into multiple slots. This approach, which is also used in traditional ETO preventive maintenance policies, allows us to capture the changing working conditions and incorporate them into the decision-making process (Liao and Tian, 2013). By considering the temporal dimension, we can effectively adapt the maintenance decisions to varying working conditions over time.

Consider the repairable systems and assume that the decision-making duration is divided into m time slots. For a specific working condition i, the features $\mathbf{x}_i$ is an $p \times m$ matrix represented as: $\mathbf{x}_i = (x_{ijk})_{1 \leq j \leq p, 1 \leq k \leq m}$. When facing high-dimensional features, where the number of features p and the number of time slots m are large, the dimensionality reduction method can be applied to reduce the feature space while preserving the essential information. For example, principal component analysis can effectively reduce the dimensionality of the features while retaining the most important patterns and relationships within the data. By applying the dimensionality reduction methods, we can decrease the size of the feature set while still capturing the relevant information necessary for making accurate maintenance decisions. This dimensionality reduction can simplify the modelling problem and alleviate the potential computational burden

associated with high-dimensional feature spaces. In addition, deep learning algorithms, including convolutional neural networks, are well-suited for handling high-dimensional input features, such as two-dimensional matrices or three-dimensional tensors.

4.4 Numerical experiments

This Section presents experiments designed to evaluate the proposed E2E maintenance decision making framework. The purpose of these experiments is to validate the framework's ability to accurately generate maintenance decisions under varying working conditions. Specifically, the experiments aim to explore the following questions: (1) How accurately are the cost rate functions approximated? (Section 4.4.2) (2) How accurately can the framework output maintenance decisions compared with conventional ETO methods? (Section 4.4.3) (3) How does the performance change when the underlying parameters are varied? (Section 4.4.4)

4.4.1 Experiments description

Experiment settings

Consider a set of identical repairable systems observed over the same time window, and the failure intensity is in the form of the power-law function with three features, as given by

$$m(t|\mathbf{x}) = \frac{a}{b}(\frac{t}{b})^{a-1} \exp(b_1 x_1 + b_2 x_2 + b_3 x_3),$$

where a is the shape parameter and b is the scale parameter. The cost of minimal repair is $c_m = 1000$ and the cost of PR is $c_p = 2000$. The decision duration is set to the range [1,200]. The values of features are randomly generated for each working condition, where x_1 and x_2 are randomly sampled from uniform distributions, $x_1 \sim U[0,1], x_2 \sim U[-1,1]$. The feature x_3 is a binary variable generated by a uniform distribution, and $x_3 \sim U\{0,1\}$. Monte Carlo simulation is used to randomly generate accumulated time to failure under 200 different working conditions. This is achieved by generating time intervals between successive failures in a non-homogeneous Poisson process, using the power-law failure intensity function with parameters listed in Table

4.1. For each working condition, 50 systems are operated.

Table 4.1: Experiment settings for the repairable systems

Parameters	c_p	c_m	a	b	b_1	b_2	b_3	x_1	x_2	x_3	J	M
Values	2000	1000	2.5	50	0.2	0.1	0.7	$U[0,2]$	$U[-1,1]$	$U\{0,1\}$	50	200

For non-repairable systems, the cost of PR is $c_r = 3000$, the cost of failure replacement $c_f = 4000$, and the decision interval is set to the range [10,100]. The number of features is also set to 3. The true non-repairable system is assumed to follow the Weibull PH model with a Weibull hazard rate, which is given as:

$$\lambda(t|\mathbf{s}) = \frac{\alpha}{\beta}(\frac{t}{\beta})^{\alpha-1}\exp(c_1 s_1 + c_2 s_2 + c_3 s_3)$$

where α is the shape parameter, β is the scale parameter and c_1, c_2, c_3 are coefficients for three features. The values of features s_1, s_2, s_3 are randomly generated using a uniform distribution within the interval $[0, 2]$. Monte Carlo simulation is used to randomly generate the time-to-failure data for each system based on the time-to-failure distribution derived from the Weibull PH model. The parameters used in the experiments for the set of non-repairable systems are presented in Table 4.2. For each of the 400 working conditions, 50 identical systems are operated.

Table 4.2: Experiment settings for the non-repairable systems

Parameters	c_r	c_f	α	β	c_1	c_2	c_3	s_1,s_2,s_3	Q	N
Values	3000	4000	3	60	0.4	0.1	0.2	$U[0,2]$	50	400

After labelling the time-to-failure data, we obtain the structural data D_{RS} and D_{NRS} for repairable and non-repairable systems respectively. A pair consisting of a working condition and its PR time in the structural data is referred to as a sample. 70% of the samples are randomly selected to train the supervised learning models, i.e. LASSO and ANN. The remaining 30% is used to evaluate the performance of the model. Min-max normalisation is applied to scale the training and test data set into the range [0,1]. As for the ANN, a three-hidden-layer architecture is employed, with the number of neurons in each layer set as 32, 16, and 1. The corresponding activation functions

for each layer are Relu, Relu and sigmoid (Muhammad and Yan, 2015). Techniques such as grid search or random search with cross-validation can systematically tune hyperparameters of LASSO and ANN. The process during the supervised learning stage is the same for both repairable and non-repairable systems.

Performance metrics

To evaluate the performance of the proposed E2E framework, we use the Mean Absolute Error (MAE) for maintenance time prediction, and for maintenance cost rate. MAE in maintenance time prediction quantifies how closely the predicted maintenance times align with the ground truth or optimal policy. This reflects the accuracy of the model in identifying when maintenance should occur. MAE in maintenance cost rate captures the average maintenance cost incurred per unit time. This metric reflects the overall economic effectiveness of the maintenance policy.

For repairable systems, the performance metrics are given as:

$$MAE_{RS}^{T} = \frac{1}{N_1}\sum_{i=1}^{N_1}\left|\tilde{TM}_i - TM_i^*\right|$$
$$MAE_{RS}^{C} = \frac{1}{N_1}\sum_{i=1}^{N_1}\left|C_{RS}(\tilde{TM}_i) - C(TM_i^*)\right|$$

where MAE_{RS}^{T} is the MAE between the predicted PR time given by E2E or ETO methods and true optimal PR time, MAE_{RS}^{C} is the MAE of cost rate based on predicted PR time and cost rate based on true optimal PR time, and N_1 is the number of evaluated samples.

For non-repairable systems, the MAE of PR time and corresponding maintenance cost rate are:

$$MAE_{NRS}^{T} = \frac{1}{N_2}\sum_{i=1}^{N_2}\left|\tilde{TR}_i - TR_i^*\right|$$
$$MAE_{NRS}^{C} = \frac{1}{N_2}\sum_{i=1}^{N_2}\left|C_{NRS}(\tilde{TR}_i) - C(TR_i^*)\right|$$

where MAE^{T}_{NRS} is the MAE between predicted PR time and optimal PR time, MAE^{C}_{NRS} is the MAE of maintenance cost rate based on predicted PR time and cost rate based on optimal PR time, and N_2 is the number of evaluated samples.

In addition, in-sample and out-of-sample performances are reported to assess the performance of the proposed E2E optimisation framework. The in-sample performance refers to evaluating the train data used to fit the supervising models, while the out-of-sample performance measures the model's ability to generate accurate results on unseen data. Lower values of MAE correspond to more accurate performance. To mitigate the effects of randomness, we conduct each experiment 20 times and then report their average performance.

4.4.2 Construction of approximated cost rate functions and maintenance decision

To construct the approximated cost rate functions for repairable systems and non-repairable systems respectively, we use the empirical expectation to approximate the expected number of failures within the given duration, and the empirical CDF to approximate the true failure CDF of the assumed model. For repairable systems, we compare the empirical expected number of failures with the true expected failure number when the number of systems $J = 10$ and $J = 50$, as shown in Figure 4.5a. Additionally, we compare the corresponding approximated maintenance cost rate and true maintenance cost rate in Figure 4.5b. Based on these results, it is observed that the empirical expected number of failures is a better approximation when more systems are operated under the same working condition. Moreover, the approximated maintenance cost rate closely aligns with the true maintenance cost rate when $J = 50$ compared to $J = 10$, which is consistent with Proposition 4.1.

Similarly, when considering the number of systems running under the same condition as $Q = 10$ and $Q = 50$, we present the empirical failure CDF and the true CDF for the Weibull PH model using the parameters provided in Table 4.2. The functions are depicted in Figure 4.6a. It can be found that the empirical CDF for $Q = 50$ closely approximates the true CDF compared to the case of $Q = 10$. After utilising the em-

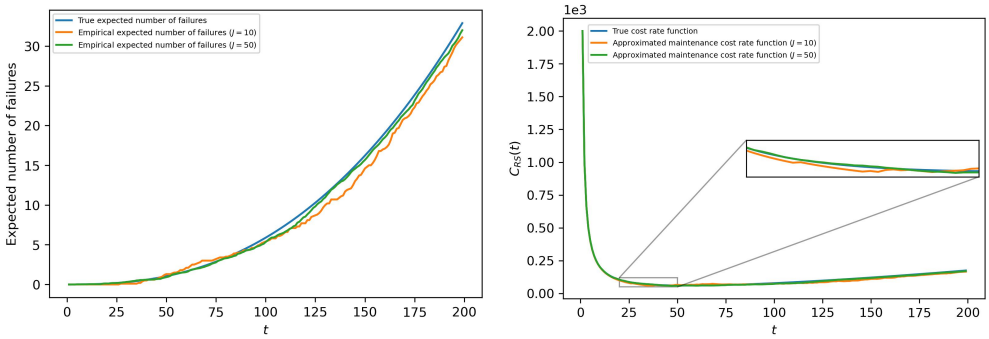

(a) Number of expected failures

(b) Approximated maintenance cost rate function

Figure 4.5: Approximation of the number of expected failures and data-driven cost rate function when $J = 10$ and $J = 50$ for repairable systems

pirical failure CDF to construct the maintenance cost rate function for non-repairable systems, we show the approximation of the data-driven maintenance cost rate function to the true maintenance cost rate function in Figure 4.6b. It is evident that there is a noticeable discrepancy between the approximated cost rate function for $Q = 10$ and the true cost rate function. In contrast, the approximation of maintenance cost rate function for $Q = 50$ exhibits a closer resemblance to the true cost function, which is consistent with the theoretical result established in Proposition 4.2.

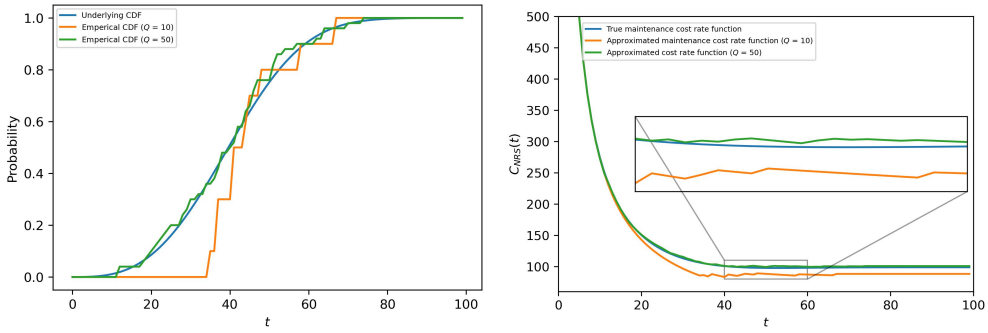

(a) Approximation of failure CDF

(b) Data-driven maintenance cost rate

Figure 4.6: Approximation of the number of expected failures and data-driven cost rate function when $J = 10$ and $J = 50$ for repairable systems

The data-driven maintenance decisions for repairable and non-repairable systems are obtained by minimising their corresponding approximated maintenance cost rate

functions. To assess the approximation of the data-driven method, we create problem setups with the number of systems that operated under the same working condition $J \in \{50, 100, 150, ..., 500\}$ and $Q \in \{50, 100, 150, ..., 500\}$. While the time-to-failure data is randomly generated from the true failure CDF, we keep the other parameters in Table 4.1 unchanged. Figure 4.7a presents the differences between the labels and optimal PR time for repairable systems across 20 repeated experiments. Here, the differences for the i^{th} randomly-generated non-repairable failure data under the working condition j are calculated as $DIFF_{ij}^{RS} = \tilde{TM}_{ij}^{*} - TM_{ij}^{*}$.

We repeat the experiment 20 times for each $J \in \{50, 100, 150, ..., 500\}$ and $Q \in \{50, 100, 150, ..., 500\}$, and evaluate the differences between the data-driven PR time and the optimal PR time. The differences are shown in the boxplot in Figure 4.7b. Figure 4.7a and Figure 4.7b suggest that as the number of systems increases, the spread of the differences between the data-driven PR time and optimal PR time decreases, due to the improved approximation of the maintenance cost function. Furthermore, the median of the differences is close to 0 regardless of the number of systems, which indicates the robustness of the E2E framework.

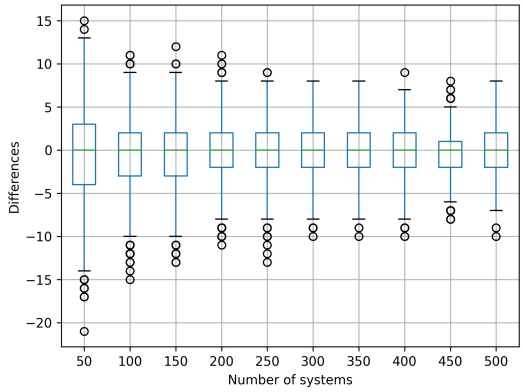

(a) Differences in repairable systems when varying J

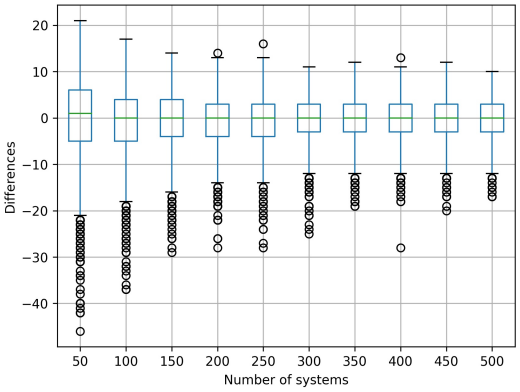

(b) Differences in non-repairable systems when varying Q

Figure 4.7: Differences of data-driven PR time and optimal PR time when varying the number of systems under the same working condition

To analyse the results provided by LASSO and ANN models, we compare the PR time predicted by LASSO, and ANN with the optimal PR time and the data-driven

PR time. For repairable systems, we vary the value of $x_1 \in \{0, 0.0.01, 0.02, ..., 1\}$, while keeping x_2 and x_3 fixed at 1 and 0. Figure 4.8a shows that the PR time suggested by LASSO decreases with an increase in x_1. For non-repairable systems, we fix $s_2 = 1$, $s_3 = 1$ and vary s_1 to illustrate the results on a 2D plot while keeping other parameters unchanged. The value of s_1 ranges from 0 to 2 with an iteration step of 0.02. This means the value of $s_1 \in \{0, 0.0.02, 0.04, ..., 2\}$. Figure 4.8b reveals that the optimal PR time approximately decreases linearly with an increase in s_1. The data-driven PR time shown in the scatter plot is distributed on both sides of the optimal solutions. The scatter plots in the two figures also reveal patterns in the overestimation and underestimation of the true optimal PR time. LASSO and ANN capture the trend of the solutions as the value of s_1 increases.

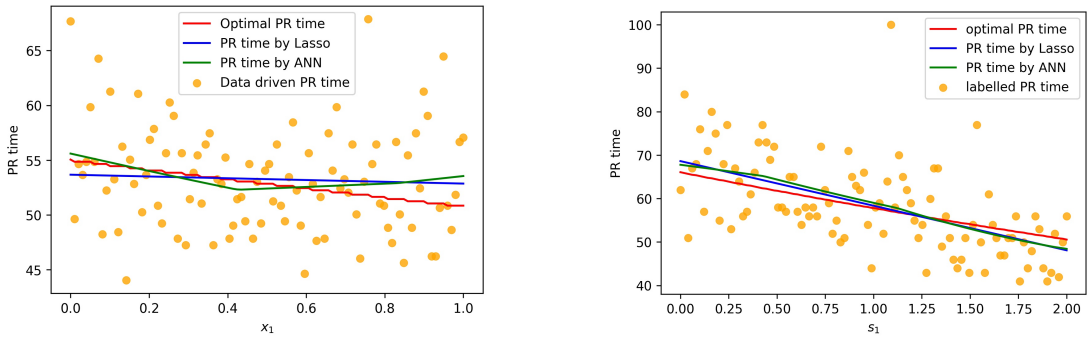

(a) Repairable systems

(b) Non-repairable systems

Figure 4.8: Comparison of decisions given by LASSO, ANN, optimal PR time and data-driven PR time

4.4.3 Comparisons with ETO methods

Comparisons with ETO methods for repairable systems

In order to compare the performance of the proposed E2E with the traditional ETO framework for repairable systems, we introduce two Cox PI based models. These serve as prediction models when the true intensity function is unknown. The Andersen-Gill (AG) model assumes that failure events are conditionally independent given covariates, with a shared baseline intensity function across events. In contrast, the Prentice-Williams-Peterson (PWP) model accounts for the order of events, allowing the baseline

intensity to vary by failure number. Using both AG and PWP models allows us to cover different assumptions about the repair process, and ensures a more robust comparison between the proposed E2E framework and traditional ETO-based approaches.

1. Andersen & Gill (AG) model: The AG model assumes independence between events and applies a PI model for repairable system i as $m_i(t|\mathbf{x}) = m_{i0}(t) \exp\left[\mathbf{b}^{\mathrm{T}}\mathbf{x}_i\right]$, where $m_0(t)$ is the baseline intensity function, $\mathbf{b}$ is the coefficient vector and $\mathbf{x}_i$ is the working condition vector of system i.

2. Prentice, Williams and Peterson (PWP) model: In this model, the baseline intensity varies from event to event, and the failure intensity function for the n^{th} failure event is expressed as $m_n(t|\mathbf{x}) = m_{0n}(t) \exp\left[\mathbf{b}_n^{\mathrm{T}}\mathbf{x}\right]$, where $m_{0n}(t)$ is the baseline intensity function for the n^{th} failure event and $\mathbf{b}_n$ is the vector of coefficient for the n^{th} failure event.

These two models have been applied in survival analysis, as well as in engineering reliability analysis (Jiang et al., 2006). The R package *survival* provides functions for predicting failure intensity based on the PWP model and AG model for recurring events. The corresponding syntax of R code can be found in existing works (Therneau, 2023; Thenmozhi et al., 2019).

To evaluate the performance, we experiment repeatedly 20 times with the parameter settings described in Table 4.1 for repairable systems. We then prepare the structural data with features and data-driven maintenance decisions. After training the supervised learning algorithms, LASSO and ANN, the average performance over 20 repeated experiments is reported in Table 4.3. The E2E methods using LASSO and ANN are denoted as E2E_LASSO and E2E_ANN in the results. The finding shows that the out-of-sample is close to the in-sample performance for repairable systems. In addition, E2E_LASSO and E2E_ANN outperform ETO methods in terms of both the MAE_{RS}^{T} and MAE_{RS}^{C}. E2E_LASSO achieves the best performance, likely due to its sparse structure and stable generalisation. E2E_ANN performs slightly worse, which may be attributed to its higher data requirement and limited training samples in the current setting. The good performance of E2E_LASSO demonstrates that simple linear models

can benefit from direct E2E learning.

Table 4.3: Comparison of E2E and ETO methods for repairable systems

	In-sample		Out-of-sample	
	MAE^T_{RS}	MAE^C_{RS}	MAE^T_{RS}	MAE^C_{RS}
E2E_LASSO	**0.783**	**0.017**	**0.762**	**0.023**
E2E_ANN	1.193	0.056	1.301	0.062
ETO_AG	1.234	0.063	1.190	0.060
ETO_PWP	1.367	0.074	1.289	0.077

Comparisons with ETO methods for non-repairable systems

We introduce two ETO methods with different time-to-failure models to compare with the proposed E2E methods for non-repairable systems. The Cox PH model estimates relative failure risk over time, while the WAFT model predicts time-to-failure directly under covariate effects. By incorporating both, we evaluate the performance of the E2E framework and ensure a fair comparison against well-established methods in the literature.

1. Cox PH model: It assumes that the features have a linear multiplicative effect on the hazard function. The hazard function is given as $\lambda(t|\mathbf{s}) = \lambda_0(t)\exp(\mathbf{c}^\intercal\mathbf{s})$, where $\lambda_0(t)$ is the baseline hazard function, and $\mathbf{c}$ denotes the vector of coefficients for each feature. Consider two individual non-repairable systems, where one has feature $\mathbf{s}'$ and the other has feature $\mathbf{s}''$. The ratio of their hazard functions is $\frac{\lambda(t|\mathbf{s}')}{\lambda(t|\mathbf{s}'')} = \frac{\exp(\mathbf{c}^\intercal\mathbf{s}')}{\exp(\mathbf{c}^\intercal\mathbf{s}'')}$ and the coefficient vector $\mathbf{c}$ can be estimated based on this ratio via MLE. The baseline hazard rate is estimated using Barlow's non-parametric method.

2. Weibull Accelerated Failure Time (WAFT) model: The WAFT model is a fully parametric model in the field of reliability analysis. It is recognised as a useful alternative tool to the Cox PH model (Volf and Timková, 2014). Unlike the Cox PH model, which assumes the effect of features is to multiply the baseline hazard rate by a constant, the WAFT model exhibits a multiplicative relationship with respect to time to failure. It quantifies how different levels of covariates either

speed up or slow down the time to failure. For the random time to failure t_i of i_{th} non-repairable system, the general form of the WAFT model is given by $\log t_i = \mathbf{a}^\intercal \mathbf{s}_i + \sigma\epsilon_i$, where $\mathbf{a}$ is the coefficient in WAFT model, ϵ_i is the random error, and σ is the scale parameter. The parameters of the WAFT model can be estimated by the MLE method.

The Cox PH model and WAFT model can be implemented using library *survival* in R or library *lifelines* in Python (Davidson-Pilon, 2019; Therneau, 2023). We repeat the experiment 20 times using the parameter settings in Table 4.2 for non-repairable systems. Subsequently, we prepare the structured data with features and labels by minimising the empirical maintenance cost function. After training the supervised learning algorithms, LASSO and ANN, we report the average MAE over 20 repeated experiments in Table 4.4. We find that E2E_LASSO exhibits the best performance, while E2E_ANN has a slightly poorer performance but is still better than ETO_WAFT. This pattern is consistent with the results observed in the repairable system setting.

Table 4.4: Comparison of E2E and ETO methods for non-repairable systems

	In-sample		Out-of-sample	
	MAE^T_{NRS}	MAE^C_{NRS}	MAE^T_{NRS}	MAE^C_{NRS}
E2E_LASSO	**1.192**	**0.015**	**1.193**	**0.015**
E2E_ANN	1.956	0.043	2.005	0.045
ETO_COXPH	1.591	0.033	1.584	0.033
ETO_WAFT	2.230	0.046	2.210	0.046

4.4.4 Sensitivity analysis

Sensitivity analysis for repairable systems

We investigate the effects of different magnitudes of shape parameter a, scale parameter b, the ratio of PR cost to minimal repair cost, the size of the data set M, and the number of systems J under the same working condition. In the sensitivity analysis, we vary the target parameter in the original settings while keeping other parameters unchanged, and we report the averaged performance metrics over 20 repeated experiments.

First, we analyse the effects of the shape parameter a in the true power-law PI

model. For each $a \in \{2.5, 3, 3.5, 4\}$, we repeat the experiment 20 times and evaluate the averaged performance. The shape parameter influences the magnitudes of the expected number of failures and thus controls the optimal PR time. Table 4.5 shows that as the value of a increases, the MAE^T_{RS} and MAE^C_{RS} for both E2E and ETO methods decrease. Among these methods, E2E_Lasso shows the best performance.

Table 4.5: Performance of the E2E and ETO methods under different shape parameter a for repairable systems

Shape parameter α	Methods	In-sample		Out-of-sample	
		MAE^T_{RS}	MAE^C_{RS}	MAE^T_{RS}	MAE^C_{RS}
$a = 2.5$	E2E_LASSO	**0.783**	**0.017**	**0.762**	**0.023**
	E2E_ANN	1.193	0.056	1.301	0.062
	ETO_AG	1.234	0.063	1.190	0.060
	ETO_PWP	1.367	0.074	1.289	0.077
$a = 3.0$	E2E_LASSO	**0.768**	**0.022**	**0.760**	**0.030**
	E2E_ANN	1.082	0.068	1.128	0.063
	ETO_AG	0.881	0.040	0.783	0.041
	ETO_PWP	1.167	0.064	1.189	0.067
$a = 3.5$	E2E_LASSO	**0.669**	**0.027**	**0.673**	**0.035**
	E2E_ANN	0.870	0.061	0.930	0.055
	ETO_AG	1.111	0.092	1.130	0.088
	ETO_PWP	1.138	0.111	1.117	0.092
$a = 4$	E2E_LASSO	**0.599**	**0.019**	**0.576**	**0.030**
	E2E_ANN	0.701	0.048	0.753	0.044
	ETO_AG	1.066	0.080	1.113	0.082
	ETO_PWP	1.133	0.094	1.252	0.095

In addition, we investigate the impact of the scale parameter of the power-law function for the failure intensity. We alter the value of the scale parameter $b \in \{25, 50, 75, 100\}$ and evaluate the average performance of the E2E and ETO methods across 20 repeated experiments. Table 4.6 presents the results, indicating that the performance of the E2E and ETO methods deteriorates with increasing values of b. Additionally, E2E_LASSO consistently shows the best performance across the range of b values.

We further examine the effects of cost ratio $\eta = \frac{c_p}{c_m}$ on the performance of E2E and ETO methods. We vary the cost ratio over the values $\eta \in \{1.5, 2.0, 2.5, 3\}$ while keeping the cost of minimal repair $c_m = 1000$. Table 4.7 displays the performance,

Table 4.6: Performance of the E2E and ETO methods under different scale parameter b for repairable systems

Scale parameter b	Methods	In-sample		Out-of-sample	
		MAE^T_{RS}	MAE^C_{RS}	MAE^T_{RS}	MAE^C_{RS}
$b = 25$	E2E_LASSO	**0.588**	**0.078**	**0.533**	**0.070**
	E2E_ANN	0.690	0.132	0.674	0.110
	ETO_AG	1.851	0.316	1.023	0.239
	ETO_PWP	138.095	284.596	136.189	281.031
$b = 50$	E2E_LASSO	**0.783**	**0.017**	**0.762**	**0.023**
	E2E_ANN	1.193	0.056	1.301	0.062
	ETO_AG	1.234	0.063	1.190	0.060
	ETO_PWP	1.367	0.074	1.289	0.077
$b = 75$	E2E_LASSO	**0.871**	**0.004**	**0.917**	**0.010**
	E2E_ANN	1.842	0.033	1.858	0.030
	ETO_AG	2.131	0.051	2.116	0.049
	ETO_PWP	2.595	0.059	2.590	0.057
$b = 100$	E2E_LASSO	**1.462**	**0.100**	**1.510**	**0.107**
	E2E_ANN	2.318	0.108	2.253	0.111
	ETO_AG	3.035	0.204	2.866	0.120
	ETO_PWP	3.938	0.397	4.610	0.425

indicating that E2E_LASSO achieves the best performance among the four methods and also shows the highest stability when the cost ratio varies.

We also show the effects of the number of working conditions on the performance of E2E and ETO methods. The number of working conditions is varied over the values $M \in \{100, 200, 300, 400\}$ while keeping other parameters unchanged. Table 4.8 reports the performance of E2E and ETO methods using the averaged performance of 20 repeated experiments when varying the value of M. The results indicate that E2E_LASSO consistently has the best performance, and increasing the number of working conditions leads to performance improvement for the E2E and ETO methods.

Finally, we examine the effects of the number of the systems on the performance of the methods by varying the number of systems over the values $J \in \{50, 100, 150, 200\}$ under the same working condition. From Figure 4.7a we know that the estimated labels become more accurate as the number of systems increases. The results presented in Table 4.9 suggest that E2E_ANN appears more sensitive to the accuracy of labels, which is closely related to the number of systems. Specifically, when $J \geq 100$, E2E_ANN

Table 4.7: Performance of the E2E and ETO methods under different cost ratio η for repairable systems

Cost ratio η	Methods	In-sample		Out-of-sample	
		MAE_{RS}^{T}	MAE_{RS}^{C}	MAE_{RS}^{T}	MAE_{RS}^{C}
$\eta = 1.5$	E2E_LASSO	**0.800**	**0.017**	**0.775**	**0.025**
	E2E_ANN	1.250	0.059	1.298	0.060
	ETO_AG	1.081	0.050	1.063	0.050
	ETO_PWP	1.210	0.062	1.228	0.066
$\eta = 2.0$	E2E_LASSO	**0.783**	**0.017**	**0.762**	**0.023**
	E2E_ANN	1.193	0.056	1.301	0.062
	ETO_AG	1.234	0.063	1.190	0.060
	ETO_PWP	1.367	0.074	1.289	0.077
$\eta = 2.5$	E2E_LASSO	**0.840**	**0.021**	**0.830**	**0.030**
	E2E_ANN	1.154	0.058	1.164	0.051
	ETO_AG	1.203	0.062	1.182	0.063
	ETO_PWP	1.471	0.081	1.433	0.083
$\eta = 3$	E2E_LASSO	**0.923**	**0.027**	**0.921**	**0.038**
	E2E_ANN	1.315	0.083	1.420	0.074
	ETO_AG	1.034	0.054	1.113	0.061
	ETO_PWP	1.564	0.102	1.524	0.099

Table 4.8: Performance of the E2E and ETO methods under different number of working conditions M for repairable systems

data size M	Methods	In-sample		Out-of-sample	
		MAE_{RS}^{T}	MAE_{RS}^{C}	MAE_{RS}^{T}	MAE_{RS}^{C}
$M = 100$	ETE_LASSO	**1.082**	**0.037**	**1.087**	**0.056**
	ETE_ANN	1.454	0.097	1.621	0.104
	ETO_AG	1.350	0.091	1.378	0.097
	ETO_PWP	1.718	0.132	1.660	0.133
$M = 200$	E2E_LASSO	**0.783**	**0.017**	**0.762**	**0.023**
	E2E_ANN	1.193	0.056	1.301	0.062
	ETO_AG	1.234	0.063	1.190	0.060
	ETO_PWP	1.367	0.074	1.289	0.077
$M = 300$	ETE_LASSO	**0.789**	**0.019**	**0.737**	**0.026**
	ETE_ANN	1.091	0.053	1.144	0.056
	ETO_AG	0.979	0.043	0.997	0.044
	ETO_PWP	1.221	0.063	1.226	0.063
$M = 400$	ETE_LASSO	**0.759**	**0.013**	**0.734**	**0.025**
	ETE_ANN	1.015	0.047	1.078	0.041
	ETO_AG	0.953	0.042	0.970	0.044
	ETO_PWP	1.099	0.053	1.108	0.053

outperforms ETO_AG and ETO_PWP.

Table 4.9: Performance of the E2E and ETO methods under different numbers of systems J for repairable systems

Number of systems J	Methods	In-sample		Out-of-sample	
		MAE^T_{RS}	MAE^C_{RS}	MAE^T_{RS}	MAE^C_{RS}
$J = 50$	E2E_LASSO	**0.783**	**0.017**	**0.762**	**0.023**
	E2E_ANN	1.193	0.056	1.301	0.062
	ETO_AG	1.234	0.063	1.190	0.060
	ETO_PWP	1.367	0.074	1.289	0.077
$J = 100$	ETE_LASSO	**0.628**	**0.014**	**0.604**	**0.016**
	ETE_ANN	0.926	0.039	1.007	0.044
	ETO_AG	0.970	0.044	0.950	0.041
	ETO_PWP	1.123	0.052	1.319	0.112
$J = 150$	ETE_LASSO	**0.710**	**0.015**	**0.664**	**0.022**
	ETE_ANN	0.904	0.039	0.967	0.039
	ETO_AG	0.951	0.0410	0.933	0.040
	ETO_PWP	1.021	0.045	1.062	0.050
$J = 200$	ETE_LASSO	**0.627**	**0.013**	**0.603**	**0.018**
	ETE_ANN	0.768	0.028	0.848	0.031
	ETO_AG	0.896	0.039	0.911	0.039
	ETO_PWP	0.935	0.040	0.973	0.042

Sensitivity analysis of non-repairable systems

We explore the effects of different magnitudes of the shape parameter α, scale parameter β, cost ratio of PR cost c_r to failure replacement cost c_f, number of working conditions N, and the number of systems Q under the same working condition.

We first examine the effects of the shape parameter α. When $\alpha = 1$, the Weibull distribution reduces to an exponential distribution, and the optimal PR time tends to be infinity, indicating that the system is replaced only at failure. We vary the value $\alpha \in \{2.5, 3, 3.5, 4\}$, and for each value of α we average the performance of 20 repeated experiments. The reason why we chose these values starting from $\alpha = 2.5$ is that when $\alpha > 2$, the slope of the hazard rate increases with time, making it applicable for seeking optimal PR time. Table 4.10 indicates that as the shape parameter α increases, the MAE^T_{NRS} and MAE^C_{NRS} for all methods decrease. E2E_LASSO shows the best performance, and E2E_ANN is better than ETO_WAFT in each case. When $\alpha = 4$,

the performance of E2E_ANN surpasses that of ETO_COXPH.

Table 4.10: Performance of the E2E and ETO methods under different shape parameter α for non-repairable systems

α	Methods	In-sample		Out-of-sample	
		MAE^T_{RS}	MAE^C_{RS}	MAE^T_{RS}	MAE^C_{NRS}
$\alpha = 2.5$	E2E_LASSO	**1.255**	**0.006**	**1.233**	**0.006**
	E2E_ANN	2.634	0.031	2.722	0.035
	ETO_COXPH	2.524	0.032	2.565	0.033
	ETO_WAFT	3.646	0.047	3.637	0.047
$\alpha = 3$	E2E_LASSO	**1.192**	**0.015**	**1.193**	**0.015**
	E2E_ANN	1.956	0.043	2.005	0.045
	ETO_COXPH	1.591	0.033	1.584	0.033
	ETO_WAFT	2.230	0.046	2.210	0.046
$\alpha = 3.5$	E2E_LASSO	**0.755**	**0.013**	**0.757**	**0.013**
	E2E_ANN	1.337	0.039	1.380	0.043
	ETO_COXPH	1.328	0.039	1.343	0.041
	ETO_WAFT	1.751	0.049	1.750	0.048
$\alpha = 4$	E2E_LASSO	**0.566**	**0.010**	**0.548**	**0.010**
	E2E_ANN	1.049	0.038	1.085	0.040
	ETO_COXPH	1.159	0.039	1.145	0.038
	ETO_WAFT	1.392	0.045	1.389	0.044

Next we examine the effects of the scale parameter β on E2E and ETO methods. The value of β varies within the range of $\{10, 20, 30, ..., 70\}$. It is important to note that, based on the characteristic of the Weibull distribution, a larger scale parameter corresponds to distribution with greater spread. Table 4.11 provides evidence that the two E2E methods attain the best performance when $\beta = 10$ and 30. E2E_LASSO still shows the best performance and E2E_ANN has smaller values of MAE^T_{NRS} and MAE^C_{NRS} compared to ETO_WAFT when β increases.

To analyse the effects of PR cost and failure replacement cost, we introduce a cost ratio γ, which is a ratio of the PR to the failure replacement cost. For each $\gamma \in \{0.2, 0.4, 0.6, 0.8\}$, the cost of failure replacement c_f remains unchanged at 4000, while the cost of PR $c_r = \gamma c_f$. A smaller cost ratio γ implies a smaller c_r relative to a fixed c_f. The value of c_r is lower than c_f because PR action helps avoid unscheduled downtime compared to a failure replacement. Table 4.12 shows that as the cost ratio γ increases, both the in-sample and out-of-sample MAE^T_{NRS} and MAE^C_{NRS} increases.

Table 4.11: Performance of the E2E and ETO methods under different scale parameter β for non-repairable systems

β	Models	In-sample		Out-of-sample	
		MAE^{T}_{NRS}	MAE^{C}_{NRS}	MAE^{T}_{NRS}	MAE^{C}_{NRS}
$\beta = 10$	E2E_LASSO	**0.221**	**0.203**	**0.230**	**0.220**
	E2E_ANN	0.287	0.384	0.283	0.417
	ETO_COXPH	4.198	4.597	3.598	4.514
	ETO_WAFT	1.353	7.071	1.349	7.223
$\beta = 30$	E2E_LASSO	**0.558**	**0.032**	**0.551**	**0.033**
	E2E_ANN	1.032	0.103	1.008	0.105
	ETO_COXPH	1.353	0.204	1.343	0.203
	ETO_WAFT	2.015	0.364	2.015	0.364
$\beta = 50$	E2E_LASSO	**0.809**	**0.013**	**0.826**	**0.013**
	E2E_ANN	1.721	0.059	1.730	0.061
	ETO_COXPH	1.700	0.064	1.684	0.064
	ETO_WAFT	2.343	0.091	2.344	0.092
$\beta = 70$	E2E_LASSO	**1.150**	**0.009**	**1.125**	**0.009**
	E2E_ANN	2.284	0.036	2.345	0.038
	ETO_COXPH	1.823	0.027	1.864	0.028
	ETO_WAFT	2.455	0.035	2.470	0.035

Furthermore, we find that E2E_LASSO maintains the smallest MAE values for the performance metrics. When $\gamma = 0.6$ and 0.8, the performance of E2E_ANN surpasses ETO_WAFT but remains inferior to ETO_COXPH.

To analyse the effects of the number of working conditions N on the performance of the proposed method, we conduct experiments with varying N, $N \in \{100, 200, 300, ...1000\}$ and keep other parameters unchanged. The averaged performance metrics of 20 repeated experiments for each number of working conditions are presented in Figure 4.9. The figure illustrates the in-sample and out-of-sample performance metrics. We observe that E2E_LASSO maintains the best performance and demonstrates stability when the sample size reaches 200. The MAE^{C}_{RS} and MAE^{T}_{RS} of E2E_ANN decrease, and are smaller than ETO_WAFT, when the number of working conditions increases from 100 to 200.

Finally, we compare the performance of the E2E and ETO methods under different numbers of systems. We evaluate the performance by averaging over 20 repeated experiments with varying the number of systems $Q \in \{50, 100, ..., 500\}$. Figure 4.10

Table 4.12: Performance of the E2E and ETO methods under different cost ratio γ

γ	Models	In-sample		Out-of-sample	
		MAE^{T}_{NRS}	MAE^{C}_{NRS}	MAE^{T}_{NRS}	MAE^{C}_{NRS}
$\gamma = 0.2$	E2E_LASSO	**0.468**	**0.026**	**0.474**	**0.028**
	E2E_ANN	0.896	0.100	0.894	0.104
	ETO_COXPH	0.792	0.074	0.815	0.080
	ETO_WAFT	0.589	0.031	0.583	0.031
$\gamma = 0.4$	E2E_LASSO	**0.539**	**0.024**	**0.553**	**0.024**
	E2E_ANN	0.943	0.075	0.974	0.080
	ETO_COXPH	0.819	0.053	0.848	0.055
	ETO_WAFT	0.819	0.037	0.826	0.037
$\gamma = 0.6$	E2E_LASSO	**0.590**	**0.014**	**0.599**	**0.014**
	E2E_ANN	1.130	0.049	1.168	0.052
	ETO_COXPH	1.101	0.045	1.110	0.046
	ETO_WAFT	1.336	0.049	1.342	0.050
$\gamma = 0.8$	E2E_LASSO	**1.011**	**0.005**	**0.990**	**0.005**
	E2E_ANN	2.045	0.023	2.089	0.023
	ETO_COXPH	1.885	0.021	1.848	0.020
	ETO_WAFT	2.996	0.042	2.975	0.041

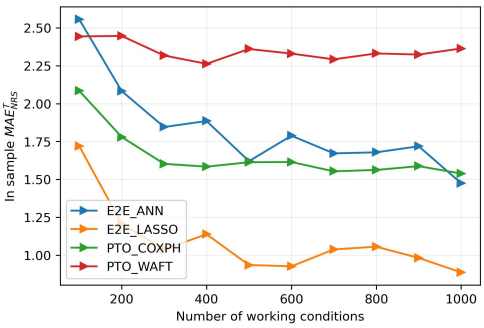

(a) In-sample MAE^{T}_{NRS}

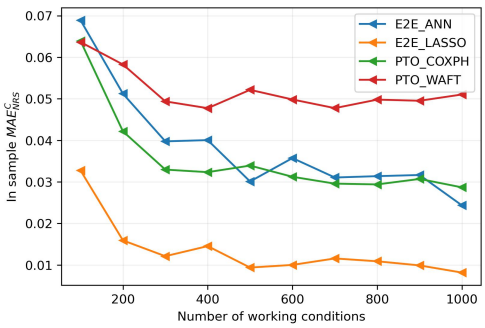

(b) In-sample MAE^{C}_{NRS}

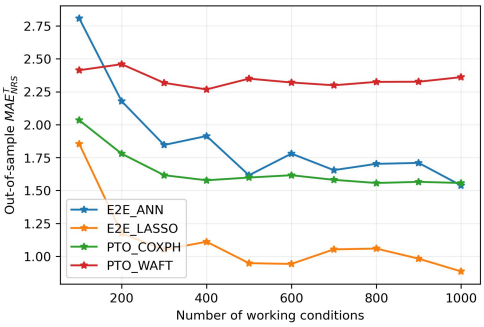

(c) Out-of-sample MAE^{T}_{NRS}

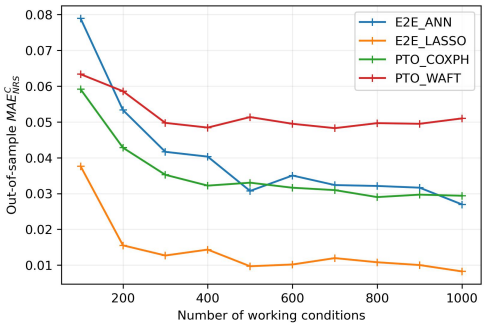

(d) Out-of-sample MAE^{C}_{NRS}

Figure 4.9: Performance of the E2E and ETO methods under different number of working conditions N for non-repairable systems

displays the comparison of MAE^{T}_{NRS} and MAE^{C}_{NRS} applying E2E and ETO methods for non-repairable systems. As the number of systems Q increases, both E2E methods show improved performance and narrow the gaps with the optimal decisions. Notably, E2E_LASSO keeps the best performance and E2E_ANN has greater improvement as the number of systems increases. These results suggest that the E2E methods have the potential to achieve better performance compared to the ETO methods for non-repairable systems, especially as the number of systems increases.

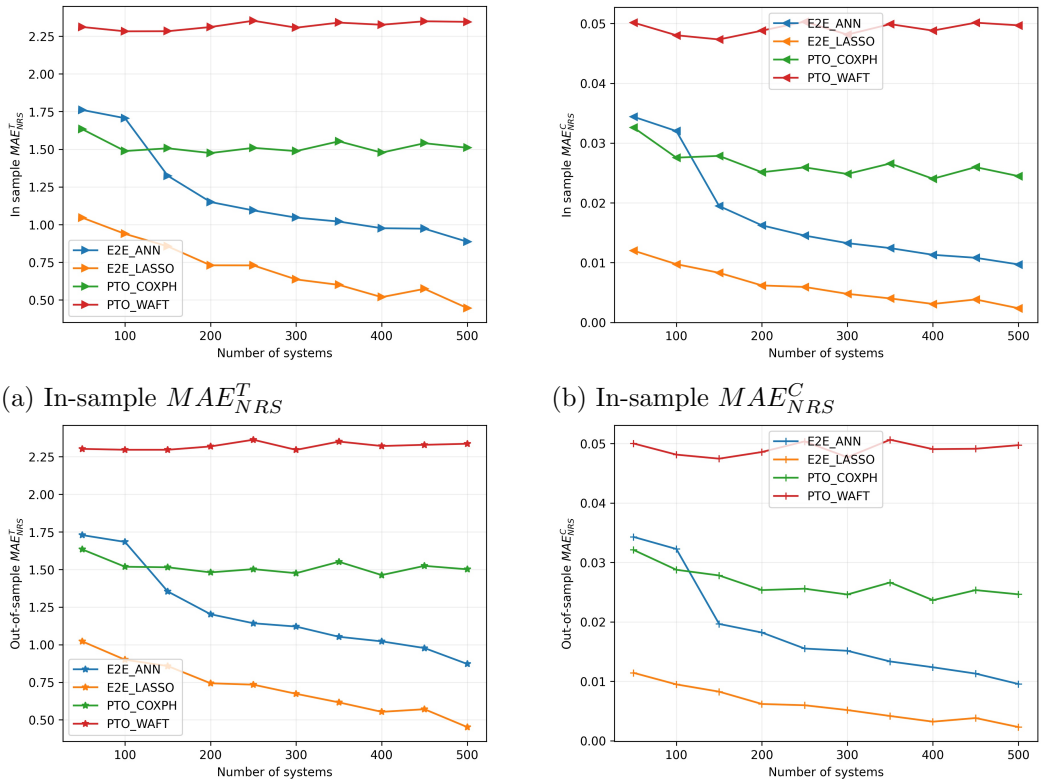

(a) In-sample MAE^{T}_{NRS}

(b) In-sample MAE^{C}_{NRS}

(c) Out-of-sample MAE^{T}_{NRS}

(d) Out-of-sample MAE^{C}_{NRS}

Figure 4.10: Comparison of MAE^{T}_{NRS} and MAE^{C}_{NRS} applying E2E and ETO methods with different number of systems for non-repairable systems

4.4.5 Discussion

The experimental comparisons and sensitivity analyses provide insights into the strengths and applicability of the proposed E2E framework relative to ETO approaches. First, across both repairable and non-repairable systems, the E2E methods consistently achieve lower MAE^T and MAE^C values than ETO methods. This suggests that directly learning maintenance decisions from operational data, without explicit model parameter estimation, reduces the error propagation in ETO pipelines. Sensitivity analyses reveal the following patterns:

- When the number of observed systems (J or Q) is small, both E2E and ETO methods suffer, but E2E methods, especially those using LASSO, show greater resilience to label estimation errors.

- As the number of systems or working conditions increases, the performance gap between E2E and ETO methods narrows. This is because larger datasets also benefit ETO models by improving parameter estimation accuracy. Nevertheless, E2E_LASSO maintains good stability and accuracy across different settings.

- When degradation processes are highly stochastic (e.g., lower shape parameters a or α), E2E methods adapt better compared to model-based ETO approaches that rely on strict distributional assumptions.

- Varying cost ratios (η or γ) affects all methods, but E2E_LASSO remains the least sensitive, suggesting its potential in applications where cost structures are uncertain or dynamic.

Moreover, it is observed that E2E_ANN, although slightly less accurate than E2E_LASSO under small sample conditions, catches up and even surpasses some ETO models as the sample size increases. This indicates that deeper architectures could benefit from larger datasets, while linear models such as LASSO remain robust under limited data. These findings collectively highlight that the proposed E2E framework shows advantages in maintenance policy making under dynamic operational conditions. The ability

to bypass explicit model assumptions allows for flexible decision-making in data-driven maintenance settings.

4.5 Summary

This Chapter proposes a data-driven E2E framework for preventive maintenance considering varied working conditions. In contrast to the traditional ETO framework that first specifies and estimates the model parameters, and then optimises the objective, the proposed E2E framework can directly determine the optimal maintenance decision in a data-driven manner. E2E framework bypasses the estimation parameters of an assumed time-to-failure model, thereby reducing the error propagation from parameter estimation to decision-making. We analytically provide the error bounds between the data-driven maintenance cost rate functions and the true maintenance cost rate functions. We also show that the data-driven maintenance cost rate functions can approximate the true maintenance cost rate functions with high probability with a large number of systems operated under the same working condition. Numerical experiments demonstrate that the proposed E2E framework outperforms the ETO framework across both repairable and non-repairable systems. In particular, E2E_LASSO consistently achieves the best performance, likely due to its robustness and sparsity in limited-data settings, while E2E_ANN performs slightly worse, possibly due to higher data requirements. These results suggest that direct policy learning can outperform prediction-based optimisation, and that a simple, linear model can be highly effective when aligned with the learning objective.

This study can be further extended in the following three directions. The first is to extend the E2E framework to CBM problems considering degradation processes and associated features, which could expand the framework's applicability to a wider range of maintenance strategies. Different to TBM, which collects time-to-failure data, CBM relies on condition monitoring data, resulting in a different labelling process based on degradation levels. Incorporating labels for maintenance actions at specific degradation levels under particular working conditions could enable future work to extend the E2E framework to CBM strategies.

Secondly, this study considers fixed values for the costs of maintenance actions, which limits the applicability of the E2E framework. Future work could address this by incorporating the costs of maintenance actions as features in the input of the supervised learning algorithm, enabling the development of an adaptive model. Additionally, feature engineering and dimensionality reduction techniques could be applied to construct a more effective set of inputs, which could further enhance the model's effectiveness.

Thirdly, we could investigate the sample complexity of the E2E framework, i.e., the number of samples required to achieve a satisfactory solution. Analysing sample complexity offers insights into the number of samples needed to achieve better performance compared to the ETO framework. The probably approximately correct learning framework could be a valuable tool for conducting a theoretical sample complexity analysis, particularly when employing machine learning algorithms.

Chapter 5

Online Learning and Control of Degradation Process under Dynamic Working Conditions

This Chapter considers a maintenance policy for discrete degradation paths caused by shock-based damage, modelled as a CP process. To address the challenges arising from component heterogeneity and dynamic working conditions, this Chapter employs Bayesian linear regression and Bayesian Poisson regression to analyse the occurrence and magnitude of shocks. Regression parameters are updated with the arrival of new observations on shock occurrences and magnitudes during online monitoring. The optimal maintenance problem is formulated as an MDP, wherein decisions are triggered when the parameters are updated and converge towards the underlying values.

The proposed framework in this Chapter is inherently data-driven: it relies on continuous condition monitoring data, including timestamped shock events and their measured magnitudes, along with contextual covariates that characterise dynamic working environments. Rather than assuming a fixed parametric form for system degradation, the model updates the posterior distributions of regression parameters directly from data. This allows the system to adaptively learn degradation dynamics specific to each component and operating condition, without relying on pre-calibrated failure models.

This Chapter not only establishes an analytically tractable degradation model that

captures the component heterogeneity and dynamic working conditions, but also investigates the structure of the optimal preventive maintenance thresholds. To ease the computation burden caused by the dynamic working conditions, a heuristic algorithm was proposed to use the most probable distribution, thereby reducing the complexity of the state space. The effectiveness of the developed approach is illustrated through comparisons in numerical experiments, and its applicability is demonstrated through an illustrative example.

5.1 Introduction

Unplanned asset failure poses a notable risk, resulting in irreversible economic loss and requiring expensive actions for restoration. To mitigate such undesirable breakdowns, it is crucial to develop an effective maintenance policy. Leveraging modern sensors and wireless communication technologies, CBM has emerged as the most advanced maintenance strategy with the availability of real-time system information. Unlike traditional time-based maintenance, which mainly addresses failure events, the CBM policy monitors the degradation status of observed components. Once the measurement exceeds the optimised critical level/threshold, maintenance actions are decided (Van der Duyn Schouten and Wartenhorst, 1994; Olsen and Tomlin, 2020).

The heterogeneity among observed components poses challenges in utilising monitoring data. In the industry, the performance of components produced from the same manufacturing process varies due to variations in raw materials and fluctuation in the manufacturing process (Yang et al., 2019). Consequently, these components may display distinct degradation paths, even when operating within the same environment (Sun et al., 2021). In such a scenario, the conventional assumption of homogeneous components is no longer suitable for accurately modelling the degradation process. In addition, neglecting latent heterogeneity may introduce errors in the maintenance optimisation stage.

The dynamic working conditions fluctuating over time, including environmental factors such as temperature, and humidity, alongside operating factors like production rate (uit het Broek et al., 2020) and operating speed, present challenges in accurately

capturing the degradation process. For example, the lifetime of lithium-ion batteries shortens when operating at lower temperatures (Kong et al., 2021). In addition, it has been verified that a higher charge/discharge rate accelerates capacity loss due to mechanical-induced damage of active particles (Snyder, 2016). Consequently, the degradation process is affected by both latent heterogeneity and varied working environments, thereby adding further challenges to the optimisation of the maintenance policy.

The traditional CBM policy involves modelling the degradation process and subsequently integrating it into the MDP (Derman, 1963; De Jonge and Scarf, 2020) or partially observable MDP (Kim, 2016; Khaleghei and Kim, 2021; Kim and Makis, 2013) to derive the optimal PM threshold, depending on whether the degradation stat can be directly observed. However, the degradation modelling and policy optimisation phases are separated in most existing works, implying first estimating the parameters of the stochastic degradation process and then optimising the PM threshold. Notably, the conventional CBM policy seldom considers the heterogeneity of components with the availability of real-time data. Furthermore, to the best of our knowledge, no research has addressed the CBM policy considering the integration of two critical factors: latent heterogeneity and dynamic working conditions.

This Chapter explicitly focuses on shock-based degradation, in which the degradation paths are driven by discrete shock events and are modelled as a CP process. This modelling choice is based on the Poisson distribution governing the number of shocks and the Normal distribution characterising the magnitude induced by each shock. To account for latent heterogeneity and dynamic working conditions, Bayesian linear regression and Bayesian Poisson regression are introduced to characterise the CP process. Regression is a statistical tool employed for delineating relationships between dependent and independent variables (Azoury and Miyaoka, 2009). The Bayesian updating approach is employed to learn the parameters of a component's degradation process using real-time data (Drent et al., 2023). The process is initiated by estimating the initial parameters based on historical data sets. Subsequently, these parameters are iteratively updated and the optimal PM threshold is determined based on the latest

observations at each epoch. The formulation of the CBM problem is modelled as a finite-horizon MDP, allowing for a comprehensive analysis of its structural properties. This approach contributes to an understanding of the adaptability of CBM policy to evolving component heterogeneity and dynamic working conditions. The contributions of this Chapter are summarised as follows:

- Latent heterogeneity and dynamic operating conditions are jointly modelled using Bayesian linear and Poisson regressions within a discrete stochastic degradation framework.
- Real-time data is incorporated to dynamically update model parameters, supporting informed maintenance decisions on the PM threshold under heterogeneous and time-varying conditions.
- An analytically tractable solution is presented, providing a structured framework for formulating and computing the CBM policy under uncertainty. It also reveals key structural properties of the policy, such as the maintenance threshold and its monotonic behaviour under certain conditions.
- A heuristic solution algorithm is designed to alleviate the computational burden due to the complexity of state space, enhancing the practical feasibility of the proposed methodology.

This Chapter is structured as follows. Section 5.2 extends the discrete stochastic degradation process, modelled as a CP process in this work, to incorporate heterogeneity and varied working conditions. Section 5.3 describes the online Bayesian updating for Bayesian linear regression and Bayesian Poisson regression. Section 5.4 presents a finite-horizon MDP model, outlining the methodology to obtain the optimal PM threshold and associated structural properties. Section 5.5 introduces the heuristic algorithm to simplify the state space. In Section 5.6, the performance of the proposed method is studied through numerical experiments and an illustrative example. In Section 5.7, the findings of this work are concluded and the future directions are discussed. Proofs are presented in the Appendix B.

5.2 Degradation model with dynamic working conditions

In this work, the CP process models degradation as a result of discrete shock events, in line with the focus on shock-based degradation. The components undergo continuous monitoring until the accumulative degradation exceeds the failure threshold, denoted as D. During the monitoring process, dedicated sensors are employed to observe the number of shocks, the magnitude of degradation caused by each shock, and the working conditions influencing the degradation process. Although the system is continuously monitored, the maintenance decision can only be made at planned downtime epochs. The maintenance decision is made at equally spaced discrete-time epochs t, $t = 0, 1, ..., T$. The interval between epoch $t-1$ and epoch t is denoted as period t. Let K be the number of random shocks that follow the Poisson process, i.e., $K \sim \text{Poisson}(\lambda)$, where λ is the expected rate of shock occurrence. The magnitude of degradation caused by each shock is denoted as X that follows the Normal distribution, $X \sim \mathcal{N}(u, \sigma^2)$, where u is the mean and σ^2 is the variance. The accumulative degradation up to epoch t is given by $Z_t = Z_{t-1} + Y_t$, where $Y_t := \sum_{i=1}^{K_t} X_{ti}$ represents the degradation increment in period t. In addition, let $\mathbf{X}_t = [X_{t1}, X_{t2}, ..., X_{tK_t}]$ be a K_t-dimensional vector comprising the degradation magnitude of each shock up to epoch t. The described accumulative shock-based degradation model is referred to as the CP process, illustrated in Figure 5.1.

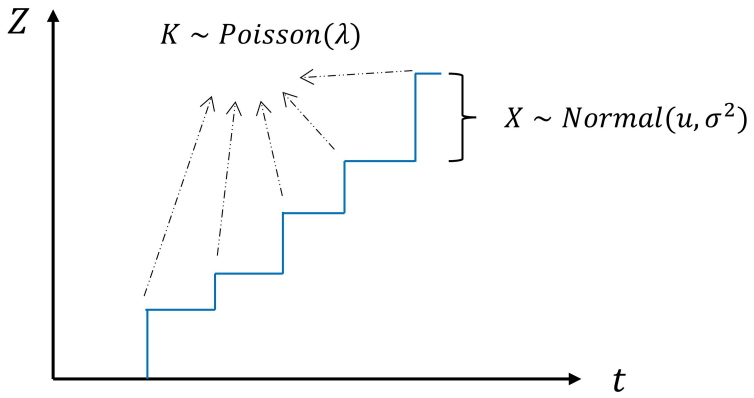

Figure 5.1: Illustration of a CP Process with Poisson-distributed shock times and normally distributed magnitude

The degradation process is influenced by the dynamic working conditions, which can be represented as covariates, also referred to as explanatory variables (Ye and Chen, 2014). Let $\mathbf{s}_t$ denote the vector $\mathbf{s}_t = [1, s_{t1}, ..., s_{tp}]$, where p is the number of covariates, and s_{ti} is the ith covariate in period t. Systems are considered to be operated in environments with well-documented historical data and real-time monitoring of working conditions, indicating future operational conditions are predictable to some extent. For a more focused investigation and to facilitate the derivation of insights, it is assumed that the decision-maker has prior knowledge of the component's working conditions. When the working conditions are partially known, they can be forecasted based on historical observations. Heterogeneity among components is modelled by incorporating covariates into the CP process that will be updated given the number and magnitudes of shocks.

Regression models estimate the relationship between a dependent variable and independent variables, enabling the analysis of how shock number and shock magnitude are influenced by working conditions. One method of accommodating regression and heterogeneity is Bayesian regression, which provides a flexible framework for modelling relationships between variables while accounting for uncertainty (Azoury and Miyaoka, 2009; Evgeniou et al., 2007). As the occurrence of random shock follows the Poisson distribution and shock magnitude follows the Normal distribution, Bayesian Poisson regression and Bayesian Normal regression are employed. These methods are applied to model the relationship among working conditions, the number of shocks and magnitude, accounting for the heterogeneity among the component population. The number of shocks in period t is represented as a Poisson regression model:

$$\lambda(\mathbf{s}_t) = e^{\mathbf{s}_t \boldsymbol{\alpha}}, K_t \sim \text{Poisson}(\lambda(\mathbf{s}_\text{t})),$$

where $\boldsymbol{\alpha}$ is a $p+1$-dimensional vector representing regression coefficients, and $\lambda(\mathbf{s}_t)$ is the link function to regression covariates with Poisson distribution. This choice of using an exponential over a linear link function is motivated by the fact that the linear function can only yield non-negative values for restricted combinations of $\mathbf{s}_t$ and $\boldsymbol{\alpha}$ (Cox and Lewis, 1966). The exponential link function is adopted in the Bayesian Poisson re-

gression model (Christiansen and Morris, 1997; Ma and Kockelman, 2006). Considering unobservable heterogeneity with random regression coefficients among components, a multivariate Normal prior on the coefficient vector $\boldsymbol{\alpha} \sim \mathcal{N}(\tilde{\boldsymbol{\alpha}}, \Sigma_\alpha)$ is introduced, where $\tilde{\boldsymbol{\alpha}}$ is the prior mean and Σ_α is covariance matrix.

Similarly, the standard Normal regression model is employed for the magnitude of each shock in period t,

$$u(\mathbf{s}_t) = \mathbf{s}_t\boldsymbol{\beta}, X_t \sim \mathcal{N}(u(\mathbf{s}_t), \sigma^2),$$

where $\boldsymbol{\beta}$ is a $p+1$-dimensional vector of regression coefficients, σ^2 is the variance and $u(\mathbf{s}_t)$ is the link function. It is worth noting that a linear link function is used in Normal regression to derive analytical solutions (Gelman et al., 1995). When the relationship between the mean of magnitude and covariates displays non-linear patterns, it is expected to be transformed into linear functions. Techniques like Taylor Expansion and piecewise linear approximations can be employed for this purpose (Ban and Rudin, 2019). For the prior information on the random vector $\boldsymbol{\beta}$, it is assumed that $\boldsymbol{\beta} \sim \mathcal{N}(\tilde{\boldsymbol{\beta}}, \Sigma_\beta)$, where $\tilde{\boldsymbol{\beta}}$ is the prior mean and Σ_β is prior covariance matrix. As pointed out by (Gelman et al., 1995), the parameter σ is typically less influential than $\boldsymbol{\beta}$, and the variance σ^2 can be estimated using sample variance, estimating the coefficient vector $\boldsymbol{\beta}$ is consequently more challenging. Given the non-negativity of shock magnitudes, it is assumed that the mean and variance satisfying $\mathbf{s}_t\tilde{\boldsymbol{\beta}} \geq 3\sigma(1 + \mathbf{s}_t\Sigma_\beta\mathbf{s}_t^\intercal)$ for $\forall t \in \{0, 1, ..., T\}$, ensuring that the probability of negative magnitude is negligible. The comprehensive framework of modelling the heterogeneous degradation process and dynamic working conditions, is depicted in Figure 5.2.

Following the Bayesian updating process, the initial task of the maintenance decision-maker involves estimating the prior parameters $\tilde{\boldsymbol{\alpha}}_0, \tilde{\boldsymbol{\beta}}_0, \Sigma_{\alpha_0}, \Sigma_{\beta_0}$ based on historical observations. MLE is applied to carry out this estimation based on the historical data set (Si et al., 2018), and the details are provided in Appendix B.6. When applying for a newly installed component, the online inference at epoch $t+1$ is leveraged to update the posterior distribution of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ and determine the conditional distribution of $(z_{t+1}|\mathbf{s}_1, ..., \mathbf{s}_t, k_1, .., k_t, \mathbf{x}_1, ..., \mathbf{x}_t)$. With the updated posterior distribution, the decision maker can reduce the uncertainty of the degradation process and make maintenance

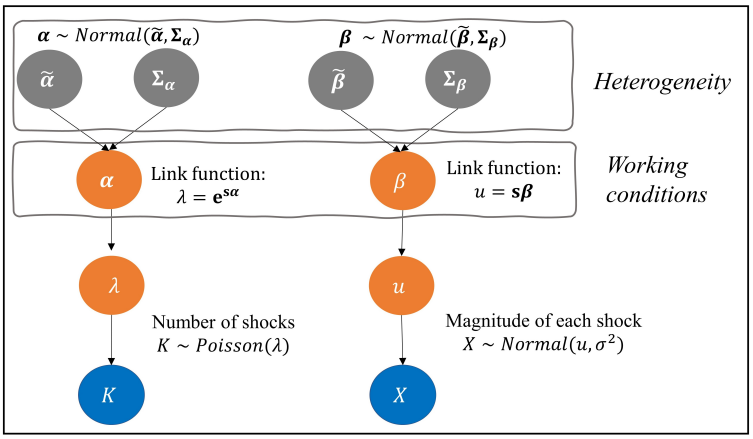

Figure 5.2: Illustration of CP process considering working condition and heterogeneity

decisions to minimise the overall maintenance cost. The described Bayesian updating and maintenance decision process is depicted in Figure 5.3.

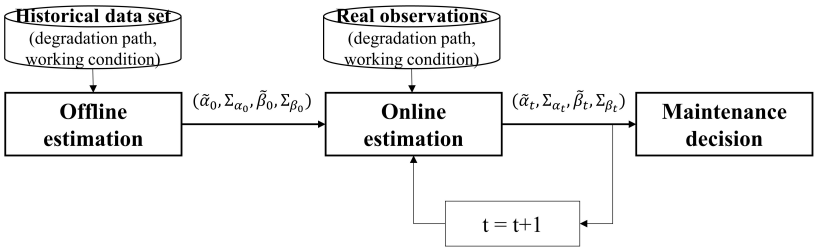

Figure 5.3: Flow chart of Bayesian updating and maintenance decision process considering heterogeneity and dynamic working conditions

5.3 Bayesian updating of degradation parameters

The initial estimates of the parameters $\tilde{\boldsymbol{\alpha}}_0, \tilde{\boldsymbol{\beta}}_0, \Sigma_{\alpha_0}, \Sigma_{\beta_0}$ are derived using MLE based on the historical degradation data set. This is also referred to as the offline update process. Subsequently, the online Bayesian updating method is utilised to update the posterior distribution on $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ by incorporating the in-situ degradation measurements and the real-time working conditions (Drent et al., 2023; Alós-Ferrer and Garagnani, 2023).

Denote the covariate matrix of working conditions as

$$\mathbf{S}_t = \begin{bmatrix} 1 & s_{11} & \cdots & s_{1p} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & s_{t1} & ... & s_{tp} \end{bmatrix}$$

and the vector of the occurrence of shocks $\mathbf{k}_t = [k_1, \cdots, k_t]^\intercal$, where s_{tp} represents the p covariate in period t. The posterior distribution $\mathbb{P}(\boldsymbol{\alpha}|\mathbf{S}_t, \mathbf{k}_t)$ is calculated with Bayes rule as follows:

$$\mathbb{P}(\boldsymbol{\alpha}|\mathbf{S}_t, \mathbf{k}_t) \propto \mathbb{P}(\mathbf{k}_t|\mathbf{S}_t, \boldsymbol{\alpha})\mathbb{P}(\boldsymbol{\alpha})$$

Since $K \sim Poisson(\lambda(\mathbf{s}))$, the likelihood is

$$\mathbb{P}(\mathbf{k}_t|\mathbf{S}_t, \boldsymbol{\alpha}) = \prod_{i=1}^{t} \frac{e^{\mathbf{s}_i\boldsymbol{\alpha}k_i}e^{-e^{\mathbf{s}_i\boldsymbol{\alpha}}}}{k_i!}. \tag{5.1}$$

As the occurrence of shocks k_i increases, Equation (5.1) can be approximated using Normal approximation (Chan and Vasconcelos, 2009). The approximation of Equation (5.1) is

$$\begin{aligned} \mathbb{P}(\mathbf{k}_t|\mathbf{S}_t, \boldsymbol{\alpha}) &\approx \prod_{i=1}^{t} \frac{1}{k_i}\mathcal{N}(\mathbf{s}_t\boldsymbol{\alpha}|\log k_i, k_i^{-1}) \\ &= \prod_{i=1}^{t} \frac{1}{(2\pi)^{\frac{t}{2}}} \exp\left[\sum_{i=1}^{t} -\frac{k_i}{2}(\mathbf{s}_i\boldsymbol{\alpha} - \log k_i)^2)\right] \\ &= \frac{|\Sigma_{k_t}|^{\frac{1}{2}}}{(2\pi)^{\frac{t}{2}}} \exp\left[-\frac{1}{2}(\mathbf{S}_t\boldsymbol{\alpha} - \log \mathbf{k}_t)^\intercal\Sigma_{k_t}^{-1}(\mathbf{S}_t\boldsymbol{\alpha} - \log \mathbf{k}_t)\right], \end{aligned} \tag{5.2}$$

where $\Sigma_{k_t} = \text{diag}(\frac{1}{k_1}, ..., \frac{1}{k_t})$. The Normal approximation is provided in the Appendix B.7. The prior distribution of $\boldsymbol{\alpha}$ is assumed to follow a multivariate Normal distribution as

$$\mathbb{P}(\boldsymbol{\alpha}) = \frac{1}{\sqrt{(2\pi)^{p+1}|\Sigma_{\alpha_0}|}} \exp\left[-\frac{1}{2}(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}_0)^\intercal\Sigma_{\alpha_0}^{-1}(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}_0)\right],$$

where $\boldsymbol{\alpha}_0$ is the prior mean vector and Σ_{α_0} is the prior covariance matrix. Historical observations of degradation and working conditions from the same or similar type of

components can be used as the prior data set to estimate $\boldsymbol{\alpha}_0$ and Σ_{α_0}. Therefore, the posterior distribution $\boldsymbol{\alpha}$ is

$$\begin{aligned}\mathbb{P}(\boldsymbol{\alpha}|\mathbf{S}_t, \mathbf{k}_t) &\propto \exp\left[-\frac{1}{2}(\mathbf{S}_t\boldsymbol{\alpha} - \log \mathbf{k}_t)^\intercal \Sigma_{k_t}^{-1}(\mathbf{S}_t\boldsymbol{\alpha} - \log \mathbf{k}_t) - \frac{1}{2}(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}_0)^\intercal \Sigma_{\alpha_0}^{-1}(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}_0)\right] \\ &\approx \mathcal{N}(\tilde{\boldsymbol{\alpha}}_t, \Sigma_{\alpha_t}),\end{aligned} \tag{5.3}$$

where $\tilde{\boldsymbol{\alpha}}_t = (\mathbf{S}_t^\intercal \Sigma_{k_t}^{-1}\mathbf{S}_t + \Sigma_{\alpha_0}^{-1})^{-1}(\mathbf{S}_t^\intercal \Sigma_{k_t}^{-1} \log \mathbf{k}_t + \Sigma_{\alpha_0}^{-1}\tilde{\boldsymbol{\alpha}}_0)$, and $\Sigma_{\alpha_t} = (\mathbf{S}_t^\intercal \Sigma_{k_t}^{-1}\mathbf{S}_t^\intercal + \Sigma_{\alpha_0}^{-1})^{-1}$. Denote $\mathbf{P}_t = \mathbf{S}_t^\intercal \Sigma_{k_t}^{-1}\mathbf{S}_t$ and $\mathbf{Q}_t = \mathbf{S}_t^\intercal \Sigma_{k_t}^{-1} \log \mathbf{k}_t$. The posterior distribution of $\boldsymbol{\alpha}$ is updated as Lemma 5.1.

Lemma 5.1 *Given the observed shocks $\mathbf{k}_t$ and working conditions $\mathbf{S}_t$, the posterior distribution of $\boldsymbol{\alpha}$ at epoch t can be approximated by multivariate Normal distribution with mean $\tilde{\boldsymbol{\alpha}}_t = (\mathbf{P}_t + \Sigma_{\alpha_0}^{-1})^{-1}(\mathbf{Q}_t + \Sigma_{\alpha_0}^{-1}\tilde{\boldsymbol{\alpha}}_0)$ and covariance matrix $\Sigma_{\alpha_t} = (\mathbf{P}_t + \Sigma_{\alpha_0}^{-1})^{-1}$, i.e., $\boldsymbol{\alpha} \sim \mathcal{N}(\tilde{\boldsymbol{\alpha}}_t, \Sigma_{\alpha_t}|\mathbf{P}_t, \mathbf{Q}_t)$, where $\mathbf{P}_t = \mathbf{S}_t^\intercal \Sigma_{k_t}^{-1}\mathbf{S}_t$ and $\mathbf{Q}_t = \mathbf{S}_t^\intercal \Sigma_{k_t}^{-1} \log \mathbf{k}_t$.*

Proofs of the Lemmas and Propositions are presented in Appendix B.1-B.5. Similarly, the posterior distribution of $\boldsymbol{\beta}$ can also be derived following the standard Bayesian process (Deisenroth et al., 2020). Denote the vector of the occurrence of shocks $\mathbf{x}_t = [x_1, \cdots, x_J]^\intercal$, where J is the total number of shocks $J = \sum_{i=1}^{t} k_i$, and let the corresponding covariate matrix

$$\mathcal{S}_t = \begin{bmatrix} 1 & s_{11} & \cdots & s_{1p} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & s_{J1} & ... & s_{Jp} \end{bmatrix}.$$

Denote $\mathbf{M}_t = \mathcal{S}_t\mathcal{S}_t^\intercal + \sigma^2\Sigma_{\beta_0}^{-1}$ and $\mathbf{N}_t = \mathcal{S}_t^\intercal \mathbf{x}_t$, the updating of the posterior of $\boldsymbol{\beta}$ is summarised in Lemma 5.2.

Lemma 5.2 *Given the observed degradation magnitudes $\mathbf{x}_t$ and working conditions $\mathcal{S}_t$, the posterior distribution of $\boldsymbol{\beta}$ at epoch t is multivariate Normal with mean $\tilde{\boldsymbol{\beta}}_t = \mathbf{M}_t^{-1}(\mathbf{N}_t + \sigma^2\Sigma_{\beta_0}^{-1}\tilde{\boldsymbol{\beta}}_0)$ and covariance matrix $\Sigma_{\beta_t} = \sigma^2\mathbf{M}_t^{-1}$, i.e., $\boldsymbol{\beta} \sim \mathcal{N}(\tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t}|\mathbf{N}_t)$, where $\mathbf{M}_t = \mathcal{S}_t\mathcal{S}_t^\intercal + \sigma^2\Sigma_{\beta_0}^{-1}$ and $\mathbf{N}_t = \mathcal{S}_t^\intercal \mathbf{x}_t$.*

Proof of Lemma 5.2 is similar to that of Lemma 5.1, therefore, the detailed proof is omitted to avoid repetition. Based on the updated regression coefficient vectors $\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t$ and covariance matrices $\Sigma_{\alpha_t}, \Sigma_{\beta_t}$, the number of shocks in period $t+1$ can be approximated as a negative binomial distribution, and the magnitude of each shock in period $t+1$ follows Normal distribution. The predictive distributions of K_{t+1} and X_{t+1} are summarised in Proposition 5.1.

Proposition 5.1 *1. Given the working conditions, the predictive shock numbers K_{t+1} in period $t+1$ is approximated as a negative binomial distribution,*

$$K_{t+1} \sim \text{NegBin}(\mathrm{e}^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}, \mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal}|\mathbf{P}_t, \mathbf{Q}_t),$$
$$\mathbb{P}(k_{t+1}|\mathbf{P}_t, \mathbf{Q}_t) = \frac{\Gamma(a+k_{t+1})}{\Gamma(k_{t+1}+1)\Gamma(a)}(\frac{1}{1+b})^a(\frac{b}{1+b})^{k_{t+1}},$$

where $a = (\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal})^{-1}$, $b = \mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal}e^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}$, $\tilde{\boldsymbol{\alpha}}_t = (\mathbf{P}_t + \Sigma_{\alpha_0}^{-1})^{-1}(\mathbf{Q}_t + \Sigma_{\alpha_0}^{-1}\tilde{\boldsymbol{\alpha}}_0)$ and $\Sigma_{\alpha_t} = (\mathbf{P}_t + \Sigma_{\alpha_0}^{-1})^{-1}$.

2. Given the working condition, the predictive shock magnitude X_{t+1} in period $t+1$ follows Normal distribution,

$$X_{t+1} \sim \mathcal{N}\left(\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t, \sigma^2(1+\mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal})|\mathbf{N}_t\right),$$

where $\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t$ is mean, $\sigma^2(1+\mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal})$ is variance, $\tilde{\boldsymbol{\beta}}_t = \mathbf{M}_t^{-1}(\mathbf{N}_t + \sigma^2\Sigma_{\beta_0}^{-1}\tilde{\boldsymbol{\beta}}_0)$ and $\Sigma_{\beta_t} = \sigma^2\mathbf{M}_t^{-1}$

In proposition 5.1, the predictive distributions of K_{t+1} and X_{t+1} are updated using the working conditions, number of shocks and the associated shock magnitude. Due to the lack of conjugacy between the Poisson likelihood and the multivariate Normal prior, the predictive distribution of K_{t+1} cannot be derived analytically. Thus, the Normal distribution is used to approximate the log-gamma distribution, as the log-gamma shape can resemble a normal distribution. This approximation is typically sufficient when the deviations between the two distributions are negligible, enabling the use of conjugacy to obtain an analytic result. Combined with the solution in Proposition 5.1, the degradation increment Y_{t+1} in period $t+1$ is further derived in Proposition 5.2.

Proposition 5.2 *Given the working condition, the predictive distribution of degradation increment Y_{t+1} in period $t+1$ is*

$$f(Y_{t+1}=y|\mathbf{P}_t,\mathbf{Q}_t,\mathbf{N}_t)=\sum_{k=0}^{\infty}\frac{\Gamma(a+k)}{\Gamma(k+1)\Gamma(a)}(\frac{1}{1+b})^a(\frac{b}{1+b})^k\left(2\pi k(\sigma^2+d)\right)^{-\frac{1}{2}}e^{-\frac{1}{2}\frac{(y-kc)^2}{k(\sigma^2+d)}},$$

where $a=(\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^\intercal)^{-1}$, $b=\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^\intercal e^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}$, $c=\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t, d=\mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^\intercal$.

Proposition 5.2 can be used to predict the degradation increment in period $t+1$ based on the working condition $\mathbf{s}_{t+1}$, the posterior distributions on $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ at epoch t. Earlier work shows that the stochastic order of the random variable Y_{t+1} is decreasing in t and increasing in the accumulative degradation when the working conditions are not considered (Drent et al., 2023). However, when considering the dynamic working conditions, the monotonic behaviour will no longer hold for the Bayesian linear regression and Bayesian Poisson regression. To hold the monotonicity property, the covariates must satisfy certain conditions, as being formalised in the Proposition 5.3.

Proposition 5.3 *Let* $a_{t+1}=(\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^\intercal)^{-1}$, $a_{t+1}^{'}=(\mathbf{s}_{t+1}\Sigma_{\alpha_t}^{\prime}\mathbf{s}_{t+1}^\intercal)^{-1}$, $b_{t+1}=\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_t^\intercal e^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}$ *and* $b_{t+1}^{\prime}=\mathbf{s}_{t+1}\Sigma_{\alpha_t}^{\prime}\mathbf{s}_{t+1}^\intercal e^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t^{\prime}}$, $c_{t+1}=\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t$ *and* $c_{t+1}^{\prime}=\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t^{\prime}$, *if* $a_{t+1}\leq a_{t+1}^{\prime}$, $b_{t+1}\leq b_{t+1}^{\prime}$ *and* $c_{t+1}\leq c_{t+1}^{\prime}$, *the stochastic order* $\langle Y_{t+1}|\mathbf{P}_t,\mathbf{Q}_t,\mathbf{N}_t\rangle\preceq\langle Y_{t+1}|\mathbf{P}_t^{\prime},\mathbf{Q}_t^{\prime},\mathbf{N}_t^{\prime}\rangle$.

The stochastic order facilitates the study of degradation behaviours to compare the expected degradation increments for two components in period $t+1$ when heterogeneity exists. It also serves as a fundamental step for the monotonicity of value functions in the MDP modelling. The stochastic order in Proposition 5.3 is illustrated through an example.

Example 5.1 *Let us consider a case where a single covariate s influences the degradation process. This example considers 2 periods, and $s_1=0.6, s_2=0.5$ in each period. At epoch 0, the prior of $\boldsymbol{\alpha}$ has mean $\boldsymbol{\alpha}_0=[0.1,0.5]$ and covariance matrix $\Sigma_{\alpha_0}=\left[\begin{smallmatrix}4&-7\\-7&14\end{smallmatrix}\right]$. The prior of $\boldsymbol{\beta}$ has mean $\boldsymbol{\beta}_0=[0.1,2]$ and covariance matrix $\Sigma_{\beta_0}=\left[\begin{smallmatrix}2&-4\\-4&10\end{smallmatrix}\right]$.*

The number of shocks follows Poisson distribution, and the magnitude of each shock follows Normal distribution with known standard deviation $\sigma^2=0.02$. Consider two cases of realisations of shocks and magnitudes in period 1. In case A, the number of

shock $k_1 = 2$*, the magnitude of each shock* $x_{11} = 1$ *and* $x_{12} = 2$*. In case* $A^{'}$*, the number of shock* $k_1^{'} = 3$*, the magnitude of each shock* $x_{11}^{'} = 2$*,* $x_{12}^{'} = 3$ *and* $x_{13}^{'} = 4$*. Given the observations of shocks in period 1, the posterior distribution of* $\boldsymbol{\alpha}$ *can be updated according to Lemma 5.1. For case A,* $\boldsymbol{\alpha} \sim \mathcal{N}(\tilde{\boldsymbol{\alpha}}_1, \Sigma_{\alpha_1})$*,* $\tilde{\boldsymbol{\alpha}}_1 = [0.048, 0.860]^\intercal$*,* $\Sigma_{\alpha_1} = \left[\begin{smallmatrix} 3.965 & -6.754 \\ -6.754 & 12.281 \end{smallmatrix}\right]$*. For case* $A^{'}$*,* $\boldsymbol{\alpha}^{'} \sim \mathcal{N}(\tilde{\boldsymbol{\alpha}}_1^{'}, \Sigma_{\alpha_1}^{'})$*,* $\tilde{\boldsymbol{\alpha}}_1^{'} = [-0.043, 1.504]^\intercal$*,* $\Sigma_{\alpha_1}^{'} = \left[\begin{smallmatrix} 3.959 & -6.712 \\ -6.712 & 11.986 \end{smallmatrix}\right]$.

Figure 5.4a displays the regression lines for the mean of random variable $\log(K)$ *in period 2, with the covariate* s_2 *in the x-axis. The vector of covariates is represented as* $\boldsymbol{s}_2 = [1, s_2]$*, incorporating the intercept in the regression. Figure 5.4b illustrates the value of* a_2 *and* $a_2^{'}$*. It is evident that* $a_2^{'} \geq a_2$ *holds for any* s_2*. In addition,* $b_2^{'} \geq b_2$ *is satisfied when* $s_2 \geq 0.144$*. Hence, in period 2, the stochastic order of* K *of case* $A^{'}$ *dominates the stochastic order of case A when* $s_2 \geq 0.144$ *in this example. For a more detailed perspective, Figure 5.4d illustrates* $\mathbb{P}(K > k)$ *during period 2 given* $s_2 = 0.5$*.*

Similarly, the stochastic order of X *in period 2 can be explored. Given the observations of shock magnitude in period 1, the posterior distribution of* $\boldsymbol{\beta}$ *can be updated according to Lemma 5.2. For case A,* $\boldsymbol{\beta} \sim \mathcal{N}(\tilde{\boldsymbol{\beta}}_1, \Sigma_{\beta_1})$*,* $\tilde{\boldsymbol{\beta}}_1 = [0.075, 2.121]$*,* $\Sigma_{\beta_1} = \left[\begin{smallmatrix} 1.952 & -3.758 \\ -3.758 & 8.788 \end{smallmatrix}\right]$*. While for case* $A^{'}$*,* $\boldsymbol{\beta}^{'} \sim \mathcal{N}(\tilde{\boldsymbol{\beta}}_1^{'}, \Sigma_{\beta_1}^{'})$*,* $\tilde{\boldsymbol{\beta}}_1^{'} = [-0.176, 3.376]$*,* $\Sigma_{\beta_1}^{'} = \left[\begin{smallmatrix} 1.935 & -3.676 \\ -3.676 & 8.378 \end{smallmatrix}\right]$*. Figure 5.5a plots the regression lines for the average magnitude of shocks in period 2. The x-axis is the covariate* s_2 *and the y-axis is* $\mathbf{s}_2\tilde{\beta}_1$*, which is the mean of* X *in period 2. It is straightforward to see that* $\mathbf{s}_2\tilde{\beta}_1 \leq \mathbf{s}_2\tilde{\beta}_1^{'}$ *holds if* $s_2 \geq 0.2$ *and fails otherwise. Because the variance of magnitude* σ *is not working conditions dependent, the stochastic order of* X *in period 2 is achieved when the means are ordered. Figure 5.4d plots* $\mathbb{P}(X > x)$ *in period 2 when* $s_2 = 0.5$*. The stochastic order of* X *of case* $A^{'}$ *dominates the stochastic order of case A if* $s_2 \geq 0.2$ *in this example. Therefore, it can be concluded that the stochastic order holds* $\langle Y_2 | \mathbf{P}_1, \mathbf{Q}_1, \mathbf{N}_1 \rangle \preceq \langle Y_2 | \mathbf{P}_1^{\prime}, \mathbf{Q}_1^{\prime}, \mathbf{N}_1^{\prime} \rangle$ *in this example if* $s_2 \geq 0.2$*, as* $Y_2 = \sum_{i=1}^{K_2} X_{2i}$*.*

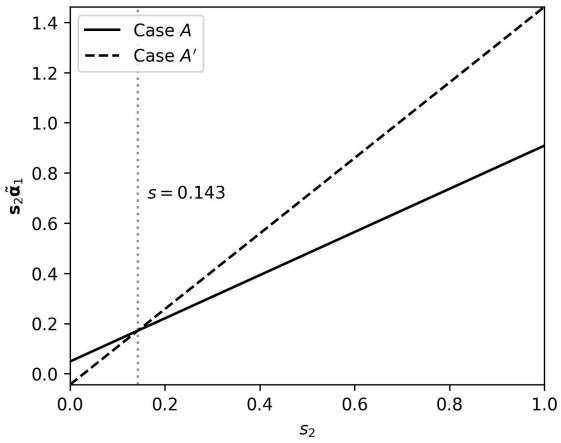

(a) Regression line for the mean of $\log(K)$

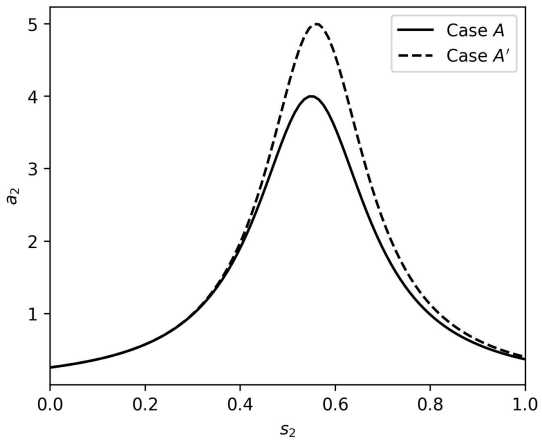

(b) a_2 and a_2' for case A and case A'

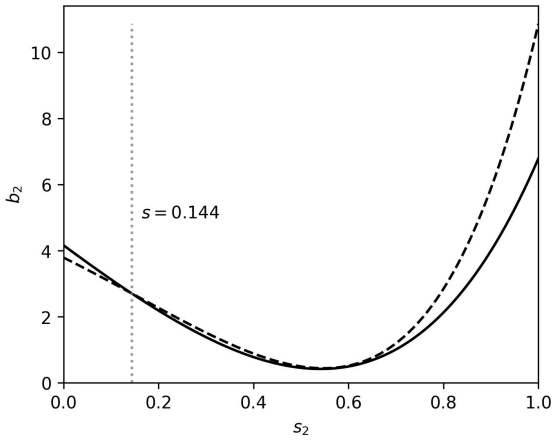

(c) b_2 and b_2' for case A and case A'

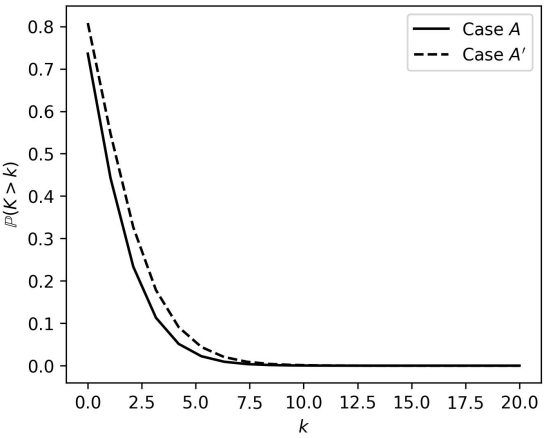

(d) $\mathbb{P}(K > k)$ when $s_2 = 0.5$

Figure 5.4: Illustrations comparing the stochastic order of the number of shocks in period 2

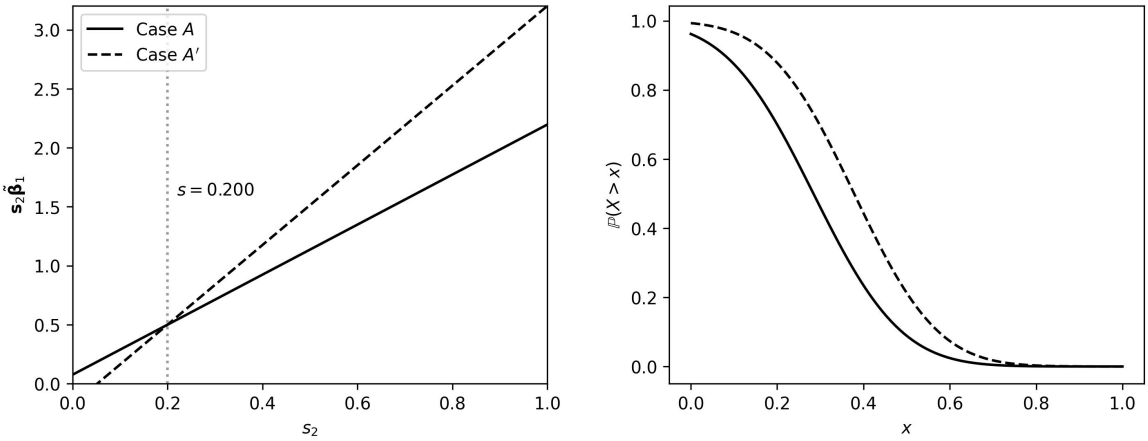

(a) Regression line for the mean of X (b) $\mathbb{P}(X > x)$ when $s_2 = 0.5$

Figure 5.5: Illustrations comparing the stochastic order of the magnitude of shocks in period 2

5.4 Maintenance modelling and optimal control policy

At each decision epoch, two maintenance actions are considered: Corrective Maintenance (CM) and PM. The component is failed when the cumulative degradation exceeds the threshold D. Upon failure at a decision epoch, CM is performed immediately incurring a cost of c_r. At epoch t, if the cumulative degradation is less than the threshold D, the component can either continue operating without any maintenance intervention or PM is implemented at a cost of c_p. It is assumed that the time required for CM and PM is negligible, compared with the operating horizon. Without loss of generality, assume $0 < c_p < c_r < \infty$, representing the economic loss of unscheduled maintenance resulting from breakdowns, such as late deliveries and shortages.

The system state at epoch t is defined by the vector $(z_t, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t, \boldsymbol{\chi}_t)$, which compactly represents the degradation history, shock behaviour, and future working condition trajectory. This choice of state representation is informed by both statistical sufficiency and predictive relevance. The variables $\mathbf{P}_t$, $\mathbf{Q}_t$, and $\mathbf{N}_t$ serve as cumulative sufficient statistics derived from past observations and covariates, and they form the basis for updating the posterior distributions of model parameters. Proposition 5.2 shows that the predictive distribution of the degradation increment Y_{t+1} depends explicitly on

the covariates $\mathbf{s}_{t+1}$ and on the posterior of these parameters, which are functions of the cumulative statistics. Therefore, these elements are not only essential for maintaining the Markov property but also serve as direct inputs for predicting future degradation dynamics under uncertain working conditions. The inclusion of z_t ensures that the absolute degradation level is tracked, which is central to maintenance triggering and cost evaluation. The future covariate sequence $\boldsymbol{\chi}_t = [\mathbf{s}_t, \mathbf{s}_{t+1}, \ldots, \mathbf{s}_T]$ enables anticipatory policy decisions by incorporating predictive knowledge about upcoming environments.

The decision maker selects one of three actions from the set {CM, PM, NULL}, where NULL indicates no maintenance action. After either CM or PM is taken, the component is restored to the as-good-as-new state with the same parameters, resulting in an updated component degradation state of $(0, \mathbf{P}_0, \mathbf{Q}_0, \mathbf{N}_0, \boldsymbol{\chi}_t)$, where $\mathbf{P}_0$ is a $(p+1) \times (p+1)$ zero matrix, and $\mathbf{Q}_0$ and $\mathbf{N}_0$ are zero vectors of dimension $p+1$.

The value function that denotes the expected maintenance cost from epoch t to epoch T is $V_{t,\boldsymbol{\chi}_t}(\boldsymbol{A}_t)$, where A_t is the state set that $\boldsymbol{A}_t = (z_t, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t)$. $V_{t,\boldsymbol{\chi}_t}(\boldsymbol{A}_t)$ represent the expected discounted maintenance cost from the current epoch t to the epoch T at state $\boldsymbol{A}_t$. The maintenance policy aims to minimise the expected value function and find the optimal action at each decision epoch. The value function follows the Bellman optimality equation as

$$V_{t,\boldsymbol{\chi}_t}(\boldsymbol{A}_t) = \begin{cases} \min\left\{c_p + V_{t,\boldsymbol{\chi}_t}(\boldsymbol{A}_0), \gamma \mathbb{E}\left[V_{t+1,\boldsymbol{\chi}_{t+1}}(\boldsymbol{A}_{t+1}|\boldsymbol{A}_t)\right]\right\}, & z_t < D, \\ c_r + V_{t,\boldsymbol{\chi}_t}(\boldsymbol{A}_0), & z_t \geq D, \end{cases} \tag{5.4}$$

where $\gamma \in (0,1)$ is the discount factor. When the state transit from $\boldsymbol{A}_t$ to $\boldsymbol{A}_{t+1}$, $\mathbf{P}_{t+1} = \mathbf{P}_t + \mathbf{s}_{t+1}^\intercal \frac{1}{k_{t+1}} \mathbf{s}_{t+1}$, $\mathbf{Q}_{t+1} = \mathbf{Q}_t + \mathbf{s}_{t+1}^\intercal \frac{1}{k_{t+1}} \log k_{t+1}$, $\mathbf{N}_{t+1} = \mathbf{N}_t + \mathbf{s}_{t+1}\mathbf{x}_{t+1}$, $\boldsymbol{\chi}_{t+1} := (\mathbf{s}_t, \mathbf{s}_{t+1}, ..., \mathbf{s}_T)$. After the last decision epoch T, there are no future values of working condition covariates. Thus, the terminal value function is defined as $V_{T+1,\boldsymbol{\chi}_{T+1}}(\cdot) = 0$ regardless of the component states. Following the Bellman equation, the optimal maintenance decision is a control limit policy, as stated in Proposition 5.4.

Proposition 5.4 *For any given working condition $\boldsymbol{\chi}_t$, the optimal policy for the problem described in Equation* (5.4) *is a control limit policy. There exists a sequence $\zeta_{t,\boldsymbol{\chi}_t}(\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t, \boldsymbol{\chi}_t)$ at epoch t, $t = 0, 1, ..., T$, such that the optimal action at state*

$(z_t, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t)$ *is PM if* $z_t \geq \zeta_{t,\boldsymbol{\chi}_t}(\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t)$.

5.4.1 Heuristic solution algorithm

In this section, backward dynamic programming is applied to solve the finite MDP formulation discussed in Equation (5.4). The degradation is discretised into L states with equal increment $\delta = \frac{D}{L}$, $L \in \mathbf{N}^+$, as the accumulative degradation amount Z_t is a continuous random variable. Te component's accumulative degradation state is denoted as $\boldsymbol{\Delta} = \{\delta, 2\delta, ..., D-\delta, D\}$. The degradation level of the component is in the state $j\delta$ when the accumulative degradation falls within the range $(j-1)\delta \leq Z < j\delta$.

Even though the degradation is discrete into a set of finite states, there is still a computational challenge caused by Bayesian online learning. With the observations up to epoch t, the posterior distributions of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are updated. However, when calculating the value function at epoch t, the number of possible shock occurrence distribution sequences and shock magnitude distribution sequences at future period $t+i$, $i = 1, ..., T-t$, are infinite because of the uncertain future realisations of shock occurrence and magnitude.

In order to conquer the computational issue, the most likely distribution is leveraged to simplify the Bayesian online updating process (Azoury and Miyaoka, 2009). In general, given the information up to epoch t, the most likely probability mass function of shock occurrences in future period $t+i$ is $K_{t+i} \sim NegBin(e^{\mathbf{s}_{t+i}\tilde{\boldsymbol{\alpha}}_t}, \mathbf{s}_{t+i}\Sigma_{\alpha_t}\mathbf{s}_{t+i}^{\intercal})$, and the most likely probability density function of shock magnitude in future period $t+i, i = 1, ..., T-t+1$ is $X_{t+i} \sim \mathcal{N}\left(\mathbf{s}_{t+i}\tilde{\boldsymbol{\beta}}_t, \sigma^2(1+\mathbf{s}_{t+i}\Sigma_{\alpha_t}\mathbf{s}_{t+i}^{\intercal})\right)$. Hence, given the observations up to epoch t, the most likely predictive distribution of degradation increment Y_{t+i} at epoch $t+i$ is $\hat{f}_{t+i,t}(Y_{t+i})$, and the corresponding cdf is defined as $\hat{F}_{t+i,t}(Y_{t+i})$. The most likely predictive distribution offers the simplification of the state space by keeping parameter $\mathbf{P}, \mathbf{Q}, \mathbf{N}$ for future period $t+i, i = 1, ..., T-t+1$ constant, aligning them with the values for the period t. Thus the state space collapses into one dimension with the accumulative degradation z_t. Let $\hat{V}_{t,t,\boldsymbol{\chi}_t}(z_t)$ denote the minimum expectation of discounted maintenance cost from epoch t to T, given the updated parameter up to epoch t. The optimality equation applying the most likely

distribution and discretisation of degradation is as follows:

$$\hat{V}_{t,t,\boldsymbol{\chi}_t}(z_t) = \begin{cases} \min\left\{c_p + \hat{V}_{t,t,\boldsymbol{\chi}_t}(0), \gamma \sum_{z_{t+1}\in\boldsymbol{\Delta}} \hat{F}_{t+1,t}(z_{t+1}-z_t)\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_{t+1})\right\}, & z_t < D, \\ c_r + \hat{V}_{t,t,\boldsymbol{\chi}_t}(0), & z_t \geq D. \end{cases} \tag{5.5}$$

The terminal function is $\hat{V}_{T+1,t,\boldsymbol{\chi}_{T+1}}(\cdot) = 0$. Details of the backward dynamic programming incorporating the most likely distribution are presented in Appendix B.8, which can return the optimal total discounted maintenance cost and maintenance actions at each epoch.

When incorporating the most likely distribution heuristic algorithm, the maintenance policy is still a control optimal policy. It can further reduce the computation burden. Based on Equation (5.5) and the stochastic order of the degradation increment in Proposition 5.3, the monotonicity property for the PM threshold is established when incorporating the most likely distribution heuristic algorithm in Proposition 5.5. From Proposition 5.3, it comes down to restrictions to the working conditions to hold the monotonicity of random variable Y_{t+1}, which implies the monotonicity of the PM threshold at epoch $t+1$.

Proposition 5.5 *For any given* $\boldsymbol{\chi}_t$, *let* $a_{t+1} = (\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal})^{-1}$, $a_{t+1}^{'} = (\mathbf{s}_{t+1}\Sigma_{\alpha_t}^{\prime}\mathbf{s}_{t+1}^{\intercal})^{-1}$, $b_{t+1} = \mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal}e^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}$ *and* $b_{t+1}^{\prime} = \mathbf{s}_{t+1}\Sigma_{\alpha_t}^{\prime}\mathbf{s}_{t+1}^{\intercal}e^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t^{\prime}}$, $c_{t+1} = \mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t$ *and* $c_{t+1}^{\prime} = \mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t^{\prime}$. *If* $a_{t+1} \leq a_{t+1}^{\prime}$, $b_{t+1} \leq b_{t+1}^{\prime}$ *and* $c_{t+1} \leq c_{t+1}^{\prime}$. *when implementing the most likely distribution heuristic algorithm, the PM threshold* $\hat{\zeta}_{t,\boldsymbol{\chi}_t}(\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t) \geq \hat{\zeta}_{t,\boldsymbol{\chi}_t}(\mathbf{P}_t^{\prime}, \mathbf{Q}_t^{\prime}, \mathbf{N}_t^{\prime})$.

5.4.2 Special case: Regression through the origin

If the relationship between the covariates and the dependent variable assumes an intercept of zero, the regression line is forced through the origin. Regression through the origin is rarely discussed, as imposing this constraint—by forcing the regression line to pass through zero—may reduce the model's flexibility and potentially weaken its ability to fit the data accurately. Nevertheless, there exist reasons in many cases to argue that the dependent variable equals zero when the covariates are at zero. In instances where the operating speed of the production machine is zero, it is appropriate to employ re-

gression analysis through the origin. This is justified by the observation that during machine shutdown periods, the absence of activity indicates no shock arrivals.

When the regression is through the origin and only one covariate is considered, for the Bayesian Poisson regression that applied to the number of shocks, the parameter $\boldsymbol{\alpha}$ collapses into slope α, which is a single parameter when the regression through the origin with one covariate s. The prior on α is Normal distributed with prior mean $\tilde{\alpha}_0$ and variance σ_{α_0}. At epoch t, the posterior of α is Normal distributed with mean $\tilde{\alpha}_t = (P_t + \sigma_{\alpha_0}^{-2})^{-1}(Q_t + \sigma_{\alpha_0}^{-2}\tilde{\alpha}_0)$, and variance $\sigma_{\alpha_t}^2 = (P_t + \sigma_{\alpha_0}^{-2})^{-1}$, where $P_t = \sum_{i=1}^{t} \frac{s_i^2}{k_i}$, $Q_t = \sum_{i=1}^{t} \frac{s_i \log k_i}{k_i}$. In a similar manner, for the Bayesian linear regression, the parameter $\boldsymbol{\beta}$ collapses into slope β, which is a single parameter when the regression through the origin with one covariate s. The prior on β is Normal distributed with prior mean $\tilde{\beta}_0$ and variance σ_{β_0}. At epoch t, the posterior of β is Normal distributed with mean $\tilde{\beta}_t = M_t^{-1}(N_t + \sigma^2\sigma_{\beta_0}^{-2}\tilde{\beta}_0)$ and variance $\sigma_{\beta_t} = \sigma^2 M_t^{-1}$, where $M_t = \sigma^2\sigma_{\beta_0}^{-2} + \sum_{i=1}^{t} k_i s_i^2$ and $N_t = \sum_{i=1}^{t}(s_i \sum_{j=1}^{k_t} x_{ij})$.

In addition, for the special case that the regression is forced through the original with one covariate, the following Corollary 5.1 can be extended from Proposition 5.5 when applying the heuristic algorithm.

Corollary 5.1 *For any given $\boldsymbol{\chi}_t$, if the regression lines are forced through the origin with one covariate s, let $a_{t+1} = (s_{t+1}\sigma_{\alpha_t})^{-1}$, $a'_{t+1} = (s_{t+1}\sigma'_{\alpha_t})^{-1}$, $b_{t+1} = s_{t+1}\sigma_{\alpha_t}e^{s_{t+1}\tilde{\alpha}}$ and $b'_{t+1} = s_{t+1}\sigma'_{\alpha_t}e^{s_{t+1}\tilde{\alpha}'}$, $c_{t+1} = s_{t+1}\tilde{\beta}_t$ and $c'_{t+1} = s_{t+1}\tilde{\beta}'_t$. If $a_{t+1} \leq a'_{t+1}$, $b_{t+1} \leq b'_{t+1}$ and $c_{t+1} \leq c'_{t+1}$, when implementing the most likely distribution heuristic algorithm, the PM threshold $\hat{\zeta}_{t,\boldsymbol{\chi}_t}(P_t, Q_t, N_t) \geq \hat{\zeta}_{t,\boldsymbol{\chi}_t}(P'_t, Q'_t, N'_t)$.*

5.5 Numerical studies

5.5.1 Experiment setting

In numerical studies, it is assumed that the underlying parameters are unknown to the decision maker, but he/she can assess the prior distribution of the heterogeneous components. For illustrative purposes, let us consider one working condition, and set up the underlying coefficient vector $\boldsymbol{\alpha} = [0.7, 1.0]^\intercal$ and the coefficient vector $\boldsymbol{\beta} = [0.1, 0.2]^\intercal$.

The prior distribution of $\boldsymbol{\alpha}$ is assumed to follow a multivariate Normal distribution with mean vector $\boldsymbol{\alpha}_0 = [0.1, 0.5]^\intercal$ and covariance matrix $\Sigma_{\alpha_0} = \text{diag}(0.2, 0.2)$. The prior distribution of $\boldsymbol{\beta}$ is assumed to follow a multivariate Normal distribution with mean vector $\boldsymbol{\beta}_0 = [0.01, 0.1]^\intercal$ and covariance matrix $\Sigma_{\beta_0} = \text{diag}(0.2, 0.2)$. The variance of the Bayesian Linear regression $\sigma^2 = 0.01$. The working conditions vary along with the epoch and are assumed as constant within one period, which is in a piecewise form illustrated in Figure 5.6. Given the underlying parameters, the degradation paths are randomly generated with dynamic working conditions.

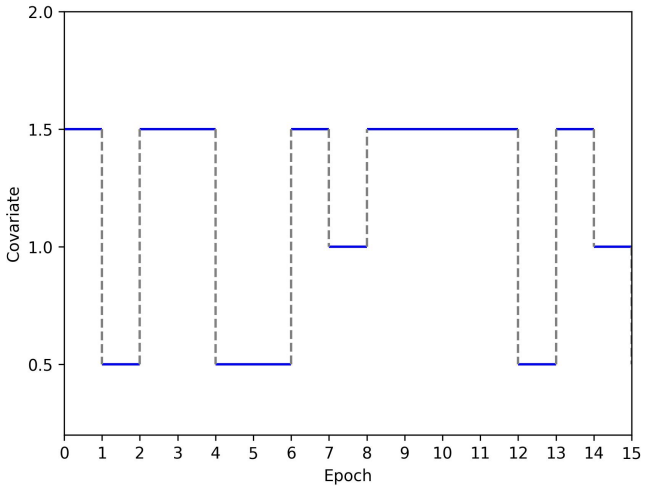

Figure 5.6: Illustration of piecewise working conditions across epochs

The accumulative degradation level of a component is discretised into 30 discrete states. A component is assumed to fail when its accumulative degradation level $Z \geq 30$. CM is carried out when a component is found to fail at the decision epoch, and PM is performed if its accumulative degradation level exceeds the PM threshold. The costs of CM and PM are given as $c_r = 400$ and $c_p = 100$. The discount parameter is set as $\gamma = 0.9$. The component is assumed to operate from epoch 0 to epoch 15 and will be discarded without any cost after the last epoch.

5.5.2 Results

First, the process of Bayesian updating upon the arrival of new observations is shown, where the underlying parameter is given in the experiment settings. Since one covariate

is considered in the numerical study, the coefficient vector $\boldsymbol{\alpha}$ comprises the intercept and coefficient associated with the covariate. Figure 5.7 displays the contour plot that illustrates the updating process of $\boldsymbol{\alpha}$ combining with the observations of the number of shocks. Figure 5.7a shows that the underlying coefficient and intercept in $\boldsymbol{\alpha}$ is far away from the prior distribution when the observation of shocks is not provided. With the increase in observations, the estimation of $\boldsymbol{\alpha}$ is becoming more centralised and approaching the true parameter. In addition, Figure 5.7 illustrates that Normal approximation can effectively be applied to estimate the parameters in the Bayesian Poisson regression.

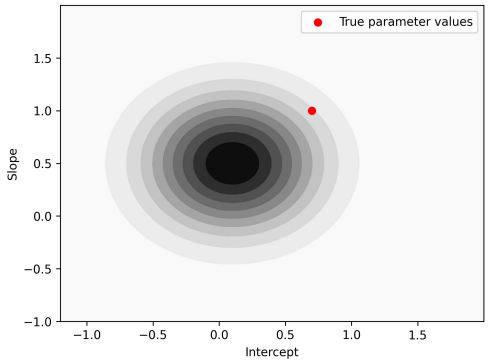

(a) Bayesian updating of $\boldsymbol{\alpha}$ using 0 observation

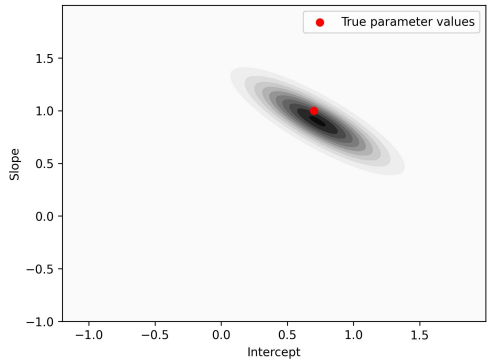

(b) Bayesian updating of $\boldsymbol{\alpha}$ using 5 observations

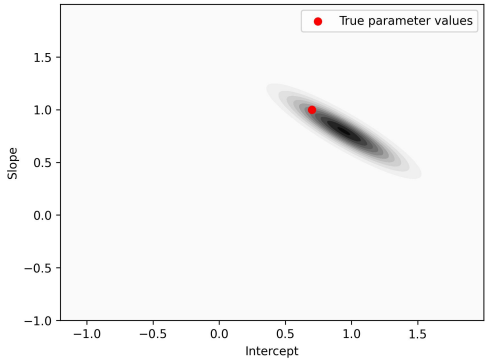

(c) Bayesian updating of $\boldsymbol{\alpha}$ using 10 observations

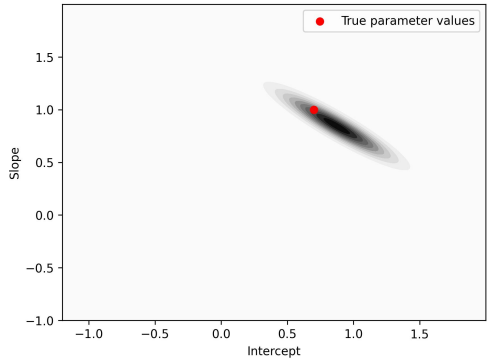

(d) Bayesian updating of $\boldsymbol{\alpha}$ using 15 observations

Figure 5.7: Bayesian updating of $\boldsymbol{\alpha}$ with increasing number of observations

Similarly, the updating process of $\boldsymbol{\beta}$ is obtained with 0, 5, 10 and 15 observations, as indicated in the contour plots of Figure 5.8. The posterior distribution of $\boldsymbol{\beta}$ is becoming

narrow as the increasing of observations. When 15 observations of the magnitude of shocks are considered, the posterior distribution converges close to the true parameters.

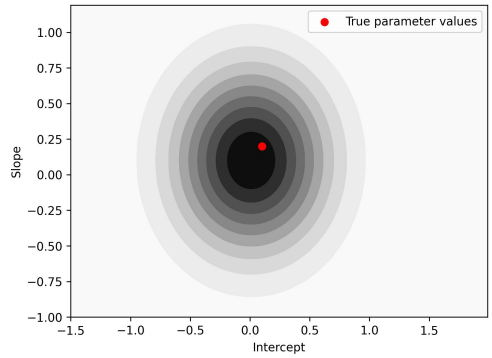

(a) Bayesian updating of $\boldsymbol{\beta}$ using 0 observation

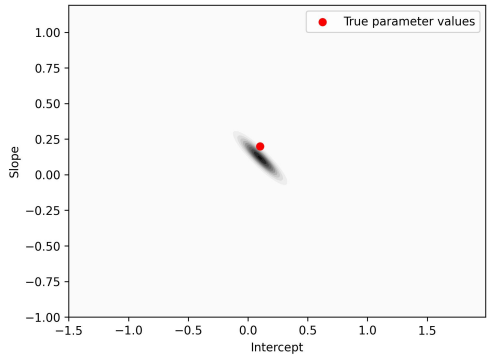

(b) Bayesian updating of $\boldsymbol{\beta}$ using 5 observations

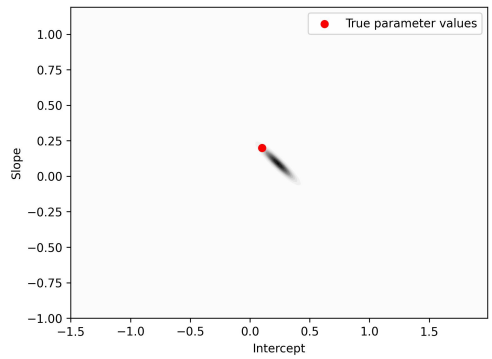

(c) Bayesian updating of $\boldsymbol{\beta}$ using 10 observations

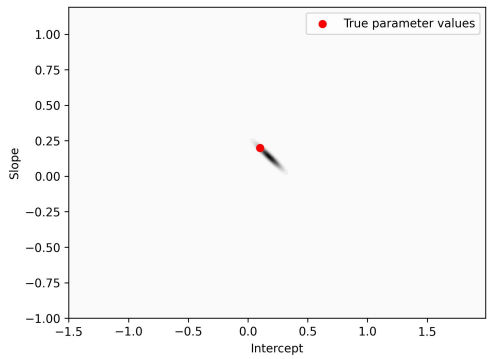

(d) Bayesian updating of $\boldsymbol{\beta}$ using 15 observations

Figure 5.8: Bayesian updating of $\boldsymbol{\beta}$ with increasing number of observations

Furthermore, the PM threshold at each epoch is presented in Figure 5.9, which is obtained by integrating Bayesian updating and dynamic programming while considering the dynamic working conditions. Combining the dynamic working conditions displayed in Figure 5.6, it can be observed that the PM threshold fluctuates along with the value of the covariate. Specifically, the PM threshold becomes lower when the value of the covariate becomes larger. This suggests that the system's sensitivity to degradation signals varies with specific operating conditions (e.g., high load, harsh environments, or elevated temperatures). Consequently, maintenance decisions should be dynamically adjusted to respond effectively to changing working conditions.

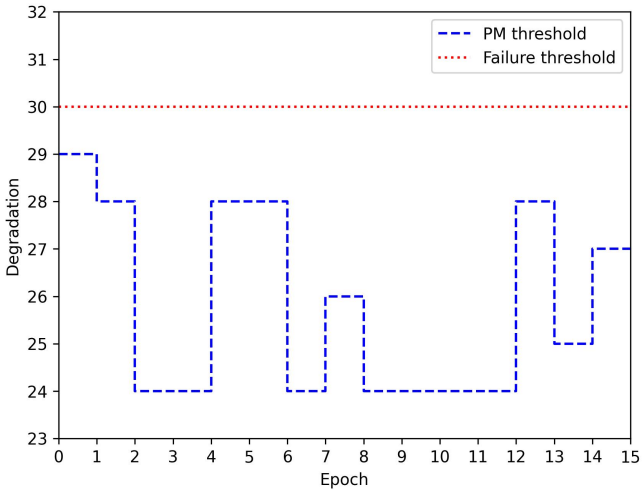

Figure 5.9: PM threshold at each epoch considering varied working conditions

5.5.3 Comparison experiments

To illustrate the advantages of the proposed Bayesian updating method considering varied working conditions, the proposed method is compared with two methods and the optimal policy given the underlying parameters. 200 degradation paths that are subject to the CP process are randomly generated using the underlying parameters of $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ and the covariate in the aforementioned settings. The average maintenance costs for the 200 degradation paths under the PM threshold given by each method are reported. It is assumed that the same component will replace the component with the same parameters of $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ when PM or CM is performed. The methods are discussed as the followings.

The first method is called the offline approach. It ignores the online Bayesian updating and only leverages the prior information to make the maintenance decisions. Compared with the integrated method proposed by this work, the offline approach directly uses the prior distribution of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ for all the epochs. The optimisation stage still solves the Bellman equation in 5.4 to minimise the expected maintenance cost over the finite horizon.

The second method is called the myopic approach. The updating of the posterior distribution of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ is considered when a new observation is available. But in the

optimisation stage, the maintenance decision is obtained by minimising the maintenance cost up to the future one epoch, instead of looking ahead of the whole decision horizon. Because of the computation efficiency, the myopic method is popular in stochastic optimisation literature (Bertsekas, 2012; Levi et al., 2007).

In Figure 5.10, the average maintenance costs under different PM thresholds are given by the offline approach, myopic approach, the proposed method in this work and the optimal policy. It can be observed that the gap between the Bayesian method and the optimal policy is the smallest among the three methods over 200 simulations. Compared with the offline approach and myopic approach, the proposed method shows better performance since the variance of the parameters is reduced with the new observations and it leverages the posterior distribution into the long-run optimisation.

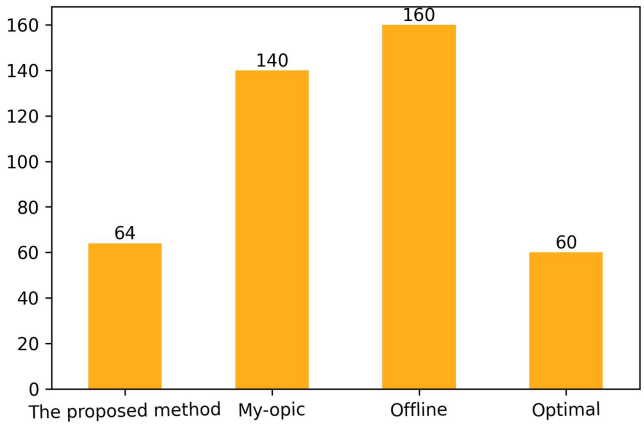

Figure 5.10: Maintenance cost comparison across the proposed new method, offline approach, myopic approach, and the optimal policy

5.6 Illustrative example

To illustrate the application of the proposed method, this section applies the experimental data related to ultra-thin gate oxide in metal-oxide-semiconductor devices (Cester et al., 2003). Ultra-thin gate oxide which serves to separate the gate electrode from the semiconductor and modulate its conductivity. The thinning of gate oxide is important to maintain adequate drive current for proper circuit operation and achieve higher performance in microelectronics (Hou et al., 2002). The term 'ultra-thin' indicates the

gate oxide layer as thin as a few nanometers or even less. Thinner gate oxide benefits in accelerating the switch speeds and reducing power consumption. However, the scaling of the gate oxide is limited by the gate leakage current, which causes unacceptable standby power consumption. The maximum acceptable gate leakage current has been suggested between 1 A/cm^2 and 10 A/cm^2, which is associated with the oxide thickness (Suehle, 2002).

The experiment conducted by Cester et al. (2003) observed a discrete degradation pattern in the leakage current of the ultra-thin gate oxide. The accelerated degradation experiments were performed on 3nm oxides under varying conditions of constant voltage stress V_{CVS} and ion irradiation ϕ. Figure 5.11 reproduces the experimental data by extracting it from the original experiment plots, with the x-axis of Figure 5.11 representing the time of the accelerated experiment, and the y-axis showing the leakage current density. It is worth noting that the constant voltage stress and iron irradiation are two main factors related to the leakage of ultra-thin gate oxide.

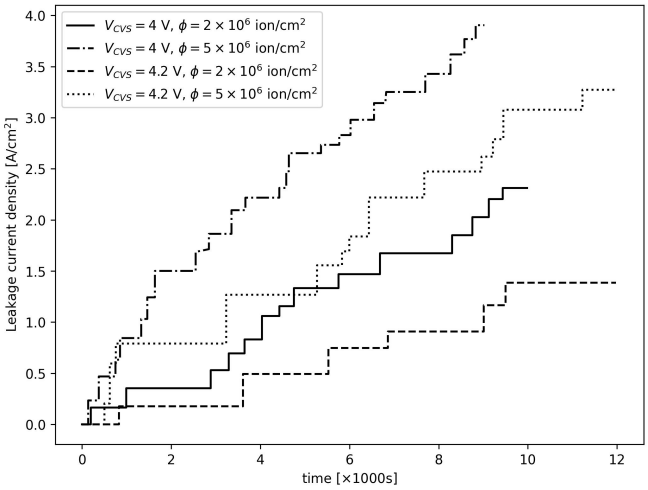

Figure 5.11: Replot of leakage current density measured under varying conditions of constant voltage stress and ion irradiation, based on the experiments from Cester et al. (2003)

In the experimental data, the covariate vector for epoch t can be represented as $s_t = [1, V_{cvs}, \phi]$, and every $1000s$ is regarded as a period. Based on the reproduced data set, the prior values of the parameters are estimated using the MLE method as described

in Appendix B.6, i.e., $\tilde{\boldsymbol{\alpha}}_0 = [1.976, -1.054, 0.667]^\intercal$, $\tilde{\boldsymbol{\beta}}_0 = [0.075, 0.009, 0.077]^\intercal$,

$$\Sigma_{\alpha_0} = \begin{bmatrix} 0.312 & -0.087 & -0.109 \\ -0.087 & 0.080 & -0.013 \\ -0.109 & -0.013 & 0.078 \end{bmatrix}, \Sigma_{\beta_0} = \begin{bmatrix} 0.001 & -0.005 & -0.004 \\ -0.005 & 0.000 & 0.000 \\ -0.004 & 0.000 & 0.000 \end{bmatrix}.$$

The prior estimates indicate that only the constant voltage stress is negatively correlated to the number of shocks. For the online estimating process with a specific ultra-thin gate oxide, the underlying coefficients are assumed as $\boldsymbol{\alpha} = [1.6, -1, 1]^\intercal$ and $\boldsymbol{\beta} = [0.05, 0.01, 0.1]^\intercal$. The illustrative example in this section is based on a limited set of experimental data, where heterogeneity is not clearly observable due to the small sample size. To reflect potential heterogeneity effects, deliberate differences between the underlying parameters and the MLE results are introduced (see the differences between $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\tilde{\boldsymbol{\alpha}}_0, \tilde{\boldsymbol{\beta}}_0)$). This setting is intended to illustrate the applicability of the proposed framework to cases involving heterogeneity, rather than to represent natural statistical estimation errors. The primary aim of this Section is to demonstrate the applicability of the proposed methodology.

The settings of constant voltage stress and ion irradiation are given in Table 5.1 for 12 periods. In Table 5.1, the operating condition 1 represents $V_{cvs} = 4V$ and 2 represents $V_{cvs} = 4.2$ V in the row of constant voltage stress, 1 represents $\phi = 2 \times 10^6$ ions/cm^2 and 2 represents $\phi = 5 \times 10^6$ ions/cm^2 in the row of ion irradiation.

Table 5.1: Settings of constant voltage stress and ion irradiation

Period	1	2	3	4	5	6	7	8	9	10	11	12
Constant voltage stress	2	2	2	1	2	2	1	2	2	2	1	1
Ion irradiation	2	1	1	1	1	2	2	2	1	2	2	2

In the row of constant voltage stress, 1 represents $V_{cvs} = 4V$ and 2 represents $V_{cvs} = 4.2$ V;
In the row of ion irradiation, 1 represents $\phi = 2 \times 10^6$ ions/cm^2 and 2 represents $\phi = 5 \times 10^6$ ions/cm^2.

At epoch 12, the parameters are updated as $\tilde{\boldsymbol{\alpha}}_{12} = [1.617, -0.904, 0.861]^\intercal$ and $\tilde{\boldsymbol{\beta}}_{12} = [0.005, 0.016, 0.094]^\intercal$, which are close to the underlying $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. The maintenance policy is illustrated in Figure 5.12, where the black solid line represents a randomly generated degradation path using the underlying parameters and the working condition

setting in Table 5.1, and the red dotted line indicates the failure threshold (10 A/cm^2). In Figure 5.12a, the leakage current density of the degradation path exceeds the PM threshold at epoch 9, while it exceeds the failure threshold at epoch 10. In addition, the blue dashed line in Figure 5.12a represents the PM threshold determined by the proposed method. In Figure 5.12b, the green dashed line represents the optimal PM threshold, determined based on the assumption of known underlying parameters. The result presents that the obtained PM threshold is close to the underlying optimal one, indicating the effectiveness of the proposed approach.

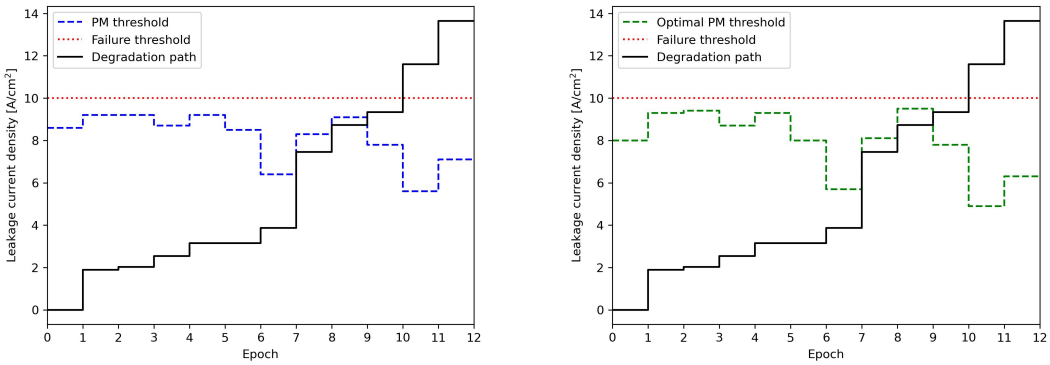

(a) PM thresholds determined by the proposed method

(b) Optimal PM thresholds determined with underlying parameters

Figure 5.12: PM thresholds of the proposed method and the optimal PM thresholds determined based on the underlying true parameters

5.7 Summary

In this Chapter, the CBM policy is investigated and applied to the component degradation following the CP process. The approach integrates both offline estimation and online updating strategies, enabling the utilisation of real-time data for capturing the dynamic behaviour of degradation. Bayesian Linear regression and Bayesian Poisson regression are employed to effectively incorporate dynamic working conditions and latent heterogeneity into the model. An analytically tractable degradation model is achieved by leveraging conjugate priors to model the heterogeneity of the component population. The optimal PM threshold is characterised in the form of a control limit policy.

Notably, our findings challenge the classic monotonicity assumption, revealing that the PM threshold is influenced by varied working conditions. The most likely heuristic and backward MDP algorithms are combined to reduce the computation time. Numerical experiments and an illustrative example are conducted to show the advancements of the developed CBM policy considering both working conditions and heterogeneity. It is important to note that the results presented are specific to degradation processes driven by random shocks. While the modelling assumptions—such as employing a Poisson process for shock occurrence and assuming normally distributed degradation magnitudes—enable controlled simulation and analysis, they may oversimplify the complexity and context-dependency of real-world scenarios. Maintenance policies for other types of degradation mechanisms could be explored by extending the framework developed in this work.

There are several directions for potential future extensions. One theoretical direction involves a more in-depth analysis of the convergence of the Bayesian updating process. This exploration will establish the error bounds between online decisions and the optimal PM threshold. Building upon existing research on the convergence rate of posterior distributions may provide valuable theoretical insights for such extension. Another potential direction is to consider systems consisting of multiple components with dependent degradation. The consideration of incorporating factors, such as varied working conditions and latent heterogeneity, in such complex systems will lead to more comprehensive and realistic maintenance policies. The degradation model developed in this Chapter is well-suited for multiple components with independent degradation processes, while future investigations will focus on dependent degradation with interactions among components.

Chapter 6

Conclusion

This thesis has investigated two key tasks in smart asset management: degradation prediction and maintenance scheduling using data-driven approaches. This chapter summarises the main findings across Chapters 3, 4, and 5, highlighting the commonalities and differences among the developed methods, and proposes potential directions for future research.

6.1 Thesis summary

This thesis addresses the challenges of degradation prediction and maintenance optimisation under varying data availability and operational conditions, and contributes a suite of methodologies within the data-driven paradigm.

Chapter 3 developed a DL framework for degradation prediction under limited observation conditions. DA methods—including jittering, scaling, window warping, and others—were employed to expand the data set. Influence functions were used to identify and filter out negative-impact synthetic samples, while a gradient-based reweighting scheme was introduced to further enhance training efficiency. To improve generalisation, a TL strategy was applied by fine-tuning the model on real samples. Numerical experiments confirmed the proposed framework's performance over baseline methods. Although these gains were moderate in the illustrative examples, the framework's effectiveness remains valuable, especially in domains where degradation data are non-stationary.

Chapter 4 proposed an E2E data-driven framework for PM decision-making under varied working conditions. Departing from the ETO approach, the E2E framework directly maped operational data to maintenance decisions, bypassing explicit model specification and parameter estimation. Theoretical results were established to measure the quality of the learned maintenance cost rate functions, with error bounds analytically derived. Among the compared methods, E2E_LASSO consistently delivered the best results—likely due to its robustness and sparsity, which make it well-suited to limited-data scenarios. In contrast, E2E_ANN showed slightly lower performance, potentially reflecting its greater data requirements. These findings highlight several insights: Direct policy learning offers a more effective alternative to prediction-based optimisation; A simple linear model, when aligned with the learning objective, can be effective and may even outperform more complex architectures; The E2E framework is adaptable to different system types, offering value for a wide range of maintenance applications.

Chapter 5 extended the investigation to CBM policies under dynamic working environments, focusing on degradation driven by shock processes modelled via a CP framework. Bayesian linear and Poisson regressions were employed for online updating, effectively capturing both working condition dynamics and latent heterogeneity across systems. A control-limit PM policy was characterised, and a most likely heuristic algorithm combined with a backward MDP was developed to reduce computational complexity. By examining the dynamic working conditions, it was observed that the PM threshold fluctuates in accordance with changes in the covariate. This behaviour suggests that maintenance decisions may benefit from being dynamically adjusted in response to changing working conditions. In such settings, static policies could be less effective, while the proposed adaptive framework provides an approach to account for variability in real time.

Across these three chapters, the requirements for data and assumptions about system behaviours varied: Chapter 3 focused on scenarios with limited degradation trajectory data, enhancing predictive models through DA and TL. Chapter 4 assumed the availability of time-to-failure data under relatively constant working conditions,

enabling direct supervised learning for TBM strategies. Chapter 5 addressed more complex settings with online degradation monitoring under time-varying conditions, integrating statistical updating and dynamic optimisation. Moreover, different perspectives of maintenance were explored: Chapter 4 emphasised TBM decisions under constant conditions. Chapter 5 developed CBM strategies accounting for real-time environmental variations and system heterogeneity.

The methodologies developed across Chapters 3 to 5 collectively contribute to building smart maintenance systems that can adapt to various operational scenarios, and address key aspects such as predictive accuracy, maintenance decision making under uncertainty.

6.2 Future research

In the summary sections of Chapters 3, 4, and 5, the potential future extensions specific to each chapter have been discussed. The final section will outline three broader avenues for future research that extend beyond the scope of these individual chapters.

Develop an explainable DL algorithm for degradation prediction.

Despite the success of DL in degradation prediction, it is often perceived as a "black box", making it less understandable and limiting its use in safety-critical applications. Typically, DL methods lack the ability to interpret learned representations, explain generated decisions, or quantify the reliability of those decisions. Future work on explainable DL for degradation prediction could focus on several key areas: analysing how input features contribute to the performance of DL networks by integrating prior knowledge of degradation mechanisms, evaluating the importance of training samples to refine data selection, and assessing model uncertainty to better support maintenance decisions.

Data-driven spare part management integrating with CBM policy.

The maintenance policy naturally interacts with spare part management, directly influencing the demand for spare parts. A key challenge in this integration is that spare parts demand is often intermittent and irregular, with extended periods of zero demand. However, by leveraging historical degradation data, it is possible to develop a data-driven approach, such as kernel-based methods, to model the relationship between preventive maintenance thresholds and spare parts demand distribution. This approach could further contribute to determining optimal maintenance thresholds and spare part inventory decisions.

Data-driven maintenance policy under the deteriorating sensor.

The CBM policy utilises sensors and communication systems to monitor equipment conditions and predict system failures based on collected measurements. In the literature, it is often assumed that sensor readings are either perfect or only subject to measurement errors. However, this assumption overlooks the fact that sensors are also prone to degradation over time due to environmental and ageing factors. To address this, a hidden Markov model could be applied to capture the dynamics between hidden states, i.e., sensor degradation and observable data. Subsequently, the maintenance problem could be modelled as a POMDP over an infinite horizon. This extension not only improves maintenance efficiency but also enhances the reliability of the monitoring system by incorporating sensor health into the maintenance scheduling process.

Appendix A

Proofs for Chapter 4

Fact 1 *(Hoeffding's inequality). Let $Z_1, ..., Z_n$ be independent bounded random variables with $Z_i \in [a, b]$ for all i, where $-\infty < a \leq b < \infty$. Let*

$$\bar{Z} = \frac{Z_1 + Z_2 + ... + Z_n}{n}$$

Then,

$$\mathbb{P}(\left|\bar{Z} - \mathbb{E}(Z)\right| \geq \epsilon) \leq 2\exp\left(\frac{-2n^2\epsilon^2}{\sum_{i=1}^{n}(b_i - a_i)^2}\right)$$

Fact 2 *(Massart's Inequality) Suppose real-valued variables $X_1, ..., X_n$ are independent and identically distributed with cumulative distribution function $F(\cdot)$. Let $\hat{F}(\cdot)$ denote the associated empirical distribution function denoted by*

$$\hat{F}(x) = \frac{1}{n}\sum_{i=1}^{n} 1[X_i \leq x], x \in \mathbb{R}$$

Then for all $\alpha \geq 0$,

$$\mathbb{P}\left[\sup_{x\in\mathbb{R}}\left|\hat{F}(x) - F(x)\right| \geq \alpha\right] \leq 2\exp^{-2n\alpha^2}$$

A.1 Proof of Proposition 4.1

Proof *For the labelling process, assume $N_i1(t), ..., N_ik(t)$ are independent and identically distributed random variables bounded by $[N_{min}, N_{max}]$. According to Hoeffding's*

inequality, with probability at least $1-2\exp\frac{-2k\omega^2}{(N_{max}-N_{min})^2}$ *then*

$$\left|\bar{N}(t)-\mathbb{E}(N(t))\right|\leq\omega.$$

Then the empirical cost rate for repairable system converges as

$$\left|C_{RS}(t;\mathbf{x}_i)-\hat{C}_{RS}(t;\mathbf{x}_i)\right|=\left|\frac{c_m(\mathbb{E}(N(t)-\bar{N}(t))}{t}\right|\leq\frac{c_m\omega}{t}$$

A.2 Proof of Proposition 4.2

Proof *According to Massart's inequality (Fact 2), for any* $t\in[T_{min}^{NRS},T_{max}^{NRS}]$ *with probability at least* $1-2\exp^{-2m\varepsilon^2}$ *then*

$$\left|\hat{F}(t;\mathbf{s}_i)-F(t;\mathbf{s}_i)\right|\leq\varepsilon \tag{A.1}$$

$$\left|\hat{R}(t;\mathbf{s}_i)-R(t;\mathbf{s}_i)\right|\leq\varepsilon \tag{A.2}$$

$$\left|\int_{t=0}^{t}\hat{R}(t;\mathbf{s}_i)dt-\int_{t=0}^{t}R(t;\mathbf{s}_i)dt\right|\leq\varepsilon t \tag{A.3}$$

Thus, for the absolute error of the approximated cost function for the non-repairable systems, we have

$$\begin{aligned}\left|\hat{C}_{NRS}(t;\mathbf{s}_i)-C_{NRS}(t;\mathbf{s}_i)\right|&=\left|\frac{(c_f-c_r)\hat{F}(t;\mathbf{s}_i)+c_r}{\int_{t=0}^{t}\hat{R}(t;\mathbf{s}_i)dt}-\frac{(c_f-c_r)F(t;\mathbf{s}_i)+c_r}{\int_{t=0}^{t}R(t;\mathbf{s}_i)dt}\right|\\&\leq\underbrace{\left|\frac{(c_f-c_r)\hat{F}(t;\mathbf{s}_i)}{\int_{t=0}^{t}\hat{R}(t;\mathbf{s}_i)dt}-\frac{(c_f-c_r)F(t;\mathbf{s}_i)}{\int_{t=0}^{t}R(t;\mathbf{s}_i)dt}\right|}_{①}\\&\quad+\underbrace{\left|\frac{c_r}{\int_{t=0}^{t}\hat{R}(t;\mathbf{s}_i)dt}-\frac{c_r}{\int_{t=0}^{t}R(t;\mathbf{s}_i)dt}\right|}_{②}\end{aligned} \tag{A.4}$$

Let $G(t;\mathbf{s}_i)$ *denote* $\int_{t=0}^{t}R(t;\mathbf{s}_i)dt$ *and* $\hat{G}(t;\mathbf{s}_i)$ *denote* $\int_{t=0}^{t}\hat{R}(t;\mathbf{s}_i)dt$*. From the property of CDF in reliability analysis, we know that* $\hat{F}(t;\mathbf{s}_i)$*,* $F(t;\mathbf{s}_i)$*,* $\hat{G}(t;\mathbf{s}_i)$*,* $G(t;\mathbf{s}_i)$ *are*

non-decreasing functions, thus we have

$$0 \leq \hat{F}(t;\mathbf{s}_i) \leq 1; 0 \leq F(t;\mathbf{s}_i) \leq 1; R(t;\mathbf{s}_i) \leq G(t;\mathbf{s}_i) \leq t; \hat{R}(t;\mathbf{s}_i) \leq \hat{G}(t;\mathbf{s}_i) \leq t.$$

We drop $\mathbf{s}_i$ *for notation convenience in the following functions. Combining with inequality given in A.1, A.2, A.3, we can analyse the item* ① *in Equation A.4:*

$$\begin{aligned}
\frac{(c_f - c_r)\hat{F}(t)}{\hat{G}(t)} - \frac{(c_f - c_r)F(t)}{G(t)} &= \frac{(c_f - c_r)[\hat{F}(t)G(t) - F(t)\hat{G}(t)]}{\hat{G}(t)G(t)} \\
&\leq \frac{(c_f - c_r)\left[\hat{F}(t)(\hat{G}(t) + \varepsilon t) - (\hat{F}(t) - \varepsilon)\hat{G}(t)\right]}{\hat{G}(t)G(t)} \\
&= \frac{(c_f - c_r)\left[\varepsilon t\hat{F}(t) + \varepsilon\hat{G}(t)\right]}{\hat{G}(t)G(t)} \leq \frac{2\varepsilon t(c_f - c_r)}{\hat{G}(t)G(t)}.
\end{aligned}$$

Similarly, we can also have

$$\frac{(c_f - c_r)\hat{F}(t)}{\hat{G}(t)} - \frac{(c_f - c_r)F(t)}{G(t)} \geq \frac{(c_f - c_r)\left[\hat{F}(t)(\hat{G}(t) - \varepsilon t) - (\hat{F}(t) + \varepsilon)\hat{G}(t)\right]}{\hat{G}(t)G(t)} \geq \frac{-2\varepsilon t(c_f - c_r)}{\hat{G}(t)G(t)}.$$

Thus, we need to find the lower bound of $\hat{G}(t)G(t)$*. The first derivative of* $\hat{G}(t)G(t)$ *is*

$$\hat{R}(t)G(t) + \hat{G}(t)R(t) \geq 0,$$

therefore $\hat{G}(t)G(t)$ *is non-decreasing over* t *for any* $t \in [T_{min}^{NRS}, T_{max}^{NRS}]$*, and*

$$\begin{aligned}
\frac{1}{\hat{G}(t)G(t)} \leq \frac{1}{\hat{G}(T_{min}^{NRS})G(T_{min}^{NRS})} &\leq \frac{1}{\hat{G}^2(T_{min}^{NRS}) - \varepsilon T_{min}^{NRS}\hat{G}(T_{min}^{NRS})} \\
&\leq \frac{1}{-\varepsilon(T_{min}^{NRS})^2 + (\hat{R}(T_{min}^{NRS})T_{min}^{NRS})^2}
\end{aligned}$$

when the value of ε *satisfied* $\varepsilon < \hat{R}^2(T_{min}^{NRS})$*.*

Thus we can obtain ① *in Equation A.4:*

$$\left|\frac{(c_f - c_r)\hat{F}(t;\mathbf{s}_i)}{\int_{t=0}^{t}\hat{R}(t;\mathbf{s}_i)dt} - \frac{(c_f - c_r)F(t;\mathbf{s}_i)}{\int_{t=0}^{t}R(t;\mathbf{s}_i)dt}\right| \leq \frac{2\varepsilon t(c_f - c_r)}{(\hat{R}^2(T_{min}^{NRS}) - \varepsilon)(T_{min}^{NRS})^2}.$$

By a similar argument, we also have item ② in Equation A.4:

$$\left| \frac{c_r}{\int_{t=0}^{t} \hat{R}(t;\mathbf{s}_i)dt} - \frac{c_r}{\int_{t=0}^{t} R(t;\mathbf{s}_i)dt} \right| \leq \frac{\varepsilon t c_r}{(\hat{R}^2(T_{min}^{NRS}) - \varepsilon)(T_{min}^{NRS})^2}.$$

So we establish the bound

$$\left| \hat{C}_{NRS}(t;\mathbf{s}_i) - C_{NRS}(t;\mathbf{s}_i) \right| \leq \frac{(2c_f - c_r)\varepsilon t}{(\hat{R}^2(T_{min}^{NRS};\mathbf{s}_i) - \varepsilon)(T_{min}^{NRS})^2} \tag{A.5}$$

with the probability of at least $1 - 2\exp^{-2Q_i\varepsilon^2}$*, where* $0 \leq \varepsilon < \hat{R}^2(T_{min}^{NRS};\mathbf{s}_i)$ *is the accuracy parameter.*

Appendix B

Proofs for Chapter 5

B.1 Proof for Proposition 5.1

Proof *We prove 1 and 2 in Proposition 5.1 respectively.*

1. *At epoch t, we can approximate the posterior distribution of $\boldsymbol{\alpha}$ using normal distribution according to Lemma 5.1. Define an intermediate random value $m = \boldsymbol{s}\boldsymbol{\alpha}$, since $\boldsymbol{\alpha}_{t+1} \sim \mathcal{N}(\tilde{\boldsymbol{\alpha}}_t, \Sigma_{\alpha_t})$, hence $m_{t+1} \sim \mathcal{N}(\boldsymbol{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t, \boldsymbol{s}_{t+1}\Sigma_{\alpha_t}\boldsymbol{s}_{t+1}^\intercal)$.*

   It is well known that the Gaussian distribution is approximately close to the log-gamma distribution. Consider a random variable $n \sim Gamma(a, b)$, then the transformed random variable $\log n$ follows the log-gamma distribution. For large a, the log-gamma distribution is close to the normal distribution, that is $\log n \sim \mathcal{N}(\log a + \log b, a^{-1})$. Hence, $e^{m_{t+1}} \sim Gamma(a, b)$, where $a = (\boldsymbol{s}_{t+1}\Sigma_{\alpha_t}\boldsymbol{s}_{t+1}^\intercal)^{-1}$ and $b = \boldsymbol{s}_{t+1}\Sigma_{\alpha_t}\boldsymbol{s}_{t+1}^\intercal e^{\boldsymbol{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}$.

   Now, we can derive the predictive distribution of k_{t+1} by marginalising the distribution of k given $e^{m_{t+1}}$ over the posterior distribution of $e^{m_{t+1}}$ given a, b

$$\begin{aligned} f(k_{t+1}|a,b) &= \int f(k_{t+1}|e^{m_{t+1}}) f(e^{m_{t+1}}|a,b) de^{m_{t+1}} \\ &= \frac{\Gamma(a+k_{t+1})}{\Gamma(k_{t+1}+1)\Gamma(a)} \hat{p}^a (1-\hat{p})^{k_{t+1}}, \end{aligned}$$

   where $\hat{p} = \frac{1}{1+b}$. Note that the negative binomial distribution is viewed as a Poisson distribution with a gamma prior on the rate parameter. Therefore, we can

conclude that the predictive distribution of the number of shocks follows a negative distribution as

$$k_{t+1}|\boldsymbol{s}_{t+1}, \tilde{\boldsymbol{\alpha}}_t, \Sigma_{\alpha_t} \sim \text{NegBin}(\mathrm{e}^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}, \mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal}).$$

2. *At epoch t, the posterior distribution of $\boldsymbol{\beta}$ follow multivariate normal distribution with $p+1$-dimensional vector according to Lemma 5.2. By marginalising the distribution of x, the predictive distribution is*

$$\begin{aligned} f(x_{t+1}|s_{t+1}, \tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t}, \sigma^2) &= \int f(x_{t+1}|s_{t+1}\boldsymbol{\beta}, \sigma^2) f(\boldsymbol{\beta}|\tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t}) d\boldsymbol{\beta} \\ &= \int \mathcal{N}(x_{t+1}|s_{t+1}\boldsymbol{\beta}, \sigma^2) \mathcal{N}_{p+1}(\boldsymbol{\beta}|\tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t}) d\boldsymbol{\beta} \\ &= \mathcal{N}\left(x_{t+1}|s_{t+1}\tilde{\boldsymbol{\beta}}_t, \sigma^2(1 + s_{t+1}\Sigma_{\beta_t}s_{t+1}^{\intercal})\right). \end{aligned}$$

This follows the linear combination property of multivariate normal distribution, see Theorem 3.1 in (Gut, 2009).

B.2 Proof for Proposition 5.2

Proof *We first discuss the predictive distribution of degradation increment given the number of shocks $K_{t+1} = k$ in period $t+1$. The predictive distribution of x_{t+1} follows the normal distribution according to Proposition 5.1. Thus, the k-fold convolution of PDF of the random variable X still follows the normal distribution. Then the condition distribution of future degradation increment Y_{t+1} is*

$$\begin{aligned} f(Y_{t+1} = y|K_{t+1} = k, \tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t}) &= f^k(\sum_{i=0}^{k} x_i|\boldsymbol{\beta}_t, \Sigma_{\beta_t}) \\ &= Normal\left(k\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t, k\sigma^2(1 + \mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal}))\right) \\ &= \left(2\pi k\sigma^2(1 + \mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal})\right)^{-\frac{1}{2}} \exp\left[\frac{-(y - k\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t)^2}{2k\sigma^2(1 + \mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal})}\right]. \end{aligned} \tag{B.1}$$

According to Proposition 5.1, the number of shocks follows the Negative Binomial

distribution, thus we have

$$\begin{aligned}\mathbb{P}(K_{t+1}=k|\tilde{\boldsymbol{\alpha}}_t,\Sigma_{\alpha_t}) &= \mathrm{NegBin}(\mathrm{e}^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t},\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal}) \\ &= \frac{\Gamma(a+k_{t+1})}{\Gamma(k_{t+1}+1)\Gamma(a)}\hat{p}^a(1-\hat{p})^{k_{t+1}},\end{aligned} \tag{B.2}$$

where where $a=(\boldsymbol{s}_{t+1}\Sigma_{\alpha_t}\boldsymbol{s}_{t+1}^{\intercal})^{-1}$, $b=\boldsymbol{s}_{t+1}\Sigma_{\alpha_t}\boldsymbol{s}_{t+1}^{\intercal}e^{\boldsymbol{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}$, *and* $\hat{p}=\frac{1}{1+b}$

Combining the Equation (B.2) *into Equations* (B.1)*, the joint probability distribution of the random variable* Y_{t+1} *and* K_{t+1} *in epoch* $t+1$ *can be obtained as:*

$$\begin{aligned}f(Y_{t+1}=y,K_{t+1}=k|\tilde{\boldsymbol{\alpha}}_t,\Sigma_{\alpha_t},\tilde{\boldsymbol{\beta}}_t,\Sigma_{\beta_t}) =& \exp\left[-\frac{1}{2}\frac{(y-k\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t)^2}{k\sigma^2(1+\mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal})}\right)\right]\frac{\Gamma(a+k)\hat{p}^a}{\Gamma(k+1)\Gamma(a)}(1-\hat{p})^k \\ &\times\left(2\pi k\sigma^2(1+\mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal})\right)^{-\frac{1}{2}}\end{aligned}$$

The predictive distribution of random variable Y_{t+1} *is*

$$f(Y_{t+1}=y|\tilde{\boldsymbol{\alpha}}_t,\Sigma_{\alpha_t},\tilde{\boldsymbol{\beta}}_t,\Sigma_{\beta_t})=\sum_{k=0}^{\infty}\frac{\Gamma(a+k)}{\Gamma(k+1)\Gamma(a)}(\frac{1}{1+b})^a(\frac{b}{1+b})^k\left[2\pi k\sigma^2(1+d)\right]^{-\frac{1}{2}}\exp\left[-\frac{1}{2}\frac{(y-kc)^2}{k\sigma^2(1+d)}\right],$$

where $a=(\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal})^{-1}$, $b=\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal}e^{\mathbf{s}_{t+1}\tilde{\boldsymbol{\alpha}}_t}$, $c=\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t$, $d=\mathbf{s}_{t+1}\Sigma_{\beta_t}\mathbf{s}_{t+1}^{\intercal}$.

B.3 Proof for Proposition 5.3

We use the following definitions to explain and compare the stochastic order of random variables (Shaked and Shanthikumar, 2007).

Definition 1 *A random variable* A *is less than or equal to a random variable* B *in the usual stochastic order, denoted as* $A\preceq B$*, if*

$$\mathbb{P}(A>x)\leq\mathbb{P}(B>x),\forall x\in(-\infty,+\infty).$$

Definition 2 *A random variable* A *is said to be smaller than or equal to a random variable* B *in the likelihood ratio order, denoted as* $A\preceq_{lr}B$ *if their PDFs* $\frac{f_A(x)}{f_B(x)}$ *is non-decreasing in* x*. Note that the likelihood ratio order is stronger than the usual stochastic order. Hence, if* $A\preceq_{lr}B$*, then* $A\preceq B$*.*

Proof *We first consider the random variable* K_{t+1} *that follows Poisson distribution with parameter* $\lambda(\mathbf{s}_{t+1}) = e^{\mathbf{s}_{t+1}\boldsymbol{\alpha}_t}$ *and* K'_{t+1} *follows Poisson distribution with parameter* $\lambda(\mathbf{s}_{t+1})' = e^{\mathbf{s}_{t+1}\boldsymbol{\alpha}'_t}$. *For Poisson distribution, the random variables* K_{t+1} *are stochastically increasing in* $\lambda(\mathbf{s}_{t+1})$, *that is,* $\lambda(\mathbf{s}_{t+1}) \preceq \lambda'(\mathbf{s}_{t+1})$ *implies* $\langle K_{t+1}|\lambda(\mathbf{s}_{t+1})\rangle \preceq \langle K'_{t+1}|\lambda'(\mathbf{s}_{t+1})\rangle$. *Since* $\lambda(\mathbf{s}_{t+1}) = e^{\mathbf{s}_{t+1}\boldsymbol{\alpha}_t}$, $e^{\mathbf{s}_{t+1}\boldsymbol{\alpha}_t}$ *approximately follows gamma distribution Gammma*(a_{t+1}, b_{t+1}) *based on the proof of Proposition 5.1, where* $a_{t+1} = (\mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal})^{-1}$ *and* $b_{t+1} = \mathbf{s}_{t+1}\Sigma_{\alpha_t}\mathbf{s}_{t+1}^{\intercal}e^{\mathbf{s}_{t+1}\tilde{\alpha}_t}$. *Hence, the stochastic order* $\lambda(\mathbf{s}_{t+1}) \preceq \lambda(\mathbf{s}'_{t+1})$ *if* $a_{t+1} < a'_{t+1}$ *and* $b_{t+1} < b'_{t+1}$.

Thus, the inequality holds

$$\begin{aligned}\langle Y_{t+1}|\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t\rangle &= \sum_{i=1}^{K_{t+1}|\mathbf{P}_t,\mathbf{Q}_t} \langle X_i|\mathbf{N}_t\rangle \\ &\preceq \langle \sum_{i=1}^{K'_{t+1}|\mathbf{P}'_t,\mathbf{Q}'_t} X_i|\mathbf{N}_t\rangle = \langle Y_{t+1}|\mathbf{P}'_t, \mathbf{Q}'_t, \mathbf{N}_t\rangle\end{aligned}$$

The random variables $\langle X_{t+1}|\mathbf{N}_t\rangle$ *is normal distributed with mean* $\mathbf{s}_{t+1}\tilde{\boldsymbol{\beta}}_t$ *and variance* $\sigma^2(1 + \mathbf{s}_j\Sigma_{\beta_t}\mathbf{s}_j^{\intercal})$, *and* $\langle X'_{t+1}|\mathbf{N}'_t\rangle$ *is normal distributed with mean* $\mathbf{s}_j\tilde{\boldsymbol{\beta}}'_{j-1}$ *and variance* $\sigma^2(1 + \mathbf{s}_j\Sigma'_{\beta_t}\mathbf{s}_j^{\intercal})$. *Since* $\Sigma_{\beta_t} = \sigma^2(\mathcal{S}_t\mathcal{S}_t^{\intercal} + \sigma^2\Sigma_{\beta_0}^{-1})^{-1}$, *it can be observed that the updating of* Σ_{β_t} *is independent with the observations. Hence* $\Sigma_{\beta_t} = \Sigma'_{\beta_t}$ *given the same working condition and* Σ_{β_0}. *That is,* X_j *and* X'_j *are normally distributed with different means but common variance. Define* $c_j = \mathbf{s}_j\tilde{\boldsymbol{\beta}}_t$ *and* $c'_j = \mathbf{s}_j\tilde{\boldsymbol{\beta}}'_t$, *where* $\tilde{\boldsymbol{\beta}}_t = \mathbf{M}_t^{-1}(\mathbf{N}_t + \sigma^2\Sigma_{\beta_0}^{-1}\tilde{\boldsymbol{\beta}}_0)$ *and* $\tilde{\boldsymbol{\beta}}'_t = \mathbf{M}_t^{-1}(\mathbf{N}'_t + \sigma^2\Sigma_{\beta_0}^{-1}\tilde{\boldsymbol{\beta}}_0)$. *If* $c_j \leq c'_j$, *it is sufficient to show the stochastic order* $\langle X_{t+1}|\mathbf{N}_t\rangle \preceq \langle X_{t+1}|\mathbf{N}'_t\rangle$. *Thus, the inequality holds*

$$\begin{aligned}\langle Y_{t+1}|\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t\rangle &= \sum_{i=1}^{K_{t+1}|\mathbf{P}_t,\mathbf{Q}_t} \langle X_i|\mathbf{N}_t\rangle \\ &\preceq \langle \sum_{i=1}^{K_{t+1}|\mathbf{P}_t,\mathbf{Q}_t} X'_i|\mathbf{N}'_t\rangle = \langle Y_{t+1}|\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}'_t\rangle\end{aligned}$$

Therefore, by merging the two inequalities, it can be concluded that

$$\begin{aligned}\langle Y_{t+1}|\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t\rangle &\preceq \langle Y_{t+1}|\mathbf{P}'_t, \mathbf{Q}'_t, \mathbf{N}_t\rangle \\ &\preceq \langle Y_{t+1}|\mathbf{P}'_t, \mathbf{Q}'_t, \mathbf{N}'_t\rangle.\end{aligned}$$

B.4 Proof for Proposition 5.4

Proof *First, we need to prove the value function is non-decreasing in z_t for all t, $t = 0, 1, ..., T$ by induction. At the last epoch $t = T$, the value function*

$$V_{T,\boldsymbol{\chi}_T}(z_T, \mathbf{P}_T, \mathbf{Q}_T, \mathbf{N}_T) = \begin{cases} 0, & z_T < D \\ c_r, & z_T \geq D. \end{cases}$$

It can be observed that $V_{T,\boldsymbol{\chi}_T}(z_T, \mathbf{P}_T, \mathbf{Q}_T, \mathbf{N}_T)$ is constant and satisfies the non-decreasing property. Assuming that the property holds at epoch $t+1$, $t = 0, ..., T-1$, we can show that it also holds at t.

Consider the case when the accumulative degradation $z_t < D$, we have $V_{t,\boldsymbol{\chi}_t}(z_t, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t) = \min\{c_p + V(0), \gamma\mathbb{E}\left[V_{t+1,\boldsymbol{\chi}_{t+1}}(z_{t+1}, \mathbf{P}_{t+1}, \mathbf{Q}_{t+1}, \mathbf{N}_{t+1})\right]\}$. As $V_{t+1}(z_{t+1}, \mathbf{P}_{t+1}, \mathbf{Q}_{t+1}, \mathbf{N}_{t+1})$ is non-decreasing in z_{t+1} and $\langle z_{t+1}|z_t^-, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t\rangle \preceq \langle z_{t+1}|z_t^+, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t\rangle$ holds given the working condition and $z_t^- \leq z_t^+$, and $\mathbb{E}\left[V_{t+1,\boldsymbol{\chi}_{t+1}}(z_t + y_{t+1}, \mathbf{P}_{t+1}, \mathbf{Q}_{t+1}, \mathbf{N}_{t+1}|z_t, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t)\right]$ is non-decreasing in z_t. Therefore, we can conclude that the value function $V_{t,\boldsymbol{\chi}_t}(z_t, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t)$ is non-decreasing in z_t when $z_t < D$. On the other hand, when the accumulative degradation $z_t \geq D$, the value function $V_{t,\boldsymbol{\chi}_t}(z_t, \mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t) = c_r + V_{t,\boldsymbol{\chi}_t}(0, \mathbf{P}_0, \mathbf{Q}_0, \mathbf{N}_0)$, which a constant and holds the non-decreasing property. Hence, the induction holds hypothesis holds for all $t, 0 \leq t \leq T$.

Then, according to the Bellman equation, PM is optimal when the inequality

$$c_p + V_{t,\boldsymbol{\chi}_t}(0, \mathbf{P}_0, \mathbf{Q}_0, \mathbf{N}_0) \leq \gamma\mathbb{E}\left[V_{t+1,\boldsymbol{\chi}_{t+1}}(z_{t+1}, \mathbf{P}_{t+1}, \mathbf{Q}_{t+1}, \mathbf{N}_{t+1})\right]$$

holds. Since the left-hand side of the inequality is constant with respect to z_t and the right-hand side is non-decreasing in z_t. Therefore, this establishes a control limit policy that there exists a threshold $\zeta_{t,\boldsymbol{\chi}_t}(\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t)$, such that the PM is performed when the

accumulative degradation $z_t \geq \zeta_{t,\boldsymbol{\chi}_t}(\mathbf{P}_t, \mathbf{Q}_t, \mathbf{N}_t)$.

B.5 Proof for Proposition 5.5

Proof *First, we prove that the maintenance policy is still a control limit policy when applying the most likely distribution. Given the observations up to epoch* t*, the parameters are updated according to Lemma 5.1 and Lemma 5.2, and the most likely predictive distribution of* Y_{t+i} *at epoch* $t+i$ *is* $\hat{f}_{t+i,t}(Y_{t+i}|\tilde{\boldsymbol{\alpha}}_t, \Sigma_{\alpha_t}, \tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t})$ *for* $i = 1, ..., T-t$*. At the last epoch* T*, the value function*

$$\hat{V}_{T,t,\boldsymbol{\chi}_T}(z_T) = \begin{cases} 0, & z_T < D \\ c_r, & z_T \geq D. \end{cases}$$

It can be observed that $\hat{V}_{T,t,\boldsymbol{\chi}_T}(z_T)$ *is constant and satisfies the non-decreasing property. Assuming that the property holds at epoch* $t+1$*,* $t = 0, ..., T-1$*, we can show that it also holds at* t*.*

Consider the case when the accumulative degradation $z_t < D$*, we have* $\hat{V}_{t,t,\boldsymbol{\chi}_t}(z_t) = \min\{c_p + V(0), \gamma\mathbb{E}\left[\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_{t+1})\right]\}$*. As* $\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_{t+1})$ *is non-decreasing in* z_{t+1} *and* $\langle z_{t+1}|z_t^-, \tilde{\boldsymbol{\alpha}}_t, \Sigma_{\alpha_t}, \tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t}\rangle \preceq \langle z_{t+1}|z_t^+, \tilde{\boldsymbol{\alpha}}_t, \Sigma_{\alpha_t}, \tilde{\boldsymbol{\beta}}_t, \Sigma_{\beta_t}\rangle$ *holds given the working condition and* $z_t^- \leq z_t^+$*, which yields* $\mathbb{E}\left[\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_t + y_{t+1}|z_t)\right]$ *is non-decreasing in* z_t*. Therefore, we can conclude that the value function* $\hat{V}_{t,t,\boldsymbol{\chi}_t}(z_t)$ *is non-decreasing in* z_t *when* $z_t < D$*. On the other hand, when the accumulative degradation* $z_t \geq D$*, the value function* $\hat{V}_{t,t,\boldsymbol{\chi}_t}(z_t) = c_r + \hat{V}_{t,t,\boldsymbol{\chi}_t}(0)$*, which is a constant and holds the non-decreasing property. Hence, the induction holds hypothesis holds for all* $t, 0 \leq t \leq T$*. Afterwards, we can further derive that the PM threshold is a control limit policy, which is similar to the proof of Proposition 5.4.*

From the control limit policy, it is known that PM is optimal when the inequality

$$c_p + \hat{V}_{t,t,\boldsymbol{\chi}_t}(0) \leq \gamma\mathbb{E}\left[\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_{t+1})\right]$$

holds. As stated in Proposition 5.3, if $a_t \leq a_t'$*,* $b_t \leq b_t'$ *and* $c_t \leq c_t'$*, then given any working condition, the stochastic order* $\langle Y_{t+1}|a_t, b_t, c_t\rangle \preceq \langle Y_{t+1}|a_t', b_t', c_t'\rangle$*. Hence, the*

right-hand side

$$\begin{aligned}\mathbb{E}\left[\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_{t+1})\right] &= \sum_{z_{t+1}\in\boldsymbol{\Delta}} \hat{F}_{t+1,t}(z_{t+1}-z_t|\tilde{\alpha}_t,\Sigma_{\alpha_t},\tilde{\beta}_t,\Sigma_{\beta_t})\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_{t+1}) \\ &\leq \sum_{z_{t+1}\in\boldsymbol{\Delta}} \hat{F}_{t+1,t}(z_{t+1}-z_t|\tilde{\alpha}'_t,\Sigma'_{\alpha_t},\tilde{\beta}'_t,\Sigma'_{\beta_t})\hat{V}_{t+1,t,\boldsymbol{\chi}_{t+1}}(z_{t+1})\end{aligned}$$

The left-hand side of the inequality is constant with respect to z_t, hence the PM threshold $\hat{\zeta}_{t,\boldsymbol{\chi}_t}(\mathbf{P}_t,\mathbf{Q}_t,\mathbf{N}_t) \geq \hat{\zeta}_{t,\boldsymbol{\chi}_t}(\mathbf{P}'_t,\mathbf{Q}'_t,\mathbf{N}'_t)$.

B.6 MLE for the initial value of the parameters

We estimate the initial values for the shock arrivals and the shock magnitude separately since they are independent in the CP process. For the degradation paths, we sort the historical shock arrivals k_i within one epoch and the corresponding working condition vectors $\boldsymbol{s}_i$ as a tuple $(k_i,\boldsymbol{s}_i), i \in I$, where $\boldsymbol{s}_i = [1, s_{i1}, ..., s_{ip}]$, p is the number of covariates, and I is the number of paths for shock arrivals. The historical magnitude of shocks and the corresponding working condition vector as a tuple $(x_i,\bar{\boldsymbol{s}}_i), i \in J$, where J is the number of path for degradation magnitude, $\bar{\boldsymbol{s}}_i$ is a $p+1$ dimensional vector.

Firstly, we show the MLE for estimating the prior distribution of $\boldsymbol{\alpha}$, which is the coefficient vector. In the Poisson regression model, the number of shocks follows the Poisson distribution $K \sim Poisson(e^{\boldsymbol{s}\boldsymbol{\alpha}})$. The likelihood function for I independent Poisson observation is a product of probability as

$$L(\boldsymbol{\alpha}) = \prod_{i=1}^{I} \frac{e^{-\boldsymbol{s}_i\boldsymbol{\alpha}}\boldsymbol{s}_i\boldsymbol{\alpha}^{k_i}}{k_i!}.$$

The log-likelihood function is

$$\log L(\boldsymbol{\alpha}) = \sum_{i=1}^{I} \left\{k_i \log(\boldsymbol{s}_i\boldsymbol{\alpha}) - \boldsymbol{s}_i\boldsymbol{\alpha}\right\}.$$

Hence, the estimated $\boldsymbol{\alpha}_0 = \arg\max_{\boldsymbol{\alpha}} \log L(\boldsymbol{\alpha})$, which can be solved by using standard numerical methods, e.g. Newton-Raphson technique.

Based on the fact that the maximum likelihood estimator $\boldsymbol{\alpha}_0$ approximates the mean of $\boldsymbol{\alpha}$, and the covariance matrix is the inverse of the information matrix, which is the minus expected value of the second derivative of the log-likelihood as $-\mathbb{E}(\frac{\partial^2 \log L(\boldsymbol{\alpha})}{\partial^2 \boldsymbol{\alpha}})$. Hence, the covariance mateix can be estimated as $\Sigma_{\alpha_0} = \boldsymbol{S}^\intercal \boldsymbol{W} \boldsymbol{S}$, where $\boldsymbol{W}$ is a diagonal matrix, $w_i = \boldsymbol{s}_i \boldsymbol{\alpha}_0, i \in I$, and $\boldsymbol{S}$ is a $I \times (P+1)$ matrix that contains the values of working conditions. The proof of covariance matrix can be found in Appendix B.5.2 of the note (Rodríguez, 2007).

For estimating the prior distribution of $\boldsymbol{\beta}$, which is the coefficient vector of working conditions for linear regression. In the Linear regression model, the shock magnitude follows normal distribution $X \sim \mathcal{N}(\bar{\boldsymbol{s}}\boldsymbol{\beta}, \sigma^2)$. The likelihood function for J independent observations is:

$$L(\boldsymbol{\beta}, \sigma) = \prod_{i=1}^{J} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_i - \bar{\boldsymbol{s}}\boldsymbol{\beta})^2}{2\sigma^2}}. \tag{B.3}$$

Differentiating the logarithm of Equation (B.3) with respect to $\boldsymbol{\beta}$, the MLE of $\boldsymbol{\beta}$ can be obtained as $\boldsymbol{\beta}_0 = (\bar{\boldsymbol{S}}^\intercal \bar{\boldsymbol{S}})^{-1} \bar{\boldsymbol{S}}^\intercal \boldsymbol{x}$, where $\bar{\boldsymbol{S}}$ is a $J \times (p+1)$ matrix that contains values of the working condition, $\boldsymbol{x}$ is a J dimensional vector of the shock magnitude. The variance-covariance matrix of the MLE estimators is then obtained by taking the negative inverse of the Hessian matrix of the log-likelihood function evaluated at the MLE, that is $\Sigma_{\beta_0} = (\bar{\boldsymbol{S}}^\intercal \bar{\boldsymbol{S}})^{-1} \sigma^2$. When the σ is unknown to the decision maker, it can be estimated by $\sigma = \frac{1}{J}(\boldsymbol{x} - \bar{\boldsymbol{S}}\boldsymbol{\beta})^\intercal (\boldsymbol{x} - \bar{\boldsymbol{S}}\boldsymbol{\beta})$ by differentiating Equation (B.3) with respect to σ. The derivate process for the mean value $\boldsymbol{\beta}_0$, the variance matrix Σ_{β_0} and variance σ can be found in Appendix B.2 and B.3 of the note (Rodríguez, 2007).

B.7 Normal approximation

Consider a random variable $X \sim$ Gamma(a, b), where a is the shape parameter and b is scale parameter, and the distribution function is:

$$\mathbb{P}(x|a, b) = \frac{x^{a-1} e^{-\frac{x}{b}}}{\Gamma(a) b^a}.$$

The logarithm $Y = \log X$ follows log-gamma distribution. When a tends to be infinite, the log-gamma distribution is normal distribution:

$$y = \log x \sim \mathcal{N}(\log a + \log b, a^{-1}).$$

By setting $a = k, k \in \{0, 1, ...\}$, and $b = 1$,

$$\mathbb{P}(x|k, 1) = \frac{x^{k-1}e^{-x}}{(k-1)!}.$$

The distribution of y is obtained as

$$\begin{aligned}
\mathbb{P}(y|k, 1) &= \mathbb{P}(x = e^y|k, 1)\frac{\partial e^y}{\partial y} \\
&= \frac{e^{ky}e^{-e^y}}{(k-1)!} \\
&\approx \mathcal{N}(y|\log k, k^{-1})
\end{aligned}$$

B.8 Heuristic algorithm

Algorithm 2: Backward dynamic programming algorithm with most likelihood distribution

Data: Initial parameters $\tilde{\boldsymbol{\alpha}}_0$, $\tilde{\boldsymbol{\beta}}_0$, Σ_{α_0}, Σ_{β_0}, cost parameter c_p, c_r, vector of covariates $\boldsymbol{\chi}_t$, last epoch T, discount factor γ

Initialise: Terminal value function $V_{T+1,t,\chi_{T+1}}(\cdot) = 0$ for all t

for $t = 0, 1, 2, ..., T$ **do**

  for $u = T - t, ..., 0$ **do**

    forall $z_{t+u} \in \Delta$ **do**

      $\hat{V}_{t+u,t,\boldsymbol{\chi}_{t+u}}(z_{t+u}) =$
      $\gamma \sum_{z_{t+u+1} \in \boldsymbol{\Delta}} \hat{F}_{t+u+1,t}(z_{t+u+1} - z_{t+u}) \hat{V}_{t+u+1,t,\boldsymbol{\chi}_{t+u+1}}(z_{t+u+1})$

      if $\hat{V}_{t+u,t,\boldsymbol{\chi}_{t+u}}(z_{t+u}) > c_p + \hat{V}_{t+u,t,\chi_{t+u}}(0)$ **then**

        $\hat{V}_{t+u,t,\boldsymbol{\chi}_{t+u}}(z_{t+u}) = c_p + \hat{V}_{t+u,t,\chi_{t+u}}(0)$

        if $u = 0$ **then**

          $\pi_{t,t,\boldsymbol{\chi}_t}(z_t) \leftarrow PM$

        end

      else

        if $u = 0$ **then**

          $\pi_{t,t,\boldsymbol{\chi}_t}(z_t) \leftarrow DN$

        end

      end

    end

  end

  Update $\tilde{\boldsymbol{\alpha}}_{t+1}$, $\tilde{\boldsymbol{\beta}}_{t+1}$, $\Sigma_{\alpha_{t+1}}$, $\Sigma_{\beta_{t+1}}$ based on new observations $\mathbf{x}_t, k_t$

end

Output: $\hat{V}_{t,t,\boldsymbol{\chi}_t}(z_t)$ and maintenance policy $\pi_{t,t,\boldsymbol{\chi}_t}(z_t)$

Bibliography

Rosmaini Ahmad and Shahrul Kamaruddin. An overview of time-based and condition-based maintenance in industrial application. *Computers & Industrial Engineering*, 63(1):135–149, 2012.

Anuoluwapo Ajayi, Lukumon Oyedele, Hakeem Owolabi, Olugbenga Akinade, Muhammad Bilal, Juan Manuel Davila Delgado, and Lukman Akanbi. Deep learning models for health and safety risk prediction in power infrastructure projects. *Risk Analysis*, 40(10):2019–2039, 2020.

Angelo Alessandri and Mauro Gaggero. Fast moving horizon state estimation for discrete-time systems using single and multi iteration descent methods. *IEEE Transactions on Automatic Control*, 62(9):4499–4511, 2017.

Carlos Alós-Ferrer and Michele Garagnani. Part-time Bayesians: Incentives and behavioral heterogeneity in belief updating. *Management Science*, 69(9):5523–5542, 2023.

Joachim Arts, Robert N Boute, Stijn Loeys, and Heletjé E van Staden. Fifty years of maintenance optimization: Reflections and perspectives. *European Journal of Operational Research*, 2024.

Katy S Azoury and Julia Miyaoka. Optimal policies and approximations for a Bayesian linear regression inventory model. *Management Science*, 55(5):813–826, 2009.

Suk Joo Bae, Way Kuo, and Paul H Kvam. Degradation models and implied lifetime distributions. *Reliability Engineering & System Safety*, 92(5):601–608, 2007.

Vilijandas Bagdonavicius and Mikhail S Nikulin. Estimation in degradation models with explanatory variables. *Lifetime Data Analysis*, 7:85–103, 2001.

Gah-Yi Ban and Cynthia Rudin. The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1):90–108, 2019.

Kasun Bandara, Hansika Hewamalage, Yuan-Hao Liu, Yanfei Kang, and Christoph Bergmeir. Improving the accuracy of global forecasting models using time series data augmentation. *Pattern Recognition*, 120:108148, 2021.

Richard Barlow and Larry Hunter. Optimum preventive maintenance policies. *Operations Research*, 8(1):90–100, 1960.

Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294, 2020.

Yoshua Bengio. Using a financial training criterion rather than a prediction criterion. *International Journal of Neural Systems*, 8(04):433–443, 1997.

Carolin Benjamins, Theresa Eimer, Frederik Schubert, Aditya Mohan, André Biedenkapp, Bodo Rosenhahn, Frank Hutter, and Marius Lindauer. Contextualize me–the case for context in reinforcement learning. *arXiv preprint arXiv:2202.04500*, 2022.

Ali Berrichi, Lionel Amodeo, Farouk Yalaoui, Eric Châtelet, and Mohamed Mezghiche. Bi-objective optimization algorithms for joint production and maintenance scheduling: application to the parallel machine problem. *Journal of Intelligent Manufacturing*, 20:389–400, 2009.

Dimitri Bertsekas. *Dynamic programming and optimal control: Volume I*, volume 4. Athena scientific, 2012.

Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020.

Linkan Bian and Nagi Gebraeel. Computing and updating the first-passage time distribution for randomly evolving degradation signals. *IIE Transactions*, 44(11):974–987, 2012.

Linkan Bian, Nagi Gebraeel, and Jeffrey P Kharoufeh. Degradation modeling for real-time estimation of residual lifetimes in dynamic environments. *IIE Transactions*, 47 (5):471–486, 2015.

Tomaž Bučar, Marko Nagode, and Matija Fajdiga. Reliability approximation using finite weibull mixture distributions. *Reliability Engineering & System Safety*, 84(3): 241–251, 2004.

Wahyu Caesarendra, Achmad Widodo, and Bo-Suk Yang. Application of relevance vector machine and logistic regression for machine degradation assessment. *Mechanical Systems and Signal Processing*, 24(4):1161–1171, 2010.

Weicheng Cai, Danwei Cai, Wenbo Liu, Gang Li, and Ming Li. Countermeasures for automatic speaker verification replay spoofing attack: On data augmentation, feature representation, classification and fusion. In *Interspeech*, pages 17–21, 2017.

Yue Cai, Ruud H Teunter, and Bram de Jonge. A data-driven approach for condition-based maintenance optimization. *European Journal of Operational Research*, 311(2): 730–738, 2023.

Zhuoran Cai, Wenxuan Ma, Xinrui Wang, Hanhong Wang, and Zhongming Feng. The performance analysis of time series data augmentation technology for small sample communication device recognition. *IEEE Transactions on Reliability*, 72(2):574–585, 2022.

C Richard Cassady and Erhan Kutanoglu. Integrating preventive maintenance planning and production scheduling for a single machine. *IEEE Transactions on Reliability*, 54(2):304–309, 2005.

Andrea Cester, Salvatore Cimino, Enrique Miranda, Andrea Candelori, Gabriella Ghidini, and Alessandro Paccagnella. Statistical model for radiation-induced wear-out

of ultra-thin gate oxides after exposure to heavy ion irradiation. *IEEE Transactions on Nuclear Science*, 50(6):2167–2175, 2003.

Antoni B Chan and Nuno Vasconcelos. Bayesian Poisson regression for crowd counting. In *2009 IEEE 12th International Conference on Computer Vision*, pages 545–551. IEEE, 2009.

J-K Chan and Leonard Shaw. Modeling repairable systems with failure rates that depend on age and maintenance. *IEEE Transactions on Reliability*, 42(4):566–571, 1993.

Mingzhi Chen, Haidong Shao, Haoxuan Dou, Wei Li, and Bin Liu. Data augmentation and intelligent fault diagnosis of planetary gearbox using ilofgan under extremely limited samples. *IEEE Transactions on Reliability*, 72(3):1029–1037, 2022.

Nan Chen, Zhi-Sheng Ye, Yisha Xiang, and Linmiao Zhang. Condition-based maintenance using the inverse Gaussian degradation model. *European Journal of Operational Research*, 243(1):190–199, 2015.

Wei Chen and Songhua Hao. Condition-based operation and maintenance strategy for load-sharing systems based on wiener process. *IEEE Transactions on Reliability*, 2025.

Nan Cheng and Tao Yuan. Nonparametric Bayesian lifetime data analysis using dirichlet process lognormal mixture model. *Naval Research Logistics (NRL)*, 60(3):208–221, 2013.

Dong-Jin Choi and Heekyung Park. A hybrid artificial neural network as a software sensor for optimal control of a wastewater treatment process. *Water Research*, 35 (16):3959–3967, 2001.

Cindy L Christiansen and Carl N Morris. Hierarchical Poisson regression modeling. *Journal of the American Statistical Association*, 92(438):618–632, 1997.

Ya-Tang Chuang and Michael Jong Kim. Bayesian inventory control: Accelerated demand learning via exploration boosts. *Operations Research*, 71(5):1515–1529, 2023.

Michele Compare, Piero Baraldi, and Enrico Zio. Challenges to IoT-enabled predictive maintenance for industry 4.0. *IEEE Internet of Things Journal*, 7(5):4585–4597, 2019.

R Dennis Cook and Sanford Weisberg. *Residuals and influence in regression*. New York: Chapman and Hall, 1982.

David Roxbee Cox and Peter AW Lewis. *The statistical analysis of series of events*. Springer, 1966.

Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.

Helen Cui. Accelerated temperature cycle test and coffin-manson model for electronic packaging. In *Annual Reliability and Maintainability Symposium, 2005. Proceedings.*, pages 556–560. IEEE, 2005.

Cuong Duc Dao and Ming J Zuo. Selective maintenance for multistate series systems with s-dependent components. *IEEE Transactions on Reliability*, 65(2):525–539, 2015.

Ajodhya Nath Das and Damodar Acharya. Age replacement of components during IFR delay time. *IEEE Transactions on Reliability*, 53(3):306–312, 2004.

Cameron Davidson-Pilon. Lifelines: Survival analysis in Python. *Journal of Open Source Software*, 4(40):1317, 2019. doi: 10.21105/joss.01317. URL `https://doi.org/10.21105/joss.01317`.

Bram De Jonge and Philip A Scarf. A review on maintenance optimization. *European Journal of Operational Research*, 285(3):805–824, 2020.

Bram de Jonge, Arjan S Dijkstra, and Ward Romeijnders. Cost benefits of postponing time-based maintenance under lifetime distribution uncertainty. *Reliability Engineering & System Safety*, 140:15–21, 2015.

Bram de Jonge, Ruud Teunter, and Tiedo Tinga. The influence of practical factors on the benefits of condition-based maintenance over time-based maintenance. *Reliability Engineering & System Safety*, 158:21–30, 2017.

Ana Debón, Andres Carrión, E Cabrera, and H Solano. Comparing risk of failure models in water supply networks using ROC curves. *Reliability Engineering & System Safety*, 95(1):43–48, 2010.

Akash Deep, Shiyu Zhou, Dharmaraj Veeramani, and Yong Chen. Partially observable Markov decision process-based optimal maintenance planning with time-dependent observations. *European Journal of Operational Research*, 311(2):533–544, 2023.

Marc Peter Deisenroth, A Aldo Faisal, and Cheng Soon Ong. *Mathematics for machine learning*. Cambridge University Press, 2020.

Sumeyra Demir, Krystof Mincev, Koen Kok, and Nikolaos G Paterakis. Data augmentation for time series regression: Applying transformations, autoencoders and adversarial networks to electricity price forecasting. *Applied Energy*, 304:117695, 2021.

Chuanjin Deng, Wenjie Zha, Kaibin Chen, and Hongfeng Lv. Research on acceleration test model and parameter determination method of integrated circuit. In *2023 24th International Conference on Electronic Packaging Technology (ICEPT)*, pages 1–5. IEEE, 2023.

Cyrus Derman. Optimal replacement and maintenance under Markovian deterioration with probability bounds on failure. *Management Science*, 9(3):478–481, 1963.

Phuc Do Van, Anne Barros, Christophe Bérenguer, Keomany Bouvard, and Florent Brissaud. Dynamic grouping maintenance with time limited opportunities. *Reliability Engineering & System Safety*, 120:51–59, 2013.

Honghui Dong, Hao Ma, Zhipeng Wang, Jie Man, Limin Jia, and Yong Qin. An online health monitoring framework for traction motors in high-speed trains using

temperature signals. *IEEE Transactions on Industrial Informatics*, 19(2):1389–1400, 2022.

Priya Donti, Brandon Amos, and J Zico Kolter. Task-based end-to-end model learning in stochastic optimization. *Advances in Neural Information Processing Systems*, 30, 2017.

Collin Drent, Melvin Drent, Joachim Arts, and Stella Kapodistria. Real-time integrated learning and decision making for cumulative shock degradation. *Manufacturing & Service Operations Management*, 25(1):235–253, 2023.

Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1):9–26, 2022.

Adam N Elmachtoub, Henry Lam, Haofeng Zhang, and Yunfan Zhao. Estimate-then-optimize versus integrated-estimation-optimization: A stochastic dominance perspective. *arXiv preprint arXiv:2304.06833*, 2023.

Alaa H Elwany and Nagi Z Gebraeel. Sensor-driven prognostic models for equipment replacement and spare parts inventory. *IIE Transactions*, 40(7):629–639, 2008.

Alaa H Elwany, Nagi Z Gebraeel, and Lisa M Maillart. Structured replacement policies for components with complex degradation processes and dedicated sensors. *Operations Research*, 59(3):684–695, 2011.

Luis A Escobar and William Q Meeker. A review of accelerated test models. *Statistical science*, pages 552–577, 2006.

Theodoros Evgeniou, Massimiliano Pontil, and Olivier Toubia. A convex optimization approach to modeling consumer heterogeneity in conjoint estimation. *Marketing Science*, 26(6):805–818, 2007.

Amir Falamarzi, Sara Moridpour, and Majidreza Nazem. Development of a tram track degradation prediction model based on the acceleration data. *Structure and Infrastructure Engineering*, 15(10):1308–1318, 2019.

Dongming Fan, Aibo Zhang, Qiang Feng, Baoping Cai, Yiliu Liu, and Yi Ren. Group maintenance optimization of subsea xmas trees with stochastic dependency. *Reliability Engineering & System Safety*, 209:107450, 2021.

Jiajie Fan, Kam Chuen Yung, and Michael Pecht. Physics-of-failure-based prognostics and health management for high-power white light-emitting diode lighting. *IEEE Transactions on device and materials reliability*, 11(3):407–416, 2011.

Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Data augmentation using synthetic data for time series classification with deep residual networks. *arXiv preprint arXiv:1808.02455*, 2018.

Flight Safety Foundation. Accident description asn database 324966. `https://asn.flightsafety.org/asndb/324966`, 2024. Accessed: 24 October 2024.

Elizabeth Fons, Paula Dawson, Xiao-jun Zeng, John Keane, and Alexandros Iosifidis. Adaptive weighting scheme for automatic time-series data augmentation. *arXiv preprint arXiv:2102.08310*, 2021.

Mitra Fouladirad, Christian Paroissin, and Antoine Grall. Sensitivity of optimal replacement policies to lifetime parameter estimates. *European Journal of Operational Research*, 266(3):963–975, 2018.

Marta A Freitas, Maria Lúiza G De Toledo, Enrico A Colosimo, and Magda C Pires. Using degradation data to assess reliability: a case study on train wheel degradation. *Quality and Reliability Engineering International*, 25(5):607–629, 2009.

Bo Fu, Wangjun Yuan, Xiaolong Cui, Tian Yu, Xilin Zhao, and Chaoshun Li. Correlation analysis and augmentation of samples for a bidirectional gate recurrent unit network for the remaining useful life prediction of bearings. *IEEE Sensors Journal*, 21(6):7989–8001, 2020.

Daniel Garcia-Pozuelo, Jorge Yunta, Oluremi Olatunbosun, Xiaoguang Yang, and Vicente Diaz. A strain-based method to estimate slip angle and tire working conditions for intelligent tires using fuzzy logic. *Sensors*, 17(4):874, 2017.

Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.

Faeze Ghofrani, Hongyue Sun, and Qing He. Analyzing risk of service failures in heavy haul rail lines: A hybrid approach for imbalanced data. *Risk Analysis*, 42(8):1852–1871, 2022.

Kevin D Glazebrook, HM Mitchell, and PS Ansell. Index policies for the maintenance of a collection of machines by a set of repairmen. *European Journal of Operational Research*, 165(1):267–284, 2005.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

Nima Gorjian, Lin Ma, Murthy Mittinty, Prasad Yarlagadda, and Yong Sun. A review on degradation models in reliability analysis. In *Engineering Asset Lifecycle Management: Proceedings of the 4th World Congress on Engineering Asset Management (WCEAM 2009), 28-30 September 2009*, pages 369–384. Springer, 2010.

Yihan Guo, Cunbao Ma, Xu Dong, Yan Liang, and Biyuan Hu. A novel health indicator based on hysteresis loop for health prediction of flight control systems. *Measurement*, 186:110076, 2021.

Allan Gut. The multivariate Normal distribution. In *An Intermediate Course in Probability*, pages 117–145. Springer, 2009.

Assaf Hallak, Dotan Di Castro, and Shie Mannor. Contextual Markov decision processes. *arXiv preprint arXiv:1502.02259*, 2015.

Zhongyang Han, Jun Zhao, Henry Leung, King Fai Ma, and Wei Wang. A review of deep learning models for time series prediction. *IEEE Sensors Journal*, 21(6): 7833–7848, 2019.

J Michael Harrison, N Bora Keskin, and Assaf Zeevi. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science*, 58(3): 570–586, 2012.

Jichao He, David W Wanik, Brian M Hartman, Emmanouil N Anagnostou, Marina Astitha, and Maria EB Frediani. Nonparametric tree-based predictive modeling of storm outages on an electric distribution network. *Risk Analysis*, 37(3):441–458, 2017.

Arild Hoff, Irina Gribkovskaia, Gilbert Laporte, and Arne Løkketangen. Lasso solution strategies for the vehicle routing problem with pickups and deliveries. *European Journal of Operational Research*, 192(3):755–766, 2009.

YT Hou, MF Li, Y Jin, and WH Lai. Direct tunneling hole currents through ultrathin gate oxides in metal-oxide-semiconductor devices. *Journal of Applied Physics*, 91(1): 258–264, 2002.

Jiawen Hu and Piao Chen. Predictive maintenance of systems subject to hard failure based on proportional hazards model. *Reliability Engineering & System Safety*, 196: 106707, 2020.

Jiawen Hu, Zuhua Jiang, and Haitao Liao. Preventive maintenance of a batch production system under time-varying operational condition. *International Journal of Production Research*, 55(19):5681–5705, 2017.

Jiawen Hu, Ancha Xu, Bo Li, and Haitao Liao. Condition-based maintenance planning for multi-state systems under time-varying environmental conditions. *Computers & Industrial Engineering*, 158:107380, 2021.

Catherine Huber-Carol and Ilia Vonta. Frailty models for arbitrarily censored and truncated data. *Lifetime Data Analysis*, 10:369–388, 2004.

Matti Huotari, Shashank Arora, Avleen Malhi, and Kary Främling. Comparing seven methods for state-of-health time series prediction for the lithium-ion battery packs of forklifts. *Applied Soft Computing*, 111:107670, 2021.

ISO. ISO 55000:2014 Asset Management – Overview, Principles and Terminology. Technical Report ISO 55000:2014, ISO, 2014. Available at: `https://www.iso.org/standard/55088.html`.

Brian Kenji Iwana and Seiichi Uchida. An empirical survey of data augmentation for time series classification with neural networks. *PLOS ONE*, 16(7):e0254841, 2021.

Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low Bellman rank are pac-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.

R Jiang. A simple approximation for the renewal function with an increasing failure rate. *Reliability Engineering & System Safety*, 95(9):963–969, 2010.

Shwu-Tzy Jiang, Thomas L Landers, and Teri Reed Rhoads. Assessment of repairable-system reliability using proportional intensity models: A review. *IEEE Transactions on Reliability*, 55(2):328–336, 2006.

Mauricio Junca and Mauricio Sanchez-Silva. Optimal maintenance policy for a compound Poisson shock model. *IEEE Transactions on Reliability*, 62(1):66–72, 2013.

Elnaz Kabir, Seth D Guikema, and Steven M Quiring. Power outage prediction using data streams: An adaptive ensemble learning approach with a feature-and performance-based weighting mechanism. *Risk Analysis*, 44(3):686–704, 2024.

Michael J Kalsher, Michael S Wogalter, RW Lim, and KR Laughery. Consumer knowledge of tire maintenance and aging hazard. *Proceedings of the Human Factors and Ergonomics Society*, 49:1757, 2005.

Zahra Karevan and Johan AK Suykens. Transductive LSTM for time-series prediction: An application to weather forecasting. *Neural Networks*, 125:1–9, 2020.

Xiao Ke, Jiawei Zou, and Yuzhen Niu. End-to-end automatic image annotation based on deep CNN and multi-label data augmentation. *IEEE Transactions on Multimedia*, 21(8):2093–2106, 2019.

Akram Khaleghei and Michael Jong Kim. Optimal control of partially observable semi-Markovian failing systems: An analysis using a phase methodology. *Operations Research*, 69(4):1282–1304, 2021.

Michael Jong Kim. Robust control of partially observable failing systems. *Operations Research*, 64(4):999–1014, 2016.

Michael Jong Kim and Viliam Makis. Joint optimization of sampling and control of partially observable failing systems. *Operations Research*, 61(3):777–790, 2013.

Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, pages 1885–1894. PMLR, 2017.

Jin-Zhen Kong, Dong Wang, Tongtong Yan, Jingzhe Zhu, and Xi Zhang. Accelerated stress factors based nonlinear Wiener process model for lithium-ion battery prognostics. *IEEE Transactions on Industrial Electronics*, 69(11):11665–11674, 2021.

Vasiliy V Krivtsov. Practical extensions to NHPP application in repairable system reliability analysis. *Reliability Engineering & System Safety*, 92(5):560–562, 2007.

Jerry Lawless and Martin Crowder. Covariates and random effects in a gamma process model with application to degradation and failure. *Lifetime Data Analysis*, 10:213–227, 2004.

Wee Lee, Nan Rong, and David Hsu. What makes some POMDP problems easy to approximate? *Advances in Neural Information Processing Systems*, 20, 2007.

Yaguo Lei, Naipeng Li, Liang Guo, Ningbo Li, Tao Yan, and Jing Lin. Machinery health prognostics: A systematic review from data acquisition to rul prediction. *Mechanical Systems and Signal Processing*, 104:799–834, 2018.

Retsef Levi, Robin O Roundy, and David B Shmoys. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research*, 32(4):821–839, 2007.

Mingyang Li and Jian Liu. Bayesian hazard modeling based on lifetime data with latent heterogeneity. *Reliability Engineering & System Safety*, 145:183–189, 2016.

Haitao Liao and Zhigang Tian. A framework for predicting the remaining useful life of a single unit under time-varying operating conditions. *IIE Transactions*, 45(9): 964–980, 2013.

Linxia Liao. Discovering prognostic features using genetic programming in remaining useful life prediction. *IEEE Transactions on Industrial Electronics*, 61(5):2464–2472, 2013.

Bin Liu, Phuc Do, Benoit Iung, and Min Xie. Stochastic filtering approach for condition-based maintenance considering sensor degradation. *IEEE Transactions on Automation Science and Engineering*, 17(1):177–190, 2019a.

Bin Liu, Mahesh D Pandey, Xiaolin Wang, and Xiujie Zhao. A finite-horizon condition-based maintenance policy for a two-unit system with dependent degradation processes. *European Journal of Operational Research*, 295(2):705–717, 2021.

Haodong Liu, Zheng Liu, Liang Tu, Jinlong Liang, and Yuhao Zhang. Fatigue life prediction and reliability assessment of cracked cfrp laminates in wtbs application: A combined approach using paris-xfem and alk techniques. *IEEE Transactions on Reliability*, 2025.

Qiaobin Liu, Wenku Shi, Zhiyong Chen, Kunheng Li, Helong Liu, and Song Li. Rubber accelerated ageing life prediction by peck model considering initial hardness influence. *Polymer Testing*, 80:106132, 2019b.

Yiqi Liu, Lingling Yuan, Shuo Huang, Daoping Huang, and Bin Liu. Integrated design of monitoring, analysis and maintenance for filamentous sludge bulking in wastewater treatment. *Measurement*, 155:107548, 2020.

Yongchao Liu, Guanjun Wang, and Peng Liu. A condition-based maintenance policy with non-periodic inspection for k-out-of-n: G systems. *Reliability Engineering & System Safety*, 241:109640, 2024.

Yuehua Liu, Wenjin Yu, Wenny Rahayu, and Tharam Dillon. An evaluative study on IoT ecosystem for smart predictive maintenance (IoT-SPM) in manufacturing:

Multiview requirements and data quality. *IEEE Internet of Things Journal*, 10(13): 11160–11184, 2023.

Hongfang Lu and Xin Ma. Hybrid decision tree-based machine learning models for short-term water quality prediction. *Chemosphere*, 249:126169, 2020.

Lu Lu, Bingxing Wang, Yili Hong, and Zhisheng Ye. General path models for degradation data with multiple characteristics and covariates. *Technometrics*, 63(3):354–369, 2021.

Jianming Ma and Kara M Kockelman. Bayesian multivariate Poisson regression for models of injury count, by severity. *Transportation Research Record*, 1950(1):24–34, 2006.

James Martens et al. Deep learning via Hessian-free optimization. In *International Conference of Machine Learning*, volume 27, pages 735–742, 2010.

Sebastian Martorell, Ana Sanchez, and Vicente Serradell. Age-dependent reliability model considering effects of maintenance and working conditions. *Reliability Engineering & System Safety*, 64(1):19–31, 1999.

Sebastián Martorell, Maryory Villamizar, Sofía Carlos, and AJRE Sánchez. Maintenance modeling and optimization integrating human and material resources. *Reliability Engineering & System Safety*, 95(12):1293–1299, 2010.

Samir Mekid. IoT for health and usage monitoring systems: mitigating consequences in manufacturing under CBM. In *2021 18th International Multi-Conference on Systems, Signals & Devices (SSD)*, pages 569–574. IEEE, 2021.

Manish Mishra and Monika Srivastava. A view of artificial neural network. In *2014 International Conference on Advances in Engineering & Technology Research (ICAETR-2014)*, pages 1–3. IEEE, 2014.

Aditya Modi, Nan Jiang, Satinder Singh, and Ambuj Tewari. Markov decision processes with continuous side information. In *Algorithmic Learning Theory*, pages 597–618. PMLR, 2018.

Iqbal Muhammad and Zhu Yan. Supervised machine learning approaches: A survey. *ICTACT Journal on Soft Computing*, 5(3), 2015.

Toshio Nakagawa and Kazumi Yasui. Note on optimal redundant policies for reliability models. *Journal of Quality in Maintenance Engineering*, 11(1):82–96, 2005.

Hoang-Phuong Nguyen, Jie Liu, and Enrico Zio. Ensemble of models for fatigue crack growth prognostics. *IEEE Access*, 7:49527–49537, 2019.

Khanh TP Nguyen, Mitra Fouladirad, and Antoine Grall. Model selection for degradation modeling and prognosis with health monitoring data. *Reliability Engineering & System Safety*, 169:105–116, 2018.

Robin P Nicolai, Rommert Dekker, and Jan M Van Noortwijk. A comparison of models for measurable deterioration: An application to coatings on steel structures. *Reliability Engineering & System Safety*, 92(12):1635–1650, 2007.

Vahid Nourani, Parisa Asghari, and Elnaz Sharghi. Artificial intelligence based ensemble modeling of wastewater treatment plant using jittered data. *Journal of Cleaner Production*, 291:125772, 2021.

Tava Lennon Olsen and Brian Tomlin. Industry 4.0: Opportunities and challenges for operations management. *Manufacturing & Service Operations Management*, 22(1): 113–122, 2020.

Afshin Oroojlooyjadid, Lawrence V Snyder, and Martin Takáč. Applying deep learning to the newsvendor problem. *IISE Transactions*, 52(4):444–463, 2020.

Iwona Paprocka. The model of maintenance planning and production scheduling for maximising robustness. *International Journal of Production Research*, 57(14):4480–4501, 2019.

Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. Specaugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019.

S Patro and Kishore Kumar Sahu. Normalization: A preprocessing stage. *arXiv preprint arXiv:1503.06462*, 2015.

Chien-Yu Peng and Sheng-Tsaing Tseng. Mis-specification analysis of linear degradation models. *IEEE Transactions on Reliability*, 58(3):444–455, 2009.

Frank Proschan. Theoretical explanation of observed decreasing failure rate. *Technometrics*, 5(3):375–383, 1963.

Meng Qi, Yuanyuan Shi, Yongzhi Qi, Chenxin Ma, Rong Yuan, Di Wu, and Zuo-Jun Shen. A practical end-to-end inventory management model with deep learning. *Management Science*, 69(2):759–773, 2023.

Mengjiao Qin, Zhihang Li, and Zhenhong Du. Red tide time series forecasting by combining ARIMA and deep belief network. *Knowledge-Based Systems*, 125:39–52, 2017.

Koosha Rafiee, Qianmei Feng, and David W Coit. Condition-based maintenance for repairable deteriorating systems subject to a generalized mixed shock model. *IEEE Transactions on Reliability*, 64(4):1164–1174, 2015.

Pramod Rajput, OS Sastry, and GN Tiwari. Effect of irradiance, temperature exposure and an arrhenius approach to estimating weathering acceleration factor of glass, eva and tedlar in a composite climate of india. *Solar Energy*, 144:267–277, 2017.

Nipat Rasmekomen and Ajith Kumar Parlikad. Condition-based maintenance of multi-component systems with degradation state-rate interactions. *Reliability Engineering & System Safety*, 148:1–10, 2016.

Lei Ren, Yuxin Liu, Di Huang, Keke Huang, and Chunhua Yang. MCTAN: A novel multichannel temporal attention-based network for industrial health indicator prediction. *IEEE Transactions on Neural Networks and Learning Systems*, 34(9):6456–6467, 2022.

Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight

examples for robust deep learning. In *International Conference on Machine Learning*, pages 4334–4343. PMLR, 2018.

G. Rodríguez. Lecture notes on generalized linear models, 2007. URL `https://grodri.github.io/glms/notes/`.

Philip A Scarf, Cristiano AV Cavalcante, Richard A Dwight, and Peter Gordon. An age-based inspection and replacement policy for heterogeneous components. *IEEE Transactions on Reliability*, 58(4):641–648, 2009.

Ameneh Forouzandeh Shahraki, Om Parkash Yadav, and Haitao Liao. A review on degradation modelling and its engineering applications. *International Journal of Performability Engineering*, 13(3):299, 2017.

Moshe Shaked and J George Shanthikumar. *Stochastic orders*. Springer, 2007.

Hui Shi and Jianchao Zeng. Real-time prediction of remaining useful life and preventive opportunistic maintenance strategy for multi-component systems considering stochastic dependence. *Computers & Industrial Engineering*, 93:192–204, 2016.

Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48, 2019.

Guojin Si, Tangbin Xia, Nagi Gebraeel, Dong Wang, Ershun Pan, and Lifeng Xi. Holistic opportunistic maintenance scheduling and routing for offshore wind farms. *Renewable and Sustainable Energy Reviews*, 207:114991, 2025.

Wujun Si, Ernie Love, and Qingyu Yang. Two-state optimal maintenance planning of repairable systems with covariate effects. *Computers & Operations Research*, 92: 17–25, 2018.

IB Sidibé, Abdelhakim Khatab, Claver Diallo, and Kondo H Adjallah. Kernel estimator of maintenance optimization model for a stochastically degrading system under different operating environments. *Reliability Engineering & System Safety*, 147:109–116, 2016.

Rajesh Siraskar, Satish Kumar, Shruti Patil, Arunkumar Bongale, and Ketan Kotecha. Reinforcement learning for predictive maintenance: A systematic technical review. *Artificial Intelligence Review*, 56(11):12885–12947, 2023.

Chelsea Snyder. *The effects of charge/discharge rate on capacity fade of lithium ion batteries*. Rensselaer Polytechnic Institute, 2016.

Shagun Sodhani, Amy Zhang, and Joelle Pineau. Multi-task reinforcement learning with context-based representations. In *International Conference on Machine Learning*, pages 9767–9779. PMLR, 2021.

Kaia Stødle, Roger Flage, Seth D Guikema, and Terje Aven. Data-driven predictive modeling in risk assessment: Challenges and directions for proper uncertainty representation. *Risk Analysis*, 43:2644–2658, 2023.

John S Suehle. Ultrathin gate oxide reliability: Physical models, statistics, and characterization. *IEEE Transactions on Electron Devices*, 49(6):958–971, 2002.

Xuxue Sun, Wenjun Cai, and Mingyang Li. A hierarchical modeling approach for degradation data with mixed-type covariates and latent heterogeneity. *Reliability Engineering & System Safety*, 216:107928, 2021.

Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.

Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. In *International Conference on Artificial Neural Networks*, pages 270–279. Springer, 2018.

Yyi Kai Teoh, Sukhpal Singh Gill, and Ajith Kumar Parlikad. IoT and fog-computing-based predictive maintenance model for effective asset management in Industry 4.0 using machine learning. *IEEE Internet of Things Journal*, 10(3):2087–2094, 2021.

Mani Thenmozhi, Visalakshi Jeyaseelan, Lakshmanan Jeyaseelan, Rita Isaac, and Rupa Vedantam. Survival analysis in longitudinal studies for recurrent events: Applications and challenges. *Clinical Epidemiology and Global Health*, 7(2):253–260, 2019.

Terry M Therneau. *A Package for Survival Analysis in R*, 2023. URL `https://CRAN.R-project.org/package=survival`. R package version 3.5-5.

OWM Thijssens and Wim JC Verhagen. Application of extended Cox regression model to time-on-wing data of aircraft repairables. *Reliability Engineering & System Safety*, 204:107136, 2020.

Zhigang Tian and Ming J Zuo. Health condition prediction of gears using a recurrent neural network approach. *IEEE Transactions on Reliability*, 59(4):700–705, 2010.

Michiel AJ uit het Broek, Ruud H Teunter, Bram de Jonge, Jasper Veldman, and Nicky D Van Foreest. Condition-based production planning: Adjusting production rates to balance output and failure risk. *Manufacturing & Service Operations Management*, 22(4):792–811, 2020.

Terry T Um, Franz MJ Pfister, Daniel Pichler, Satoshi Endo, Muriel Lang, Sandra Hirche, Urban Fietzek, and Dana Kulić. Data augmentation of wearable sensor data for parkinson's disease monitoring using convolutional neural networks. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pages 216–220, 2017.

John S Usher, Ahmed H Kamal, and Wasim Hashmi Syed. Cost optimal preventive maintenance and replacement scheduling. *IIE Transactions*, 30(12):1121–1128, 1998.

Florin Vaida and Ronghui Xu. Proportional hazards model with random effects. *Statistics in Medicine*, 19(24):3309–3324, 2000.

Frank A Van der Duyn Schouten and Pieter Wartenhorst. Transient analysis of a two-unit standby system with Markovian degrading units. *Management Science*, 40(3): 418–428, 1994.

Telli Van der Lei, Paulien Herder, and Ype Wijnia. *Asset management*. Springer, 2012.

Timothy S Vaughan. Failure replacement and preventive maintenance spare parts ordering policy. *European Journal of Operational Research*, 161(1):183–190, 2005.

James W Vaupel, Kenneth G Manton, and Eric Stallard. The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography*, 16(3):439–454, 1979.

Petr Volf and Jana Timková. On selection of optimal stochastic model for accelerated life testing. *Reliability Engineering & System Safety*, 131:291–297, 2014.

Dong Wang, Kwok-Leung Tsui, and Qiang Miao. Prognostics and health management: A review of vibration based bearing and gear health indicators. *IEEE Access*, 6: 665–676, 2017.

Hongzhou Wang. A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3):469–489, 2002.

Jia Wang, Guanghan Bai, Zhigang Li, and Ming J Zuo. A general discrete degradation model with fatal shocks and age-and state-dependent nonfatal shocks. *Reliability Engineering & System Safety*, 193:106648, 2020.

Jue Wang. Optimal Bayesian demand learning over short horizons. *Production and Operations Management*, 30(4):1154–1177, 2021.

Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *2011 International Conference on Computer Vision*, pages 1457–1464. IEEE, 2011.

KS Wang, FS Hsu, and PP Liu. Modeling the bathtub shape hazard rate function in terms of reliability. *Reliability Engineering & System Safety*, 75(3):397–406, 2002.

Wenbin Wang. A stochastic model for joint spare parts inventory and planned maintenance optimisation. *European Journal of Operational Research*, 216(1):127–139, 2012.

Xiaofei Wang, Bing Xing Wang, Yili Hong, and Pei Hua Jiang. Degradation data analysis based on gamma process with random effects. *European Journal of Operational Research*, 292(3):1200–1208, 2021.

Xiaolin Wang, Narayanaswamy Balakrishnan, and Bo Guo. Residual life estimation based on a generalized Wiener degradation process. *Reliability Engineering & System Safety*, 124:13–23, 2014a.

Zhaoqiang Wang, Changhua Hu, Wenbin Wang, Zhijie Zhou, and Xiaosheng Si. A case study of remaining storage life prediction using stochastic filtering with the influence of condition monitoring. *Reliability Engineering & System Safety*, 132: 186–195, 2014b.

Zhihua Wang, Huimin Fu, and Yongbo Zhang. Analyzing degradation by an independent increment process. *Quality and Reliability Engineering International*, 30(8): 1275–1283, 2014c.

Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big Data*, 3(1):1–40, 2016.

Qingsong Wen, Liang Sun, Fan Yang, Xiaomin Song, Jingkun Gao, Xue Wang, and Huan Xu. Time series data augmentation for deep learning: A survey. *arXiv preprint arXiv:2002.12478*, 2020.

Tangbin Xia, Lifeng Xi, Xiaojun Zhou, and Jay Lee. Condition-based maintenance for intelligent monitored series system with independent machine failure modes. *International Journal of Production Research*, 51(15):4585–4596, 2013.

Tangbin Xia, Xiaoning Jin, Lifeng Xi, and Jun Ni. Production-driven opportunistic maintenance for batch production based on MAM–APB scheduling. *European Journal of Operational Research*, 240(3):781–790, 2015.

Tangbin Xia, Ya Song, Yu Zheng, Ershun Pan, and Lifeng Xi. An ensemble framework based on convolutional bi-directional LSTM with multiple time windows for remaining useful life estimation. *Computers in Industry*, 115:103182, 2020.

Huan Yan, Peitang Wei, Pengliang Zhou, Ling Chen, Huaiju Liu, and Caichao Zhu. Experimental investigation of crack growth behaviors and mechanical properties degradation during gear bending fatigue. *Journal of Mechanical Science and Technology*, 36(3):1233–1242, 2022.

Ang Yang, Yu Wang, Fangfang Yang, Dong Wang, Yanyang Zi, Kwok Leung Tsui,

and Bin Zhang. A comprehensive investigation of lithium-ion battery degradation performance at different discharge rates. *Journal of Power Sources*, 443:227108, 2019.

Zhi-Sheng Ye and Min Xie. Stochastic modelling and analysis of degradation for highly reliable products. *Applied Stochastic Models in Business and Industry*, 31(1):16–32, 2015.

Zhi-Sheng Ye, Yu Wang, Kwok-Leung Tsui, and Michael Pecht. Degradation data analysis using Wiener processes with measurement errors. *IEEE Transactions on Reliability*, 62(4):772–780, 2013.

Zhisheng Ye and Nan Chen. The inverse Gaussian process as a degradation model. *Technometrics*, 56(3):302–311, 2014.

Meng-Lai Yin, John E Angus, and Kishor S Trivedi. Optimal preventive maintenance rate for best availability with hypo-exponential failure distribution. *IEEE Transactions on Reliability*, 62(2):351–361, 2013.

Yong Yu, Xiaosheng Si, Changhua Hu, and Jianxun Zhang. A review of recurrent neural networks: LSTM cells and network architectures. *Neural Computation*, 31(7): 1235–1270, 2019.

Kimiya Zakikhani, Fuzhan Nasiri, and Tarek Zayed. A review of failure prediction models for oil and gas pipelines. *Journal of Pipeline Systems Engineering and Practice*, 11(1):03119001, 2020.

Hongtao Zeng, Tian Lan, and Qiming Chen. Five and four-parameter lifetime distributions for bathtub-shaped failure rate using perks mortality equation. *Reliability Engineering & System Safety*, 152:307–315, 2016.

Fan Zhang and Andrew KS Jardine. Optimal maintenance models with minimal repair, periodic overhaul and complete renewal. *IIE transactions*, 30(12):1109–1119, 1998.

Fengxia Zhang, Haitao Liao, Jingyuan Shen, and Yizhong Ma. Optimal maintenance over a finite time horizon for a system under imperfect inspection and dynamic working environment. *Reliability Engineering & System Safety*, 246:110057, 2024a.

Liang Zhang, Jin Wen, Yanfei Li, Jianli Chen, Yunyang Ye, Yangyang Fu, and William Livingood. A review of machine learning in building load prediction. *Applied Energy*, 285:116452, 2021.

Linmiao Zhang, Yong Lei, and Houcai Shen. How heterogeneity influences condition-based maintenance for gamma degradation process. *International Journal of Production Research*, 54(19):5829–5841, 2016.

Qin Zhang, Yu Liu, Boyuan Zhang, and Hong-Zhong Huang. Selective maintenance optimization under limited maintenance capacities: A machine learning-enhanced approximate dynamic programming. *IEEE Transactions on Reliability*, 2024b.

Qing Zhang, Cheng Hua, and Guanghua Xu. A mixture Weibull proportional hazard model for mechanical system failure prediction utilising lifetime and monitoring data. *Mechanical Systems and Signal Processing*, 43(1-2):103–112, 2014.

Zhengxin Zhang, Xiaosheng Si, Changhua Hu, and Yaguo Lei. Degradation data analysis and remaining useful life estimation: A review on Wiener-process-based methods. *European Journal of Operational Research*, 271(3):775–796, 2018.

Xiujie Zhao, Olivier Gaudoin, Laurent Doyen, and Min Xie. Optimal inspection and replacement policy based on experimental degradation data with covariates. *IISE Transactions*, 51(3):322–336, 2019.

Xufeng Zhao, Khalifa N Al-Khalifa, Abdel Magid Hamouda, and Toshio Nakagawa. Age replacement models: A summary with new perspectives and methods. *Reliability Engineering & System Safety*, 161:95–105, 2017.

Xiaojun Zhou, Lifeng Xi, and Jay Lee. Opportunistic preventive maintenance scheduling for a multi-unit series system based on dynamic programming. *International Journal of Production Economics*, 118(2):361–366, 2009.

Yapeng Zhou, Miaohua Huang, Yupu Chen, and Ye Tao. A novel health indicator for on-line lithium-ion batteries remaining useful life prediction. *Journal of Power Sources*, 321:1–10, 2016.