A data-driven approach for condition-based maintenance optimization

Cai, Yue; Teunter, Ruud; de Jonge, Bram

Published in:
European Journal of Operational Research

DOI:
10.1016/j.ejor.2023.05.002

European Journal of Operational Research 311 (2023) 730–738

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

A data-driven approach for condition-based maintenance optimization

Yue Cai[a,b,*], Ruud H. Teunter[a], Bram de Jonge[a]

[a] Department of Operations, Faculty of Economics and Business, University of Groningen, P.O. Box 800, Groningen 9700 AV, the Netherlands
[b] School of Economics and Management, University of Chinese Academy of Sciences, Beijing 100190, China

ARTICLE INFO

Article history:
Received 8 November 2022
Accepted 1 May 2023
Available online 19 May 2023

Keywords:
Reliability
Condition-based maintenance
Data-driven
Deterioration process
Control limit policy

ABSTRACT

Developments in sensor techniques enable the continuous monitoring of the health of an operating system. The resulting condition data provides an opportunity for better prediction of failures and thereby for improving maintenance decisions. In this study, we consider condition-based maintenance for a single unit with an unknown, non-decreasing deterioration process and unknown failure behavior. Building on, but different from the existing maintenance optimization literature, we present the first fully data-driven approach, where the condition threshold triggering maintenance is based purely on past condition data and failures. Numerical results for a gamma deterioration process show that the maintenance threshold resulting from our data-driven approach converges to the optimal threshold. The threshold is set higher during the initial runs-to-failure, and this helps to explore the deterioration process. An encouraging result is that the convergence is especially fast during the first few runs-to-failure so that the expected cost rate quickly converges to the minimum cost rate.

1. Introduction

Maintenance activities are an integral part of the production strategy, and increasingly important for modern manufacturing using complex and expensive critical assets. Savings of 20–30% of the total operational costs can be realized by optimizing maintenance activities (Syan & Ramsoobag, 2019). Therefore, effective maintenance strategies are essential to ensure high system availability and consistent product quality, thereby increasing profitability and competitiveness of manufacturing firms.

Systems typically experience gradual and irreversible deterioration that can be observed over time (Elwany et al., 2011). Based on this system condition, condition-based maintenance (CBM) can be conducted to decide whether maintenance operations should be implemented, which has received great interest from both academia and practice. Essentially, the goal of CBM is to maintain the system or component at the exact right time using the actual health status of the system or component. The rapid development of sensor techniques makes it possible to identify the actual state of the system by collecting deterioration data. Moreover, the collection of these data is increasingly affordable due to the advances in deployment of sensors (Liu et al., 2017; Walter & Flapper, 2017). These developments have led to increased popularity and use of CBM, which is considered to be superior to traditional time-based maintenance (De Jonge et al., 2017; Liu et al., 2021a). When maintenance for a single unit with a single deterioration indicator is considered, a control-limit policy is typically applied. Under such a policy, maintenance activities are carried out when the deterioration level reaches or exceeds some preventive maintenance threshold (De Jonge, 2021). This policy is simple to implement in practice and is generally optimal when there is no noise in the condition measurements (De Jonge et al., 2017). In more complicated settings it can often be shown that optimal policies have a threshold structure by formulating the problem as a Markov decision process (Drent et al., 2023; Kurt & Kharoufeh, 2010; Liu et al., 2021a; 2017; Zhang & Zhang, 2022).

Generally, the operation of CBM consists of three connected steps: condition data acquisition, reliability estimation and decision making (Liu et al., 2017). Obviously, interpreting condition data, translating it into reliability estimations, and integrating it into maintenance decision making is difficult, most of all at the start of using new equipment when limited data is available. Interestingly and perhaps surprisingly, there are very few studies that have considered maintenance decision making given limited data. Nguyen & Medjaher (2019) propose a data-driven maintenance framework and consider the complete process from reliability estimation to maintenance decision making, using deep learn-

* Corresponding author at: Department of Operations, Faculty of Economics and Business, University of Groningen, P.O. Box 800, Groningen 9700 AV, the Netherlands.
E-mail address: yue.cai@rug.nl (Y. Cai).

https://doi.org/10.1016/j.ejor.2023.05.002
0377-2217/

ing methods. This method can only be applied if large amounts of historical data are available. To the best of our knowledge, all CBM research for new equipment so far assumes that the type of deterioration process is known exactly, e.g., a gamma process, and that the exact failure level is known. In this paper, we consider a new approach where both the reliability estimation and maintenance decision making are data driven. Furthermore, this approach only requires limited data.

As we are the first to take such a ‘full' data driven approach, we consider a basic system with a single unit and a single measurable condition. We measure the deterioration level periodically, and apply maintenance according to the control-limit policy. Instead of assuming a specific deterioration process such as a gamma process, we directly estimate and continuously update the failure probability (reliability) as new condition and failure data becomes available. Subsequently, we propose a data-driven maintenance decision-making rule that aims to minimize the long-run average cost per time period, consisting of corrective maintenance and preventive maintenance costs. Furthermore, we numerically (for a gamma deterioration process) explore the impact of the amount of data on the performance of the data-driven approach to investigate the method's applicability and practicality. The results are insightful and encouraging. The data-driven method converges to the optimal maintenance threshold quickly, especially during the first ten runs-to-failure, and uses a larger maintenance threshold during that initial phase which facilitates learning the failure process quicker.

The remainder of this paper is organized as follows. In Section 2, related literature is discussed. In Section 3, we introduce the problem formally. In Section 4, an optimality condition is derived for the case with a known deterioration process. In Section 5, this is used to develop our data-driven approach. In Section 6, a numerical analysis is presented to illustrate the performance of the proposed approach. Finally, we conclude and provide suggestions for future research in Section 7.

2. Literature review

Much research has been devoted to maintenance optimization based on system conditions (De Jonge & Scarf, 2020). This is known as condition-based maintenance (CBM). Typically, deterioration processes represent changing physical phenomena over time, which can be observed by monitoring the state of the system during operation (Elwany & Gebraeel, 2008). These physical transitions observed by monitoring sensors, such as accelerated vibration and temperature variations, can fully capture the underlying health state of a focal system (Elwany et al., 2011).

Deterioration processes can be modeled by using either a discrete set of deterioration states, or a continuous set. In the discrete case, Markov chains are widely used to describe the deterioration process (e.g., Andersen et al., 2022; Kurt & Kharoufeh, 2010; Zhang & Zhang, 2022). The Markov chain models are particularly effective in the case that deterioration states of a unit or system cannot be fully captured (Chen et al., 2015). Examples of deterioration processes with a continuous state are the Wiener process, gamma process, and inverse Gaussian process (Alaswad & Xiang, 2017). The Wiener process is appropriate in describing non-monotonic deterioration (Liu et al., 2017). However, for many systems deterioration is irreversible, i.e., the deterioration path is monotone such as for crack growth and wears (Chen et al., 2015). For such situations, the gamma process and the inverse Gaussian process (Chen et al., 2015; Fang et al., 2022) are more appropriate as they are non-decreasing. The gamma process is the most common choice for modeling monotonic deterioration due to its physical interpretation (Van Noortwijk, 2009). Continuous-state deterioration processes can be discretized for analytical tractability (Andersen et al., 2022; De Jonge, 2019; Liu et al., 2021a; Nguyen et al., 2019).

Most of the existing studies on condition-based maintenance optimization make the unrealistic assumption that the deterioration process is fully known. Some studies do assume an unknown or uncertain deterioration process though. In the remainder of this review we review these, and we conclude this section by explaining how the current study contributes to this stream of research.

Elwany & Gebraeel (2008) consider the optimal moment to order a new spare component as well as the optimal replacement time. Deterioration of the component is modeled using some functional form (linear or exponential) with one or more uncertain parameters. Because of measurement noise according to a Brownian motion, the parameters of the deterioration process remain uncertain. This uncertainty is modeled using a prior distribution and is updated in a Bayesian manner. Elwany et al. (2011) consider the exponential functional form and provide structural properties of the optimal preventive replacement policy using a Markov decision process formulation.

Si et al. (2018) consider deterioration according to a Wiener process with an unknown drift parameter. The unit is inspected periodically, and based on the observed deterioration levels the maximum likelihood estimation of the drift parameter is determined. Based on this estimation and on the observed deterioration level at an inspection, it is decided whether or not preventive maintenance is carried out.

Flage et al. (2012) consider a single unit with a deterioration process with uncertain parameters. Inspections are needed to learn the deterioration level, and are scheduled sequentially. A linear function is used to determine the time until the next inspection, based on the current deterioration level. Preventive maintenance is carried out if the deterioration level exceeds some threshold. Failures are observed immediately and are followed by corrective maintenance. Again, the uncertainty in the unknown parameters is modeled using a prior distribution, and is updated in a Bayesian manner. Mosayebi Omshi et al. (2020) consider a similar setting, but assume that failures are not self-announcing. They schedule the next inspection such that the probability of failure before this inspection is below some threshold.

Walter & Flapper (2017) consider a complex multi-unit system and its reliability block diagram. Each unit is either working or defective, but the combined states of all units together is considered as the condition of the system. All units have a Weibull lifetime distribution with an unknown scale parameter. A Bayesian approach is used to update the distributions that model this uncertainty. Expert knowledge is incorporated in the prior distribution.

Uncertainty in the deterioration process also exists if deterioration is influenced by an unobservable environment or if a heterogeneous set of spare components is used. Fouladirad et al. (2008) consider a unit that deteriorates according to a continuous-state stochastic process and that is inspected periodically. The unit operates in a dynamic environment and the real environment must be inferred from inspection data. Van Oosterom et al. (2017) consider a single-unit system that deteriorates according to a discrete-time Markov chain. Replacement components come from a heterogeneous population comprising multiple component types that differ in their transition probability matrix. The type of the current component should be inferred from the observed deterioration information.

Summarizing, existing studies that consider an uncertain deterioration process typically consider a specific parametric form with uncertainty in its parameters. A prior distribution is then used to model the initial uncertainty in these parameters, and Bayesian learning is applied to update these uncertainties. When the environment or the type of a spare unit is considered, it is assumed that the deterioration process is known given the environment or the spare unit type. In contrast, the main contribution of our study is developing a purely data-driven approach to specify a replacement rule for a single-unit system. The probability of system failure is estimated only based on limited observed deterioration data. Based on this, a control-limit (replacement threshold) is determined by these estimated probabilities of failures as a function of the deterioration level. This enables significant use of condition data observed by sensors and allows maintenance decisions to depend on the actual health state of the focal system. Additionally, few studies consider the impact of condition data volume on the obtained control-limit performance. The present paper is also intended to bridge this gap to explore how the data size affects the result.

3. Problem description

We consider condition-based maintenance for a single-unit system. The state of the system is assumed to be fully captured by a single condition parameter, also referred to as the deterioration level of the system. We use the term 'condition data' to refer to a set of observed/measured deterioration levels. The system deteriorates according to a stochastic deterioration process with a continuous state space. The stochastic deterioration process is unknown. We only assume that it has the Markov property, meaning that the future development of the process only depends on the current deterioration level and not on the history, and that it is non-decreasing, meaning that the condition cannot improve without maintenance. Commonly used processes for deterioration modeling such as the gamma process and the inverse Gaussian process satisfy these properties. We let deterioration level 0 represent the as-good-as-new state. The deterioration level of the system is observed at discrete, equally spaced moments in time, and for each deterioration level there is a probability of failure in the next time step. These probabilities are also unknown; we only make the intuitive assumption that the failure probability is increasing in the deterioration level. If a failure occurs, then no deterioration level is observed anymore and the system is just identified as failed. Based on the observed deterioration levels and failures, the deterioration and failure behavior of the system need to be estimated, and maintenance decisions need to be made.

We distinguish two types of maintenance actions, preventive maintenance and corrective maintenance. Corrective maintenance is required after failure of the system. Preventive maintenance can be carried out at any time before the system has failed to avoid unexpected in-service failure. Typically, the cost incurred by corrective maintenance c_{cm} is higher than the cost of preventive maintenance c_{pm}. Both maintenance types are assumed to be instantaneous, i.e., they require negligible time, and perfect, i.e., they bring the deterioration level back to 0.

The condition-based maintenance policy that we consider is the so-called control-limit policy. This policy has a single decision variable, the preventive maintenance threshold M. When an observed deterioration level exceeds this threshold M, preventive maintenance is carried out. It is intuitive that when preventive maintenance is carried out at a certain deterioration level, it is also carried out at higher deterioration level, because this implies a further deteriorated system and a higher chance of failure. The threshold M should not be chosen too low as that results in too frequent maintenance, but it should also not be chosen too high as that results in too many failures.

Given the deterioration process, the failure behavior, and the cost parameters c_{pm} and c_{cm} of the system, the long-run cost rate $\eta(M)$ as a function of the preventive maintenance threshold M can be determined and minimized. However, in this study we make the more realistic assumption that the deterioration process and failure behavior are unknown and assume that the preventive maintenance threshold M needs to be specified only based on limited observed condition data and past failures. We refer to a run-to-failure as one cycle from the system being as-good-as-new until its failure. Such a run-to-failure results in a sequence of observed deterioration levels. Typically, if we have condition data for more runs-to-failure, we are able to make better maintenance decisions. In this study we also investigate how the quality of the decisions is influenced by the number of runs-to-failure for which we have condition data.

4. Optimality condition for known deterioration process

Our approach for specifying a preventive maintenance threshold M based on condition data and past failures builds on De Jonge (2019, 2021). These studies consider condition-based maintenance in a discrete-time setting for a unit that deteriorates according to a known Markov chain with m deterioration states before failure and failure state $m+1$. The $(m+1)\times(m+1)$ transition probability matrix of the Markov chain is denoted by P, which can be written as

$$P = \begin{pmatrix} Q & r \\ 0 & 1 \end{pmatrix}.$$

Herein, r is a column vector of length m with failure probabilities. The corresponding fundamental matrix is $R = \sum_{k=0}^{\infty} Q^k = (I_m - Q)^{-1}$, with I_m the $m \times m$ identity matrix. If the process starts in deterioration state 1 (i.e., after maintenance is carried out), $R_{1,j}$ is the expected number of time periods that it is in state $j \in \{1,\dots,m\}$ before reaching a state higher than j.

At the moment that the deterioration level exceeds the preventive maintenance threshold M maintenance is carried out. Maintenance is corrective if the deterioration level has jumped from a state below M to failure state $m+1$. For any preventive maintenance threshold $M \in \{1,\dots,m\}$, the expected number of time periods until maintenance equals $D(M) = \sum_{j<M} R_{1,j}$. Furthermore, the probability that maintenance is corrective is $\sum_{j<M} R_{1,j} r_j$, implying that the mean cost per maintenance action equals $C(M) = c_{\text{pm}} + (c_{\text{cm}} - c_{\text{pm}}) \sum_{j<M} R_{1,j} r_j$. Based on renewal reward theory it follows that the cost rate $\eta(M)$ as the function of the preventive maintenance threshold M equals

$$\eta(M) = \frac{C(M)}{D(M)} = \frac{c_{\text{pm}} + (c_{\text{cm}} - c_{\text{pm}}) \sum_{j<M} R_{1,j} r_j}{\sum_{j<M} R_{1,j}}. \tag{1}$$

We refer to De Jonge (2019, 2021) for more details regarding this analysis.

4.1. Optimality condition

Whereas De Jonge (2019, 2021) only determines optimal preventive maintenance thresholds M^* numerically by comparing the cost rate $\eta(M)$ for all values of M, we continue to derive a closed-form optimality condition for M^* based on the cost rate expression. The analysis in Section 5 for the data-driven situation builds on this optimality condition. Because failure is more likely at a higher deterioration level, we have that $r_{M+1} > r_M$ for any M. From this it easily follows that $\eta(M)$ is unimodal, i.e., $\eta(M+1) > \eta(M)$ implies $\eta(M+2) > \eta(M+1)$ for any M. Specifically, we have $\eta(M+1) > \eta(M)$ if and only if

$$\sum_{j<M} R_{1,j}(r_M - r_j) > \frac{c_{\text{pm}}}{c_{\text{cm}} - c_{\text{pm}}},$$

and we have $\eta(M+2) > \eta(M+1)$ if and only if

$$\sum_{j<M+1} R_{1,j}(r_{M+1} - r_j) > \frac{c_{\text{pm}}}{c_{\text{cm}} - c_{\text{pm}}},$$

or, equivalently, if and only if

$$\sum_{j<M} R_{1,j}\left(r_{M+1}-r_j\right)+R_{1,M}(r_{M+1}-r_M) > \frac{c_{\mathrm{pm}}}{c_{\mathrm{cm}}-c_{\mathrm{pm}}}.$$

Since $r_{M+1} > r_M$, we have that $\eta(M+2) > \eta(M+1)$ holds when $\eta(M+1) > \eta(M)$. This implies that $\eta(M)$ is indeed unimodal, and therefore that the optimal threshold M^* is given by

$$M^* = \min\{M : \eta(M+1) > \eta(M)\}.$$

Based on the above, we derive

$$M^* = \min\left\{M : \sum_{j<M} R_{1,j}(r_M - r_j) > \frac{c_{\mathrm{pm}}}{c_{\mathrm{cm}}-c_{\mathrm{pm}}}\right\}. \tag{2}$$

Obviously, this cannot be applied directly to the data-driven situation, because there (a) we assume a continuum of deterioration levels rather than a finite number of levels, and (b) we assume the failure probability as a function of the deterioration level to be unknown.

5. Data-driven approach

We will now present our data-driven approach. After introducing some notation, we adapt Eq. (2) to an equation that can be used in our data-driven framework, and then derive the estimated failure probabilities (for all observed deterioration levels) needed to implement it. Let us assume data for K runs-to-failure is available. Each run-to-failure contains a number of observations over time, consisting of a deterioration level and whether or not that resulted in a failure in the next time step. Let us denote the total number of observations for all runs-to-failure together by n. We sort all observed deterioration levels in ascending order and let them be denoted by $x_1, x_2, \ldots, x_n$. Furthermore, we let the binary variables $y_1, y_2, \ldots, y_n \in \{0,1\}$ indicate whether a deterioration observation has resulted in a failure in the next time step (1) or not (0).

As an example, let us assume that we have data for 2 runs-to-failure, namely '0–5.2–6.4-Failure' and '0–4.6–6.8–7.2-Failure'. This means that we have $n = 7$ observed deterioration levels, namely $x_1 = x_2 = 0$, $x_3 = 4.6$, $x_4 = 5.2$, $x_5 = 6.4$, $x_6 = 6.8$, $x_7 = 7.2$, and that we have failure indicators $y_1 = y_2 = y_3 = y_4 = y_6 = 0$ and $y_5 = y_7 = 1$. For deterioration processes with a continuous state space (such as the commonly used gamma process), each x_i typically has a unique value, except for deterioration level 0 that we observe at the beginning of each cycle. This also applies to the presented example.

5.1. Adjusted optimality condition

The estimated failure probability in the next time step for each deterioration level x_i is denoted as $r(x_i)$. In Section 5.2 we will explain how we will obtain these estimations based on the data. Besides the estimated failure probability $r(x_i)$, we also need the estimated number of time periods per run-to-failure that each deterioration state is visited, which is denoted by $R_{1,j}$ in the matrix algebra approach. Because we have data for K runs-to-failure, and because each deterioration state x_i is observed in one of them, we empirically estimate the number of times per run-to-failure that each observed deterioration state x_i is visited as $\frac{1}{K}$. Subsequently, the cost rate function $\eta(M)$ in (1) can be adjusted into a data-driven formula $\bar{\eta}(M)$ as

$$\bar{\eta}(M) = \frac{c_{\mathrm{pm}} + \left(c_{\mathrm{cm}}-c_{\mathrm{pm}}\right)\frac{1}{K}\sum_{i:x_i<M} r(x_i)}{\frac{1}{K}\sum_i I_{x_i<M}}, \quad M \in \{x_1, \ldots, x_n\},$$

in which I_A is the indicator function that equals 1 if A is true, and 0 if not. Note that we choose to let $\bar{\eta}(M)$ only be defined for preventive maintenance thresholds M equal to one of the observed deterioration levels $x_1, \ldots, x_n$. Other values for M could also be plugged in, but extending the domain of $\bar{\eta}(M)$ to $\mathbb{R}_+$ would result in a piece-wise constant function. Because this is not realistic and because we do not have knowledge about the behavior of the deterioration process in between the observed levels of deterioration, we choose to restrict the domain of $\bar{\eta}(M)$ to $\{x_1, \ldots, x_n\}$. The preventive maintenance threshold $\bar{M}^*$ can be determined through the corresponding optimality condition, i.e.,

$$\bar{M}^* = \min\left\{M \in \{x_1, \ldots, x_n\} : \frac{1}{K}\sum_{i:x_i<M}(r(M)-r(x_i)) > \frac{c_{\mathrm{pm}}}{c_{\mathrm{cm}}-c_{\mathrm{pm}}}\right\}, \tag{3}$$

in which $r(x)$ equals the estimated probability of failure in the next time step when the current deterioration level is x. In Section 6 we investigate how the quality of a static preventive maintenance threshold $\bar{M}^*$ depends on the number of runs-to-failure K. When using the approach dynamically, the preventive maintenance threshold $\bar{M}^*$ can be updated after each new observation of a deterioration level.

5.2. Failure probability estimation

We next develop a way to estimate the failure probability in the next time period $r(x_i)$ given an observed deterioration level x_i. One could use the 'empirical' estimate by setting $r(x_i) = y_i$, but this would lead to failure probability estimates of either 0 or 1, where each estimate would be based on a single observation. Moreover, this reliability estimate as a function of the deterioration level would lead to strange up-and-down patterns. Therefore, we opt for a different approach, namely to find the function $r : \{x_1, x_2, \ldots, x_n\} \to [0,1]$ that maximizes the likelihood of the observations, but under the intuitive restriction that r is non-decreasing, i.e., that a higher deterioration level implies a higher failure probability. We remark that for ease of presentation, in what follows we simply refer to maximum likelihood (ML) estimates without repeating this restriction.

To that end, some notations and definitions are introduced first. All observations are grouped into blocks B_i, $i = 1, 2, \ldots$ The first block B_1 contains all observations before the first one that resulted in a failure (in the next time step). The last block starts in the time step after the last non-failure observation. The blocks in between start with a failure observation after a non-failure observation and end before that next happens. This is illustrated Fig. 1. If the underlying deterioration process allows the same deterioration level to be observed multiple times, we let the failure observations precede the equal non-failure observations, meaning that these observations of the same deterioration level will be in the same block. Additionally, a number of consecutive blocks can be combined into an interval I_i, as also illustrated in Fig. 1, but an interval can also consist of a single block.

The ML estimates are 0 for all observations in the first block and 1 for all observations in the last block, since the likelihood of those observations is then 1 and this does not restrict the estimates for all other observations. The ML estimates for the other observations can vary between 0 and 1.

In what follows, we will show that ML estimates are always the same for observations within a block, and sometimes the same for consecutive blocks, which is why we defined these as intervals. Lemma 5.1 states that if we keep the failure probability estimate in an interval constant, then the joint likelihood for all observations in that interval is maximized when that estimate is equal to the fraction of failures for that interval. Its proof and proofs of lemmas that will follow can be found in Appendix A.

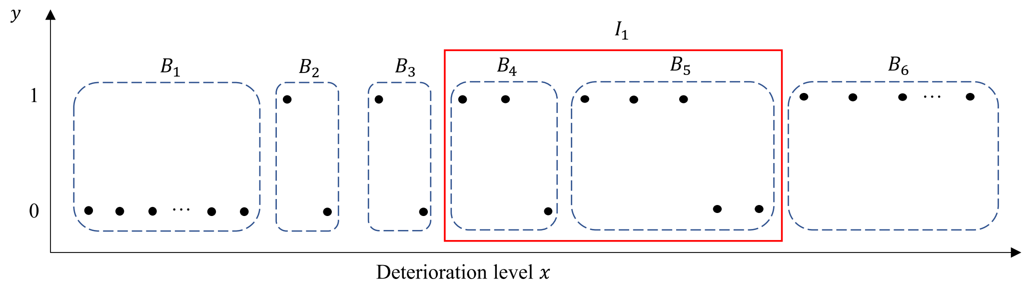

Fig. 1. Illustration of blocks and an interval.

Lemma 5.1. *If the failure probability is kept constant at value p for all observations of an interval with u failures and v non-failures, their joint likelihood is maximized by* $p = \frac{u}{u+v}$.

Next, Lemma 5.2 states that observations in the same block have the same ML estimate. The intuitive explanation is that the observed likelihoods within a block are non-increasing and that only non-decreasing estimates are allowed.

Lemma 5.2. *Any sequence of ML estimates has the same probability for all observations in a block.*

From Lemma 5.2, we know that only sequences of estimates with a constant probability per block need to be considered. Moreover, from Lemma 5.1 we have that the probability equal to the ratio of failures maximizes the likelihood. Using the same intuition as before, if a block has a lower ratio of failures than the previous block, then it must be best to set the same constant probability for both blocks, as the probabilities are required to be non-decreasing. In this way, we can extend the constant-probability property from a single block to a two-block interval. Indeed, the same logic is applied to combine larger intervals, which is formalized in Lemma 5.3.

Lemma 5.3. *Consider two consecutive intervals (I_1 and I_2) that each have constant ML estimates. If the failure ratios (r_1 and r_2) of these two intervals are decreasing ($r_1 > r_2$), then the ML estimates for both intervals are the same, implying that they can be combined to a single interval.*

Lemma 5.3 can be applied repeatedly, until a set of intervals is obtained with non-decreasing failure ratios. From Lemma 5.1, we then have that the joint likelihood is maximized when the probabilities for observations in an interval are equal to the failure ratio for that interval.

So, the ML solution is found by repeatedly combining adjacent intervals with decreasing failure ratios until we have a solution with non-decreasing failure ratios. Indeed, it is obvious that if more than two consecutive intervals have decreasing failure ratios, then we can combine them in one step. In fact, the procedure of finding the optimal solution can be further sped up by simultaneously combining all subsequences of intervals with decreasing probabilities. This leads to the following procedure for finding the optimal solution:

Step 1: Determine the failure ratio for each (new) block/interval.
Step 2: Search for consecutive intervals with decreasing probabilities and combine these into new intervals. Go back to Step 1 if intervals were indeed combined. Otherwise go to Step 3.
Step 3: Calculate the failure ratio for each interval, and set the corresponding failure probability per deterioration observation.

Finally, we get a non-decreasing function r, which is a mapping from the set of observed deterioration levels $\{x_1, x_2, \ldots, x_n\}$ to the set of probabilities [0,1].

We next illustrate this procedure for the example shown in Fig. 1. Based on the definition of a block, all observations in this example are grouped into 6 blocks with respective failure ratios 0, $\frac{1}{2}$, $\frac{1}{2}$, $\frac{2}{3}$, $\frac{3}{5}$, and 1. Because the failure ratio of the fifth block is lower than that of the fourth block, these two blocks are combined to an interval with failure ratio $\frac{5}{8}$. The result is an increasing sequence of failure ratios, thus the procedure stops. The estimated failure probability $r(x)$ for each observed deterioration level x is as shown in Fig. 2.

6. Numerical analysis

We continue to analyze the performance of our approach and explore the impact of data size on it. We first briefly introduce the gamma deterioration process, which we use to simulate deterioration data. Subsequently, a base case is explored, where the maintenance policy determined by our data-driven approach is compared with the benchmark policy that assumes that the deterioration process is known. After that, a sensitivity analysis is carried out to test the stability and robustness of our approach.

6.1. Deterioration process

We use the gamma process to model deterioration, which is a non-decreasing process that is rather flexible and applicable to model a wide variety of deterioration behaviors (Uit Het Broek et al., 2021; Van Noortwijk, 2009). We assume the gamma process to start at level 0, and during each time period with unit length the additional amount of deterioration follows a gamma distribution with shape parameter α, scale parameter β, and density function

$$f_{\alpha,\beta}(x) = \frac{x^{\alpha-1}\exp\left(\frac{-x}{\beta}\right)}{\Gamma(\alpha)\beta^{\alpha}},$$

where $\Gamma(\alpha) = \int_0^\infty x^{\alpha-1}\exp(-x)\,dx$ denotes the gamma function. We assume that failure occurs when a certain fixed failure threshold L is exceeded. The foregoing defines a stationary gamma process (increments are independent of the current deterioration level); later on we also consider non-stationary gamma processes in which we let the shape parameter depend on the current deterioration level.

6.2. Base case system

In the base case, we consider a stationary gamma process with $\alpha = 4$, $\beta = 2$, and $L = 100$. This results in a mean time to failure of 12.62, which means there are around 12 to 13 observations before failure occurs. Additionally, the preventive maintenance cost is set to $c_{pm} = 1$ and the corrective maintenance cost to $c_{cm} = 5$. We

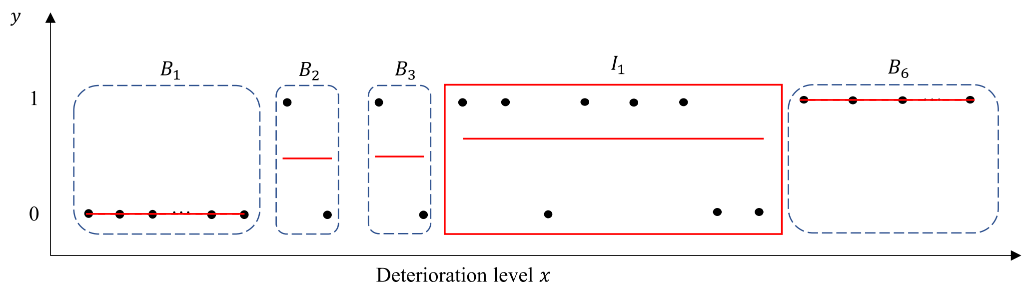

Fig. 2. The resulting estimated failure probability for each block/interval.

Table 1
Parameter values in the base case.

Parameter	Value	Interpretation
α	4	Shape parameter
β	2	Scale parameter
L	100	Failure threshold
c_{pm}	1	Preventive maintenance cost
c_{cm}	5	Corrective maintenance cost
K	$1, \ldots, 100$	Number of runs-to-failure

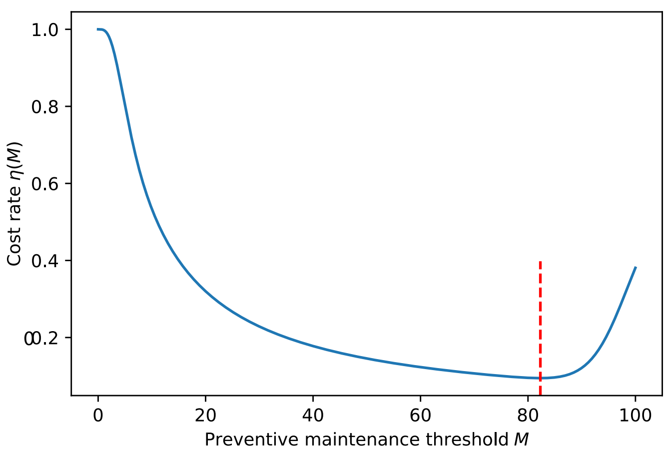

Fig. 3. Cost rate $\eta(M)$ as a function of the preventive maintenance threshold M.

consider the performance of our approach given condition data for $K = 1$ until 100 runs-to-failure. The parameter values in the base case are summarized in Table 1.

6.2.1. Benchmark policy

The discretization approach and matrix algebra approach provided by De Jonge (2019, 2021) are used to determine the long-run cost rate given any preventive maintenance threshold M (based on a discretization with $m = 20,000$ states), given that the deterioration process is known. The resulting cost rate function is shown in Fig. 3. The optimal preventive maintenance threshold is $M^* = 82.30$, and the resulting cost rate equals $\eta(M^*) = 0.0946$. This 'oracle' optimal policy is taken as our benchmark.

6.2.2. Data-driven policy

We continue to assess the performance of our data-driven approach as described in Section 5. Fig. 4 shows an example of the estimated failure probability $r(x)$ as a function of the deterioration level x based on simulated data for $K = 50$ runs-to-failure. Although the function $r(x)$ is only defined at deterioration levels x that we have observed, we provide its graph as a step function for clarity. The estimated failure probability is 0 for low deterioration levels (at which no failures have occurred), and increases

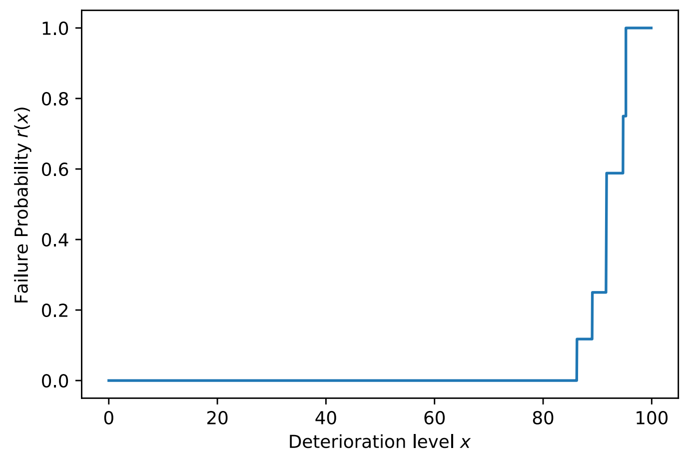

Fig. 4. Estimated failure probability $r(x)$ as a function of the deterioration level x for $K = 50$ runs-to-failure.

quite rapidly to 1 at higher deterioration levels. This is consistent with the natural deterioration mechanism of a system with lower probability of failing when the system is in a good state and high probability of failing in severe states.

Based on data of K runs-to-failure, the resulting preventive maintenance threshold $\bar{M}_K^*$ follows from the optimality criterion (3). The corresponding cost rate equals $\eta(\bar{M}_K^*)$. To get a clear impression of the performance of our approach, we repeat the simulation approach for getting $\bar{M}_K^*$ and $\eta(\bar{M}_K^*)$ 1000 times for each number of runs-to-failure K. We evaluate the performance of our approach by considering means and 95% confidence intervals (2.5th–97.5th percentile) of both the preventive maintenance threshold $\bar{M}_K^*$ and the corresponding cost rate $\eta(\bar{M}_K^*)$. The confidence intervals indicate the volatility in the decisions and in the corresponding costs when our data-driven approach is applied.We note that, to fairly investigate the effect of the amount of data on the performance of our approach, we consider the long-run cost rate of choosing a fixed preventive maintenance threshold $\bar{M}_K^*$ based on data for K runs-to-failure, and repeat this for all values of K.

Fig. 5 shows the effect of K on $\bar{M}_K^*$ and on $\eta(\bar{M}_K^*)$. On average, the preventive maintenance threshold $\bar{M}_K^*$ resulting from our approach is higher than the optimal preventive maintenance threshold. This results from the fact that the estimated failure probability $r(x)$ makes sudden steps upwards at observed deterioration levels, whereas the real failure probability function would already gradually increase towards these levels. According to the optimality condition (3), this causes the preventive maintenance threshold to also move up to an observed deterioration level. For a small number of runs-to-failure K, both the preventive maintenance threshold $\bar{M}_K^*$ and the resulting cost rate $\eta(\bar{M}_K^*)$ are likely to be far from the optimal values. However, they both converge to the optimal values

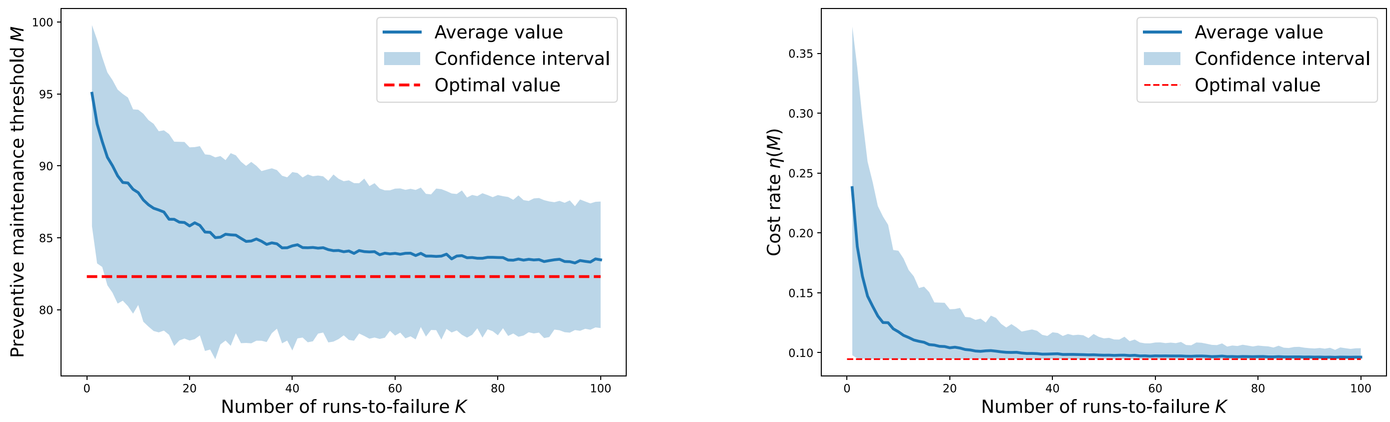

(a) $\bar{M}_K^*$ (b) $\eta(\bar{M}_K^*)$

Fig. 5. Effect of the number of runs-to-failure K on the preventive maintenance threshold $\bar{M}_K^*$ and corresponding cost rate $\eta(\bar{M}_K^*)$.

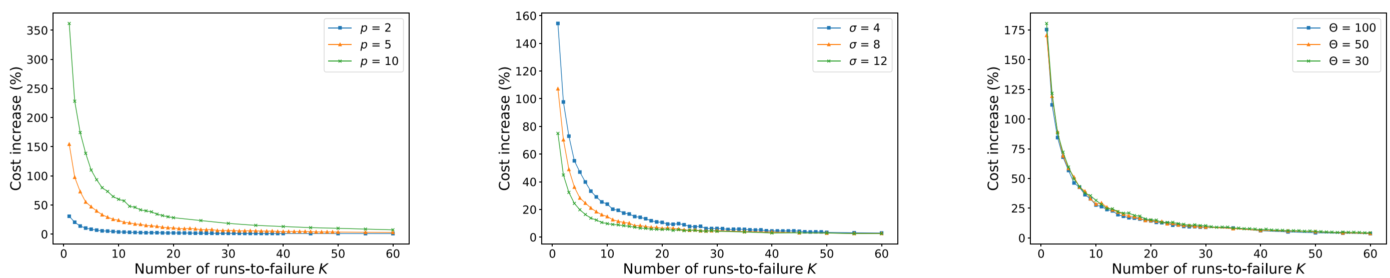

(a) (b) (c)

Fig. 6. Effect of various parameters on the percentage cost increase of the data-driven policy compared with the benchmark policy.

quite quickly as the number of runs-to-failure increases, especially after the first few runs-to-failure. We also observe that the cost rate $\eta(\bar{M}_K^*)$ converges to the optimal level much faster than the preventive maintenance threshold $\bar{M}_K^*$. This is because the cost rate function $\eta(M)$ is quite flat around its optimum (Fig. 3). Thus, small remaining deviations in $\bar{M}_K^*$ from the optimal value only have a minor impact on the costs.

6.3. Sensitivity analysis

To assess the effects of varying parameter values on our data-driven policy, we continue to perform a sensitivity study. We consider the effect of the ratio between the corrective and the preventive maintenance cost, of the volatility of the deterioration process, and of non-stationarity of the deterioration process. The parameter values are adjusted one by one while the other parameters keep the same values as in the base case. We again use 1000 simulations for each K we consider, and report percentage increases in the cost rate of our approach compared to the benchmark policy.

6.3.1. Maintenance cost

Fig. 6(a) shows the effect of the cost ratio $p = c_{\text{cm}}/c_{\text{pm}}$ on the effectiveness of the data-driven policy for $p = 2$, 5, and 10 ($p = 5$ in the base case). The cost increase of our approach is increasing in the cost ratio p, and this increase is most significant for a low number of runs-to-failure K. In other words, our approach performs less well if failures are more severe. In such situations with a high corrective maintenance cost, especially a maintenance threshold that is chosen too high leads to excessive costs. This is likely to occur if not much data is available, and it is strengthened by the fact that, on average, our approach leads to a too high maintenance threshold. For all cost ratios p the cost rate converges to the optimal cost rate if K increases, but for a higher corrective maintenance cost much more data is needed to reach the same performance for our approach.

6.3.2. Volatility in the deterioration process

Next, we consider the influence of the volatility of the deterioration process. As a measure for this volatility we use the standard deviation $\sigma = \sqrt{\alpha}\beta$ of the deterioration increments per period. We consider $\sigma = 4$, 8, and 12 ($\sigma = 4$ in the base case) and use combinations of values for α and β such that the mean time to failure is fixed at approximately 12.62; see De Jonge et al. (2017) for details.

The results are shown in Fig. 6(b). Interestingly, the percentage cost increase of our approach is highest for a more stable (i.e., less volatile) deterioration process. This can be explained by the fact that the cost rate around the optimum is more stable for a more volatile deterioration process; see Fig. 7. In such cases with more volatile deterioration, effective scheduling of preventive maintenance is difficult anyhow, implying that deviations from the optimum are less severe. On the other hand, if the deterioration process is less volatile, the optimal policy is able to carry out preventive maintenance just before most of the failures. Uncertainty in the failure behavior then has more severe consequences, especially if maintenance is carried out too late. Again, the performance is improving in the number of runs-to-failure K, and for a less volatile deterioration process more data is needed to achieve the same performance.

6.3.3. Non-stationary gamma process

We continue with deterioration according to a non-stationary gamma process. Instead of a fixed shape parameter α, we use shape function $\alpha(x) = a\exp(\frac{x}{\theta})$ with parameters a and θ to let

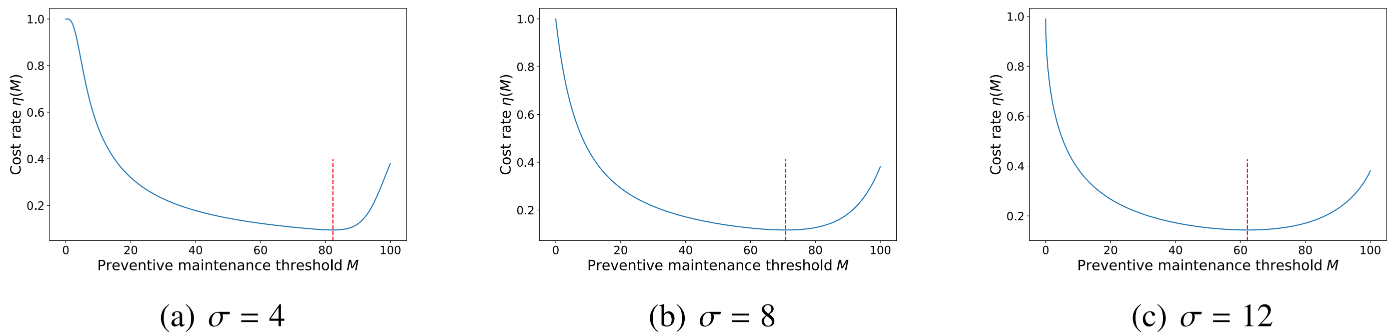

Fig. 7. Cost rate of the benchmark policy for different levels of volatility σ.

the deterioration increments depend on the current deterioration level x. This parametric form is also considered by Guida & Pulcini (2013). The scale parameter β does not vary with the deterioration level. We set $a = 1$, $\beta = 2$ and compare the cases $\theta = 30$, 50, and 100. When θ is lower, the deterioration process accelerates faster in the deterioration level.

Fig. 6(c) shows the results. It turns out that the percentage cost increase is almost not affected by varying θ, indicating that our approach is robust for different degrees of non-stationarity in the deterioration process. When the deterioration process accelerates faster in the deterioration level, the optimal maintenance threshold gets lower. When using our approach, the estimated failure probabilities at higher deterioration levels also get larger and, on average, the resulting maintenance threshold is reduced accordingly.

7. Conclusion and future extensions

We have developed a purely data-driven approach for making condition-based maintenance decisions for a single-unit system with an unknown deterioration process and unknown failure behavior. We use condition and failure data to estimate the failure probability of the system as a function of its deterioration level. Based on an optimality condition we derive for the case with a known deterioration process, we obtain a closed-form expression for the preventive maintenance threshold for the data-driven setting.

In order to assess the performance of our approach, we consider a gamma deterioration process. If only very limited data is available, the costs of using our policy is on average much higher than that of the optimal policy. However, the performance quickly improves when data for more runs-to-failure become available. Another observation is that the threshold that follows from our approach is on average higher than the optimal threshold, i.e., maintenance is carried out later than optimal on average. This is beneficial as it still ensures sufficient learning of the failure behavior once the policy resulting from our approach is being implemented.

A sensitivity analysis has shown that the effect of more data is robust under different maintenance costs, different levels of volatility in the deterioration process, and non-stationarity of the deterioration process. In all cases the resulting policy converges to the optimal policy when more data becomes available. Furthermore, our data-driven approach performs better when failures are less severe and when the deterioration process is more volatile. Non-stationarity of the deterioration process does not have a significant effect on the maintenance decisions and costs that follow from our approach.

There are numerous opportunities for future research in this direction. We consider a single condition to represent the state of a system, but in practice there are often multiple sensors resulting in multiple deterioration values. In such cases our approach needs to be extended. We need to learn, for instance, how the probability of failure depends on the combination of the different deterioration levels (Liu et al., 2021b; Mercier & Pham, 2012). Additionally, our approach suggests a preventive maintenance threshold that is higher than optimal, and this has the added benefit of additional learning. Future research could study the trade-off between more expensive learning (exploration) and its long-run advantages (exploitation) in more detail. The few existing studies in this area typically consider two or three component types with corresponding prior probabilities, and a known lifetime distribution or deterioration process given each component type (De Jonge et al., 2015; Dursun et al., 2022; Van Oosterom et al., 2017). However, a fully data-driven approach has not been developed.

The result of our approach is an estimated failure probability as a function of the deterioration level that is a step function, whereas in practice this function is likely to increase gradually. It is of interest to explore which continuous functions fit well with the data, and how these influence the performance of the data-driven approach. Another promising direction for future research seems the incorporation of a planning time for maintenance, since scheduling repairmen and ordering spare parts often require time (Kiesmüller & Sachs, 2020; Topan et al., 2020). It is valuable to analyze the effect of a planning time on the preventive maintenance threshold when the deterioration process is uncertain, and how this is affected by the amount of data. Furthermore, when the delivery time of spare components is not known, this should also be learned from data. A final direction is to study non-monotonic deterioration processes considering external factors such as varying environments and varying loads (Hu et al., 2021). In such settings it needs to be learned how the deterioration process and the failure behavior depend on the environment or the load.

Appendix A. Proof of lemmas

A1. Proof of Lemma 5.1

We consider an interval consisting of u observations resulting in failure, v observations not resulting in failure, and a constant failure probability p for all observations. The joint likelihood of the observations equals

$$L(p) = p^u (1-p)^v.$$

This is easily seen to be concave in p and the first order condition gives

$$\frac{dL}{dp} = u p^{u-1}(1-p)^v - v p^u (1-p)^{v-1} = 0,$$

which is solved by $p = \frac{u}{u+v}$.

A2. Proof of Lemma 5.2

Consider a block starting at observation x_a, consisting of u deterioration observations $x_a, \ldots, x_{a+u-1}$ that did result in failure, followed by v deterioration observations $x_{a+u}, \ldots, x_{a+u+v-1}$ that did not result in failure. The joint ML estimation of the non-decreasing sequence of failure probabilities p_a, …, $p_{a+u+v-1}$ is equal to the solution of

$$\begin{aligned} \min \quad & p_a \times \ldots \times p_{a+u-1} \times (1 - p_{a+u}) \times \ldots \times (1 - p_{a+u+v-1}), \\ \text{s.t.} \quad & p_a \leq \ldots \leq p_{a+u-1} \leq p_{a+u} \leq \ldots \leq p_{a+u+v-1}. \end{aligned}$$

We will prove by contradiction that all these probabilities p_a, …, $p_{a+u+v-1}$ are the same. Take any feasible solution with non-constant probabilities for the considered block. Then, there must exist some $j \in \{a, \ldots, a+u+v-2\}$ such that $p_j < p_{j+1}$. Obviously, y_{j+1} is either 0 (non-failure) or 1 (failure). We will show for both cases that an alternative feasible policy with higher likelihood exists, implying that the ML estimates must in fact be constant for a block. If observation $j+1$ is a non-failure, then lowering p_{j+1} to p_j (and leaving all other estimates unchanged) is still feasible and increases the likelihood of that observation and so also of all observations jointly (as the likelihoods of all other observations are unchanged). If observation $j+1$ is a failure, then raising p_j to p_{j+1} leads to a feasible solution with increased likelihood for observation j and so also for all observations jointly.

A3. Proof of Lemma 5.3

Let us consider any feasible (i.e., non-decreasing) sequence of failure probabilities with constant probabilities p_1 and p_2 for intervals I_1 and I_2, respectively. For the sequence to be feasible, we must have $p_1 \leq p_2$. We will show that any sequence with $p_1 < p_2$ cannot be optimal by constructing a better and still feasible alternative sequence with modified failure probabilities p'_1 and p'_2 for intervals I_1 and I_2, respectively, such that $p'_1 = p'_2$. The failure ratios of these two intervals are denoted as r_1 and r_2, respectively, with $r_1 > r_2$. We distinguish three cases: (i) $p_1 < r_1 < p_2$, (ii) $p_1 < p_2 \leq r_1$ and (iii) $p_1 \geq r_1$. Note that either case (i) or (ii) applies if $p_1 < r_1$ (since $p_1 < p_2$ for the considered sequence), and so the three cases together cover all possibilities.

Case (i): Set $p'_1 = p'_2 = r_1$. That is, increase p_1 to r_1 and decrease p_2 to r_1. This obviously increases the likelihood of I_1, but also of I_2 since $r_2 < r_1$ (applying Lemma 5.1). Moreover, the sequence must remain feasible, since the likelihood of I_1 is not decreased and that of I_2 is not increased, and so the likelihoods of all intervals are still non-decreasing as required.

Case (ii): Set $p'_1 = p'_2 = p_2$. That is, increase p_1 to p_2 and leave p_2 unchanged. This increases the likelihood of I_1, since $p_2 \leq r_1$ (applying Lemma 5.1). Using a similar argument as for case (i), it must also retain feasibility.

Case (iii): Combining $p_1 \geq r_1$ from the case definition with $p_1 < p_2$ and $r_1 > r_2$ as assumed, we get $p_2 > p_1 \geq r_1 > r_2$. Therefore, decreasing p_2 to $p'_2 = p_1$ and keeping p_1 unchanged ($p'_1 = p_1$) again increases the likelihood and retains feasibility.

This completes the proof.

References

Alaswad, S., & Xiang, Y. (2017). A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability Engineering and System Safety, 157*, 54–63.

Andersen, J. F., Andersen, A. R., Kulahci, M., & Nielsen, B. F. (2022). A numerical study of Markov decision process algorithms for multi-component replacement problems. *European Journal of Operational Research, 299*(3), 898–909.

Chen, N., Ye, Z.-S., Xiang, Y., & Zhang, L. (2015). Condition-based maintenance using the inverse gaussian degradation model. *European Journal of Operational Research, 243*(1), 190–199.

De Jonge, B. (2019). Discretizing continuous-time continuous-state deterioration processes, with an application to condition-based maintenance optimization. *Reliability Engineering and System Safety, 188*, 1–5.

De Jonge, B. (2021). Condition-based maintenance optimization based on matrix algebra. *Operations Research Letters, 49*(5), 741–747.

De Jonge, B., Dijkstra, A. S., & Romeijnders, W. (2015). Cost benefits of postponing time-based maintenance under lifetime distribution uncertainty. *Reliability Engineering and System Safety, 140*, 15–21.

De Jonge, B., & Scarf, P. A. (2020). A review on maintenance optimization. *European Journal of Operational Research, 285*(3), 805–824.

De Jonge, B., Teunter, R. H., & Tinga, T. (2017). The influence of practical factors on the benefits of condition-based maintenance over time-based maintenance. *Reliability Engineering and System Safety, 158*, 21–30.

Drent, C., Drent, M., Arts, J., & Kapodistria, S. (2023). Real-time integrated learning and decision making for cumulative shock degradation. *Manufacturing and Service Operations Management, 25*(1), 235–253.

Dursun, I., Akcay, A., & Van Houtum, G. J. (2022). Age-based maintenance under population heterogeneity: Optimal exploration and exploitation. *European Journal of Operational Research, 301*(3), 1007–1020.

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

Elwany, A. H., Gebraeel, N. Z., & Maillart, L. M. (2011). Structured replacement policies for components with complex degradation processes and dedicated sensors. *Operations Research, 59*(3), 684–695.

Fang, G., Pan, R., & Wang, Y. (2022). Inverse Gaussian processes with correlated random effects for multivariate degradation modeling. *European Journal of Operational Research, 300*(3), 1177–1193.

Flage, R., Coit, D. W., Luxhoj, J. T., & Aven, T. (2012). Safety constraints applied to an adaptive Bayesian condition-based maintenance optimization model. *Reliability Engineering and System Safety, 102*, 16–26.

Fouladirad, M., Grall, A., & Dieulle, L. (2008). On the use of on-line detection for maintenance of gradually deteriorating systems. *Reliability Engineering and System Safety, 93*(12), 1814–1820.

Guida, M., & Pulcini, G. (2013). The inverse gamma process: A family of continuous stochastic models for describing state-dependent deterioration phenomena. *Reliability Engineering and System Safety, 120*, 72–79.

Hu, J., Xu, A., Li, B., & Liao, H. (2021). Condition-based maintenance planning for multi-state systems under time-varying environmental conditions. *Computers and Industrial Engineering, 158*, 1–11.

Kiesmüller, G. P., & Sachs, F. E. (2020). Spare parts or buffer? How to design a transfer line with unreliable machines. *European Journal of Operational Research, 284*, 121–134.

Kurt, M., & Kharoufeh, J. P. (2010). Monotone optimal replacement policies for a Markovian deteriorating system in a controllable environment. *Operations Research Letters, 38*(4), 273–279.

Liu, B., Pandey, M. D., Wang, X., & Zhao, X. (2021a). 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.

Liu, B., Pandey, M. D., Wang, X., & Zhao, X. (2021b). A finite-horizon condition-based maintenance policy for a two-unit system with dependent degradation processes. *European Journal of Operational Research, 295*, 705–717.

Liu, B., Wu, S., Xie, M., & Kuo, W. (2017). A condition-based maintenance policy for degrading systems with age-and state-dependent operating cost. *European Journal of Operational Research, 263*(3), 879–887.

Mercier, S., & Pham, H. H. (2012). A preventive maintenance policy for a continuously monitored system with correlated wear indicators. *European Journal of Operational Research, 222*(2), 263–272.

Mosayebi Omshi, E., Grall, A., & Shemehsavar, S. (2020). A dynamic auto-adaptive predictive maintenance policy for degradation with unknown parameters. *European Journal of Operational Research, 282*(1), 81–92.

Nguyen, K. T. P., Do, P., Huynh, K. T., Bérenguer, C., & Grall, A. (2019). Joint optimization of monitoring quality and replacement decisions in condition-based maintenance. *Reliability Engineering and System Safety, 189*, 177–195.

Nguyen, K. T. P., & Medjaher, K. (2019). A new dynamic predictive maintenance framework using deep learning for failure prognostics. *Reliability Engineering and System Safety, 188*, 251–262.

Si, X., Li, T., Zhang, Q., & Hu, X. (2018). An optimal condition-based replacement method for systems with observed degradation signals. *IEEE Transactions on Reliability, 67*(3), 1281–1293.

Syan, C. S., & Ramsoobag, G. (2019). Maintenance applications of multi-criteria optimization: A review. *Reliability Engineering and System Safety, 190*, 106520.

Topan, E., Eruguz, A. S., Ma, W., Van Der Heijden, M. C., & Dekker, R. (2020). A review of operational spare parts service logistics in service control towers. *European Journal of Operational Research, 282*, 401–414.

Uit Het Broek, M. A. J., Teunter, R. H., De Jonge, B., & Veldman, J. (2021). Joint condition-based maintenance and load-sharing optimization for two-unit systems with economic dependency. *European Journal of Operational Research, 295*(3), 1119–1131.

Van Noortwijk, J. M. (2009). A survey of the application of gamma processes in maintenance. *Reliability Engineering and System Safety, 94*(1), 2–21.

Van Oosterom, C. D., Peng, C.-Y., & Van Houtum, G. J. (2017). Maintenance optimization for a Markovian deteriorating system with population heterogeneity. *IISE Transactions, 49*(1), 96–109.

Walter, G., & Flapper, S. D. (2017). Condition-based maintenance for complex systems based on current component status and Bayesian updating of component reliability. *Reliability Engineering and System Safety, 168*, 227–239.

Zhang, H., & Zhang, W. (2022). Analytical solution to a partially observable machine maintenance problem with obvious failures. *Management Science*.