Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria

doi: 10.3850/ESREL2026061419_esrel26-p26339-cd

Uncertainty-aware fault mode diagnostics for condition-based maintenance of bearing

Juseong Lee
Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, the Netherlands.
E-mail: j.lee@tue.nl

Claudia Fecarotti
Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, the Netherlands.
E-mail: c.fecarotti@tue.nl

Condition-based maintenance (CBM) relies on diagnostic information inferred from condition monitoring data to support maintenance decisions. While diagnostics are inherently uncertain in practice, most CBM frameworks either assume perfectly observable degradation states or implicitly assume fully reliable diagnostics. This limitation becomes more critical in systems with multiple fault modes, where diagnostic ambiguity is unavoidable and may strongly affect maintenance performance. This paper studies CBM of bearings under multiple fault modes. First, a physics-based kinematic model is developed to simulate vibration signals associated with outer race, inner race, and ball defects, extending an existing bearing model to include ball defects. Based on the simulated data, a data-driven fault diagnostics model is constructed using a 1D convolutional neural network with Monte Carlo dropout to estimate the posterior probability of each fault mode. The estimated diagnostic posterior is then integrated into maintenance decision-making through threshold-based condition-based replacement policies. Two policies are proposed: one triggers preventive replacement only when diagnostic confidence is sufficiently high, and the other triggers inspection when diagnostic uncertainty is high. A numerical case study demonstrates that relying on point-estimate diagnostics increases maintenance costs, whereas explicitly accounting for diagnostic uncertainty substantially reduces costs. The results highlight the importance of uncertainty-aware diagnostics for cost-effective condition-based maintenance, particularly when inspections are costly and diagnostics are imperfect.

Keywords: condition-based maintenance, fault diagnostics, diagnostic uncertainty, Monte Carlo dropout, bearing fault simulation.

1. Introduction

Condition-based maintenance (CBM) is an effective approach to maintain systems subject to increasing degradation, where maintenance decisions are taken based on the system conditions. Ideally, CBM enables just-in-time maintenance, where unnecessary interventions (early replacements) and failure risk are reduced Alaswad and Xiang (2017). However, its effectiveness hinges on how accurately and reliably the system health can be inferred from condition monitoring data de Jonge (2017). Accordingly, diagnostics directly influence the effectiveness of CBM policies. Nevertheless, most CBM frameworks proposed in the literature either assume perfectly observable degradation states Kıvanç et al. (2025), or implicitly assume fully reliable diagnostics. However, in practice, diagnostics are uncertain, and this uncertainty should be accounted for when making maintenance decisions.

When degradation paths are diverse with multiple possible fault modes, namely, multi-fault-modes diagnostics, diagnostics become more challenging, and their uncertainty becomes more crucial. Only a few studies consider diagnostic uncertainty in CBM, Lee and Mitici (2023); Lee et al. (2022); Eggertsson et al. (2023), and mostly rely on the simplified assumption of a single fault mode. Thus, integrating multi-fault mode diagnostics in CBM, considering the diagnostics uncertainty, is needed.

This work considers CBM of bearings in rotating machinery. The bearing consists of outer race, inner race, and balls (rolling elements). Depending on the location of the defect, it has three *fault modes* as shown in Fig. 1. Although bearings are often replaced as a whole rather than repaired

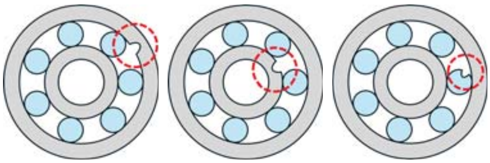

Fig. 1.: Three defect locations (fault modes) of a bearing: (a) outer race, (b) inner race, and (c) ball.

by individual elements, identifying the fault mode is essential because different fault modes have different operational impacts Randall and Antoni (2011).

Bearings are among the most-studied mechanical components for diagnostics, but the integration within CBM decision frameworks remains underexplored. For example, Zheng et al. (2024) propose a physics-based nonlinear dynamic model for early-stage bearing fault evolution. Taghiyarrenani and Berenji (2022) propose a data-driven method with data augmentation to identify bearing fault modes. However, estimation of diagnostic uncertainty and its use in CBM need further investigation in the context of maintenance optimization.

In this work, we integrate the uncertainty-aware diagnostics for multiple fault modes into a condition-based maintenance (CBM). As a preliminary step, we develop a dynamics model of a bearing to simulate faulty vibration signals from multiple bearing faults. This extends the previous work of Patil et al. (2010) to ball defects, which are critical contributors to failures in rotating machinery. We propose a data-driven diagnostics model to quantify the posterior probability of each fault mode using Monte Carlo dropout. The posterior distribution provides richer information for better CBM decisions than traditional point estimates of fault modes. Then we propose a threshold-based CBM strategy that integrates uncertainty-aware diagnostics into maintenance decisions. Using a numerical case study, we demonstrate that explicitly estimating and accounting for diagnostic uncertainty reduces maintenance costs.

2. Problem description: Condition-based maintenance of bearing using uncertainty-aware diagnostics

We consider a single-component system subject to increasing degradation. The system can be in a finite number of discrete states, and we denote the degradation state at time step t with $S_t \in \mathcal{S}$. The set $\mathcal{S}$ includes the healthy state H, the failed state F, and a number of intermediate mutually exclusive fault modes. The number and type of fault modes depend on the system at hand. The system state S_t is not directly observable, unless an *inspection* is performed at cost C_I. If the system fails ($S_t = F$), it is self-reporting, and we perform *corrective replacement* to return the system to the healthy state at cost C_C. Before the system fails ($S_t \neq F$), we can perform *preventive replacement* to return the system to the healthy state at cost C_P. We consider an infinite planning horizon, and divide it into discrete time steps t. At every time step t, we can perform maintenance action $a_t \in \{\emptyset, \mathrm{IN}, \mathrm{PM}, \mathrm{IN}+\mathrm{PM}\}$, where IN and PM denote inspection and preventive replacement, respectively. Here, corrective replacement is not a decision because it is automatically triggered by the self-reporting failure.

The objective is to minimize the long-run average maintenance cost $g(\pi)$, defined as,

$$g(\pi) = \lim_{T\to\infty} \frac{1}{T} \mathbb{E}_\pi \left[\sum_{t=1}^{T} C(a_t, S_t) \right], \quad (1)$$

where π is a maintenance policy and $C(\cdot)$ is the maintenance cost.

A condition monitoring tool (e.g., vibration sensors) is available to support maintenance decisions. However, the true state of the system is hidden, except for the failed state. Every time step t, the condition monitoring data (e.g., vibration sensors) of the system, X_t, is obtained with negligible cost. X_t is an unknown, noisy, physical function of the true state S_t. We aim to infer the posterior distribution of the hidden state given the observation,

$$h(s, X_t) = \mathbb{P}(S_t = s|X_t), s \in \mathcal{S}. \quad (2)$$

The diagnostics output serves as the information source for condition-based maintenance decisions.

This study proposes an uncertainty-aware diagnostics model h for a bearing with multiple fault modes, and a condition-based maintenance policy π using the uncertain diagnostics results.

3. Methodology

The methodology includes a dynamics model for bearing simulation, an uncertainty-aware diagnostic method, and a condition-based maintenance policy.

3.1. *Vibration simulation of faulty bearing using a dynamics model*

A dynamics model of bearing vibration under various defect types is proposed. This model is based on bearing kinematics and solves the equation of motion of the inner shaft due to Hertzian contact deformation caused by displacement and defects Patil et al. (2010).

The geometry of the bearing is defined by several parameters. The pitch diameter is D, and the ball diameter is d. The radial clearance is C_r. The number of balls is Z, and their indices are j.

Figure 2 shows the kinematics of the bearing assuming non-slip conditions among the inner race, outer race, and balls. The outer race is assumed to be fixed $\omega_o = 0$, and the rotational speed of the inner race is $\omega_i > 0$. The rotational speed of the bearing cage is $\omega_c = \frac{\omega_i}{2}(1 - \frac{d}{D}\cos\alpha)$, where α is the contact angle. The spin speed of the ball is $\omega_b = \frac{\omega}{2}\frac{D}{d}(1 - \frac{d^2}{D^2}\cos^2\alpha)$. The angular position of the ball j at time t is $\theta_j(t) = \omega_c t + \theta_j(0)$ as the ball rotates with the cage. The spin angle of the ball at time t is $\omega_b t$, and its spin angle with respect to the radial axis $\hat{r}$ is $\zeta(t) = \omega_b t + \theta_j(t)$.

The deformation of the bearing is caused by the displacement of the inner race and defects. Let the coordinates of the center of the inner race be (x, y). Its displacement larger than the clearance C_r causes the following deformation of ball j,

$$\delta_{j,\mathrm{disp}} = x\cos\theta_j(t) + y\sin\theta_j(t) - C_r \tag{3}$$

A defect is modeled as a dent on the surface of the outer race, inner race, and balls. The depth of a dent is modeled as a trigonometric form Patil

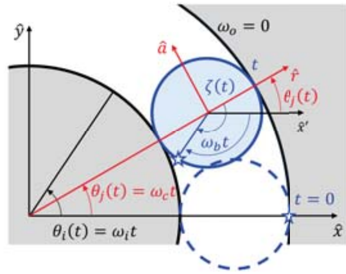

Fig. 2.: Kinematics of bearing

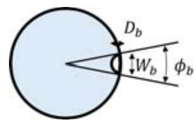

Fig. 3.: Geometry of a ball surface defect.

et al. (2010). Fig. 3 shows the geometry of a ball surface defect with depth D_b and width W_b. Its angular width is $\phi_b = \frac{W_b}{d/2}$. Its angular position at time t with respect to the radial axis $\hat{r}$ in Fig. 2 is obtained as,

$$\zeta_{j,b}(t) = \omega_b t + \omega_c t + \theta_j(0) + \psi_b, \tag{4}$$

where ψ_b is the angular position of the defect with respect to a reference point on the ball surface.

The defect on ball j contacts with outer race if and only if $-\frac{\phi_b}{2} \leq \zeta_{j,b}(t) \leq \frac{\phi_b}{2}$. Assuming its trigonometric shape, the contact reduces the deformation of the ball by $\delta_{j,\mathrm{ball}}$,

$$\delta_{j,\mathrm{ball}} = D_b \cos\left(\frac{\zeta_b(t)}{\phi_b}\pi\right) \tag{5}$$

Similarly, the defect contacts with inner race if and only if $-\frac{\phi_b}{2} \leq \zeta_b(t) + \pi \leq \frac{\phi_b}{2}$, and reduces deformation $\delta_{j,\mathrm{ball}} = D_b \cos\left(\frac{\zeta_{j,b}(t)+\pi}{\phi_b}\pi\right)$.

The defect on outer race surface and inner race surface are modeled in Patil et al. (2010). Let the reduced deformation due to outer race and inner race defect be $\delta_{j,\mathrm{out}}$ and $\delta_{j,\mathrm{in}}$, respectively. Then, the total deformation of ball j is,

$$\delta_j = \max(0, \delta_{j,\mathrm{disp}} - \delta_{j,\mathrm{out}} - \delta_{j,\mathrm{in}} - \delta_{j,\mathrm{ball}}) \tag{6}$$

The deformation exerts Hertzian force on the shaft,

$$F_j = K\delta_j^{3/2}, \tag{7}$$

Table 1.: Geometric parameters of bearing

Parameter	Value
K	$8.375 \times 10^9\,\mathrm{N/m^{3/2}}$
D	$43.39\,\mathrm{mm}$
d	$11.27\,\mathrm{mm}$
Z	7
C_r	$11.29\,\mu\mathrm{m}$
α	$0\,^\circ$

Source: Patil et al. (2010)

where K is Hertzian contact elastic deformation constant. Randall (2021) The net force on the shaft is the summation of F_j for all balls.

$$\begin{aligned} F_X &= \sum_{j=0}^{Z-1} F_j \cos\theta_j(t), \\ F_Y &= \sum_{j=0}^{Z-1} F_j \sin\theta_j(t) \end{aligned} \tag{8}$$

The vibration of the shaft (inner race) follows the equation of motion,

$$\begin{aligned} M\ddot{x} + c\dot{x} + F_X &= L_X \\ M\ddot{x} + c\dot{x} + F_Y &= L_Y, \end{aligned} \tag{9}$$

where M is mass on the shaft, c is damping coefficient of bearing in radial motions, and L_X and L_Y are external loads. Eq. (9) can be solved numerically using Runge-Kutta 4th order method. The solution $(\ddot{x}(t), \ddot{y}(t))$ simulates the 2 axes vibration signals,

$$\begin{aligned} a_x(t) &= \ddot{x}(t) + \epsilon_x \\ a_y(t) &= \ddot{y}(t) + \epsilon_y, \end{aligned} \tag{10}$$

where ϵ_x and ϵ_y are Gaussian sensor noise.

To account for operational uncertainty, the rotational speed of the inner shaft and the external loads are perturbed. The inner shaft speed ω_i is sampled from a Gaussian distribution with mean $1200\,\mathrm{RPM}$ and standard deviation $12\,\mathrm{RPM}$ (i.e., 1% of the nominal speed). The external loads are modeled as Gaussian random variables, with $L_X \sim \mathcal{N}(-100, 10^2)$, $L_Y \sim \mathcal{N}(0, 10^2)$.

3.2. *Uncertainty-aware diagnostics for multiple fault modes*

3.2.1. *Data pre-processing*

The vibration monitoring data $a_x(t)$ and $a_y(t)$ are transformed into the frequency domain using Fast Fourier Transform (FFT) to separate fault-related spectral components from noise Randall (2021). The FFT yields magnitude values at discrete frequency bins $f_k = \frac{k}{N} f_\mathrm{s}$, where f_s is sampling frequency and N is the number of samples. Only frequency components up to a predefined cut-off frequency f_cutoff were retained, as higher frequency components primarily capture measurement noise and hardly carry physically meaningful information for the bearing fault. f_cutoff is set sufficiently higher than the characteristic bearing fault frequencies (e.g., BPFO, BPFI, FTF, and BSF) Randall (2021). Finally, the magnitudes are normalized per frequency bin f_k based on the mean and standard deviation of the training data. Let $m_{y,k}$ and $m_{y,k}$ be the magnitude values of a_x and a_y at frequency bin f_k, and $X = \{(m_{x,k}, m_{y,k}) | 0 \le f_k \le f_\mathrm{cutoff}\}$ be the input feature for the fault diagnostics.

3.2.2. *1D CNN with MC dropout*

An uncertainty-aware diagnostics model is developed based on a 1-dimensional convolutional neural network (1D CNN) with Monte Carlo (MC) Dropout.

A 1D CNN extracts fault-related patterns localized along the frequency axis. CNN layers learn local spectral patterns by kernels, while an increasing number of channels capture features at increasing levels of abstraction. The network consists of three convolutional blocks, each composed of a 1D CNN layer, a ReLU activation, and max pooling. Max pooling progressively reduces the frequency resolution, focusing peak outputs. The numbers of channels and kernels are shown in Table 2

After the final CNN block, each channel is averaged into a single scalar by Global Average Pooling (GAP). GAP yields a compact channel-wise feature vector and significantly reduces the number of learnable parameters, preventing over-

Table 2.: Architecture of the 1D CNN for fault mode diagnostics

Layer	Type	Kernel / Pool	Channels
Input	–	–	2
Conv1	Conv1D + ReLU	7	16
Pool1	MaxPool1D	4	–
Conv2	Conv1D + ReLU	5	32
Pool2	MaxPool1D	4	
Conv3	Conv1D + ReLU	3	64
Pool3	MaxPool1D	4	–
GAP	Global Avg. Pool	–	64
FC1	Linear + ReLU	–	32
Dropout	Dropout	–	–
FC2	Linear (logits)	–	4

fitting and enhancing model generalization.

Lastly, the pooled features are passed to small fully connected layers with ReLU activation, followed by dropout layers. The final layer outputs class-wise logits $l(s)$ denoting the possible states $s \in \mathcal{S}$.

The model is trained using a cross-entropy loss for multi-class classification, augmented with an L_2 regularization term.

3.2.3. *Uncertainty estimation via MC dropout*

The model uncertainty of the diagnostics is quantified using MC dropout during inference. Multiple stochastic forward passes with dropout activated generate a set of predictive outputs for a given input feature X. This procedure yields a Monte Carlo approximation of the predictive posterior, which can be interpreted as approximate Bayesian inference under a variational approximation equivalent to a deep Gaussian process Gal and Ghahramani (2015).

$$\mathbb{P}\left(S = s|X\right) \approx \frac{1}{K}\sum_{k=1}^{K}\mathbb{P}\left(S = s|X, \mathbf{w}^{(k)}\right) \tag{11}$$

Here, $\mathbf{w}^{(k)}$ is the effective network weight of the k-th dropout realization.

The point-estimate of the state $\hat{s}_t$ is obtained by the maximum posterior.

$$\hat{s}_t = \arg\max_{s\in\mathcal{S}} \mathbb{P}(s|X_t) \tag{12}$$

3.3. *Condition-based maintenance using a posterior of fault diagnostics*

The true state of a bearing S_t at time t, is modeled as a discrete time Markov chain in Fig. 4.

$$S_t \in \mathcal{S} := \{H, O, I, B, F\}, \tag{13}$$

where H is the healthy state, O, I, and B are three fault modes, and F is failed state. Let $p_{i,j} = \mathbb{P}(S_t = j|S_{t-1} = i)$ denote the state transition probabilities. The degradation dynamics are characterized by the following assumptions. From the healthy state, the component either remains healthy or transits to one of the fault modes. The fault modes are mutually exclusive, and transitions between different fault modes do not happen. Each fault mode persists for a random duration and eventually transit to failure. The transition matrix P is,

$$P = \begin{bmatrix} p_{H,H} & p_{H,O} & p_{H,I} & p_{H,B} & 0 \\ 0 & p_{O,O} & 0 & 0 & p_{O,F} \\ 0 & 0 & p_{I,I} & 0 & p_{I,F} \\ 0 & 0 & 0 & p_{B,B} & p_{B,F} \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \tag{14}$$

The transition matrix P is unknown and unavailable for maintenance decisions, as is the case in a real maintenance process.

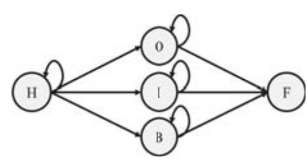

Fig. 4.: Markov chain model of bearing degradation

We propose two condition-based replacement (CBR) policies using the posterior inference $\mathbb{P}(S_t = s|X_t)$ based on the condition monitoring data X_t. A threshold-based approach is used.

The first policy (CBR1) is to trust diagnostics only if the posterior exceeds a threshold η_1. If a fault mode is inferred with sufficiently large posterior, i.e., $\mathbb{P}(\hat{s}|X_t) > \eta_1$, a preventive replacement (PM) is performed. Otherwise, the diagnostics result is considered insufficient to suspect a defect, and the system is let to run (do-nothing, $\emptyset$).

The second policy (CBR2) is to perform inspection when the posterior is less than a threshold η_2. If $\max_{s\in\mathcal{S}} \mathbb{P}(s|X_t) < \eta_2$, then an inspection is performed. The inspection is assumed to be perfect, so the true state s_t is obtained after the inspection. If $s_t \in \{O, I, B\}$, we perform preventive maintenance.

4. Numerical Case Study

4.1. *Uncertainty-aware fault mode diagnostics*

The proposed 1D CNN with MC dropout is trained on simulated condition monitoring data. The training data consists of 135 healthy data, and 135 faulty data. The faulty samples comprise 45 distinct defects, with 3 samples per defect. Each defect is characterized by its type (O, I, B), width (0.3 mm, 0.5 mm, 0.8 mm), and angular positions (-90°, -45°, 0°, 45°, and 90°). An example of a preprocessed feature X_t is shown in Fig. 5. The testing data set contains 15 healthy samples and 45 faulty samples with random sizes and positions.

First, the trained model is tested with its point-estimate diagnostics $\hat{s}_t$ in Eq. (12). The confusion matrix is shown in Fig. 6. The model detects the healthy, inner race defect, and ball defect nearly perfect, while the outer race defect has high false negative rate.

The high false negative rate for outer race defects is partly due to their physical nature. While the inner race defect and ball defect rotate rapidly with the bearing, the outer race defect is fixed. If the outer race defect is located outside of the loading zone, it does not cause detectable vibration because the bearing does not interact with it. For example, when vertical downward loading is applied, the outer race defect located in the upper side (Fig. 1c) is less likely to cause vibration because all the balls and the inner race are pushed downward. In fact, the false negative samples of outer race defect correspond to angular positions above -45°.

Second, the posterior distributions of the diagnostics are investigated. The size of each bar in Fig. 7 corresponds to the posterior of each state. The diagnostics exhibit ambiguity primarily between the pairs (H, O) and (I, B), as reflected by similarly distributed posterior across these classes. In particular, for samples with outer race defects (indices 16-30), the misclassified cases still assign a non-negligible posterior probability to state O. This indicates that these errors are not overconfident misclassifications. Such a quantified posterior enables maintenance decisions to account for diagnostic ambiguity, rather than relying solely on

Fig. 6.: Confusion matrix of test data

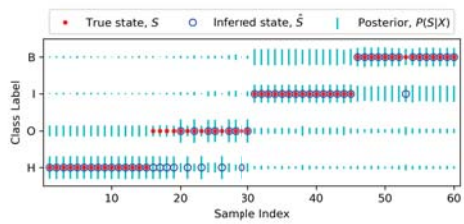

Fig. 7.: Posterior of the diagnostics of each test samples.

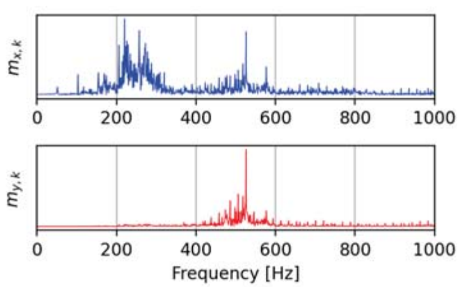

Fig. 5.: An example of feature X of a bearing with an inner race defect of size 0.5 mm.

Table 3.: Maintenance cost rate

Policy	Inspections	Prev. Rep.	Corr. Rep.	ECC	ECL	Cost rate	
TBR ($\tau_R = 4$)	-	0.71	0.28	21.47	3.76	5.71	0.0%
TBI ($\tau_I = 1$)	3.99	1.00	0.00	17.98	3.99	4.51	-3.2%
CBR0	-	0.97	0.03	11.20	1.99	5.63	+20.8%
CBR1 ($\eta_1 = 0.4$)	-	0.59	0.41	26.40	6.01	4.39	-5.7%
CBR2 ($\eta_2 = 0.4$)	3.44	1.00	0.00	16.88	4.22	4.00	-14.1%

point-wise classification.

4.2. *Condition-based maintenance*

We assume CBM setup where corrective maintenance is more expensive than preventive replacement ($C_C > C_P$) and inspection cost is not negligible ($C_I > 0$). In this numerical example, $C_I = 2$, $C_P = 10$, and $C_C = 50$.

We assume an arbitrary degradation model reflecting the operational characteristics of three fault modes. Outer and inner race defects are more common than ball defect, i.e., $p_{H,O} = p_{H,I} = 0.1$, and $p_{H,B} = 0.05$. The risk of critical bearing failure is largest with ball defect, while lest with outer race defect, i.e., $p_{O,F} = 0.2$, $p_{I,F} = 0.3$, and $p_{B,F} = 0.4$.

The two proposed condition-based replacement policies (CBR1 and CBR2), which use the diagnostic posterior, are compared to three typical maintenance policies. (1) Time-based replacement (TBR): A bearing is replaced after a fixed time interval τ_R. (2) Time-based inspection (TBI): A bearing is inspected after a fixed time interval τ_I. If the inspection finds a defect, the bearing is preventively replaced. (3) Condition-based replacement with point-estimate diagnostics (CBR0): At every time step t, the diagnostics result is fully trusted and if $\hat{S}_t$ is faulty (O, I, or B) then it is preventively replaced.

Monte Carlo simulation is conducted to estimate the cost rate (g) defined in Eq. (1). As the maintenance process is a renewal reward process, where both preventive and corrective replacements are renewal points, g can be calculated as the ratio of the expected cycle cost (ECC) to the expected cycle length (ECL). Table 3 shows the result of 100 simulations. The parameters τ_R and τ_I are numerically optimized.

Blindly relying on the diagnostics (CBR0) actually increases the cost rate. This is because any false-positive fault detection will immediately trigger preventive replacement. Too frequent preventive replacements reduce ECL unnecessarily and increase the cost rate.

By taking only the confident diagnostic results, CBR1 significantly reduces the cost rate compared to TBR, TBI, and CBR0. In particular, ECL is significantly higher under CBR1 than under CBR0, which is the main driver of cost reduction. As a trade-off, more corrective replacements are performed. Overall, this reduces the cost rate by 5.7%.

By paying for an inspection to get the true state when the diagnostics are less confident, CBR2 reduces the cost rate by 14.1%. This is achieved by two mechanisms. First, most corrective replacements are prevented by detecting defective states through inspections, especially when diagnostic uncertainty is high. Second, fewer inspections are performed than in TBI because inspections are requested only when diagnostic uncertainty is high.

Overall, Table 3 demonstrates that explicitly considering diagnostic uncertainty is essential for cost-effective condition-based maintenance.

5. Conclusion

We investigated condition-based maintenance (CBM) under uncertain multi-class fault diagnostics and demonstrated its benefits for reducing maintenance costs. First, we proposed a data-driven fault diagnostics model that quantifies uncertainty by using an approximate Bayesian posterior through Monte Carlo dropout. Second, we proposed two CBM methods that integrate diagnostic posterior information into preventive replacement and inspection decisions. The numerical case study showed that the proposed CBM policies reduce costs by preventing expensive corrective replacement and unnecessarily early preventive replacement, and by conducting inspections only when necessary. Finally, this work takes an initial step to link data-driven diagnostics and maintenance policy optimization, two fields of research that have been disjointed.

References

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

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

Eggertsson, R., R. Basten, and G.-J. v. Houtum (2023). Maintenance optimization for capital goods when information is incomplete and environment-dependent. *IISE Transactions ahead-of-print*(ahead-of-print), 1–16.

Gal, Y. and Z. Ghahramani (2015). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of The 33rd International Conference on Machine Learning*, arXiv.

Kıvanç, , C. Fecarotti, N. Raassens, and G.-J. v. Houtum (2025). Condition-based maintenance for multi-component systems: A scalable optimization model with two thresholds. *Reliability Engineering & System Safety 254*, 110634.

Lee, J. and M. Mitici (2023). Deep reinforcement learning for predictive aircraft maintenance using probabilistic remaining-useful-life prognostics. *Reliability Engineering & System Safety 230*, 108908.

Lee, J., I. d. Pater, S. Boekweit, and M. Mitici (2022). Remaining-useful-life prognostics for opportunistic grouping of maintenance of landing gear brakes for a fleet of aircraft. In *PHM Society European Conference*, Volume 7, pp. 278–285.

Patil, M., J. Mathew, P. Rajendrakumar, and S. Desai (2010). A theoretical model to predict the effect of localized defect on vibrations associated with ball bearing. *International Journal of Mechanical Sciences 52*(9), 1193–1201.

Randall, R. B. (2021). *Vibration-based condition monitoring: industrial, automotive and aerospace applications*. John Wiley & Sons.

Randall, R. B. and J. Antoni (2011). Rolling element bearing diagnostics—a tutorial. *Mechanical Systems and Signal Processing 25*(2), 485–520.

Taghiyarrenani, Z. and A. Berenji (2022). Noise-robust representation for fault identification with limited data via data augmentation. *PHM Society European Conference 7*(1), 473–479.

Zheng, L., Y. Xiang, and N. Luo (2024). Nonlinear dynamic modeling and vibration analysis for early fault evolution of rolling bearings. *Scientific Reports 14*(1), 23687.