

Article

A Digital Twin for Brake Wear Predictive Maintenance

Luke van Eyk , Johannes Moes , Brian Ellis , Stephan Schmidt  and Stephan Heyns 

Centre for Asset Integrity Management, University of Pretoria, Pretoria 0002, Gauteng, South Africa; johannes.moes@up.ac.za (J.M.); brian.ellis@live.co.za (B.E.); stephan.schmidt@up.ac.za (S.S.)

* Correspondence: luke.vaneyk@up.ac.za (L.v.E.); stephan.heyns@up.ac.za (S.H.)

Abstract

Brake pad wear is governed by coupled thermo-mechanical interactions in which degradation alters braking behaviour. This altered braking behaviour, in turn, affects vehicle dynamics, which, in turn, influences future wear evolution. Existing diagnostic and prognostic approaches typically neglect this dynamics–wear coupling, potentially limiting their ability to accurately predict remaining useful life under evolving operating conditions. This work proposes a digital-twin framework for brake wear prognosis that integrates state estimation, parameter adaptation, and dynamics-aware degradation modelling. The approach combines brake pad volume estimation from vehicle operational data, online identification of the brake wear coefficient through inverse modelling, and forward propagation of degradation using a wear-dependent dynamics model. The proposed digital-twin predictive maintenance framework is evaluated using simulated run-to-failure datasets for a mining load-haul-dumper and compared against data-driven and physics-based predictive maintenance models. Results show that the digital-twin approach improves the accuracy and stability of remaining useful life prediction, particularly under anomalous degradation conditions, and achieves earlier convergence to practically useful predictions. These findings demonstrate that accurate brake wear prognosis requires integrating degradation modelling with vehicle dynamics and online parameter updating. The proposed digital twin provides a practical pathway towards more reliable predictive maintenance in systems where degradation and system behaviour are strongly coupled.

Keywords: digital twin; brake wear; prognosis; hybrid methods

1. Introduction

Brake pad wear is a critical phenomenon due to its implications for vehicle safety, performance, and non-exhaust particulate emissions [1,2]. Wear processes are governed by complex tribological interactions influenced by operating conditions such as load, sliding velocity, pressure, material properties, surface condition, and particularly temperature, which strongly affect friction behaviour, thermal response, and material degradation [3–7]. Thermal brake studies further show that temperature evolution depends on contact pressure, convective cooling, thermal conductivity, and material composition, emphasising the importance of temperature-aware modelling in brake wear analysis [8,9].

Based on this understanding, extensive research has focused on diagnosing brake condition through estimation of wear rate and friction characteristics. These approaches include physics-based models capturing thermo-mechanical interactions [10], data-driven methods that learn relationships between operating conditions and wear behaviour [11–14], and hybrid approaches combining both paradigms [15].



Academic Editor: Sebastian Skatulla

Received: 25 May 2026

Revised: 13 July 2026

Accepted: 27 July 2026

Published: 1 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

While these methods provide estimates of current brake condition, they do not directly address the prediction of future behaviour. Prognostic approaches have therefore been developed, including Bayesian degradation models [16,17], data-driven extrapolation of braking behaviour [18,19], and multiphysics frameworks for life prediction [20]. However, these methods typically rely on pre-calibrated degradation models and assume that wear evolves independently of changes in system behaviour, limiting their ability to capture nonlinear end-of-life effects. This is particularly limiting because commonly used wear formulations, including Archard-type models, depend on empirical wear coefficients that can vary between materials, contact conditions, and operating regimes [21].

A more fundamental limitation across both diagnostic and prognostic approaches is the lack of coupling between brake wear and system dynamics. In practice, degradation alters braking performance, which in turn influences future wear evolution. Existing methods generally predict wear based on local conditions or historical trends, without accounting for this coupled interaction. Although digital-twin concepts have been explored, current implementations primarily replicate wear under known operating conditions rather than enabling long-term prognosis under evolving system behaviour [22]. These models rely solely on pre-calibrated trends, assuming wear evolves independently from the system dynamics.

This work addresses these limitations by proposing a digital-twin framework that models the coupled evolution of brake wear and vehicle dynamics over the full asset life. The proposed framework differs from existing brake wear studies in three important ways. First, rather than focusing only on diagnosis or wear-rate estimation, it performs prognosis by predicting future brake degradation and remaining useful life. Second, rather than relying on a fixed degradation parameter or pre-calibrated wear trend, it continuously updates the latent wear coefficient from operational data. Third, rather than propagating wear independently of future system response, it uses a wear-dependent brake atlas so that changes in brake volume influence future temperature, braking response, and degradation evolution. This enables remaining useful life to be predicted in a manner that accounts for both evolving brake condition and the resulting change in vehicle behaviour.

The remainder of this paper is structured as follows: Section 2 presents the data-generation methodology, the construction of various predictive maintenance methods (including the digital twin), and describes the simulation-based investigations to test these methods. Section 3 reports the results of these investigations, Section 4 discusses the findings of these results, and Section 5 concludes the work.

2. Materials and Methods

This study follows a structured methodology comprising brake degradation dataset construction, model development, and model evaluation.

Brake degradation data are generated using a physics-based simulation of a mining vehicle coupled with an Archard wear formulation. The use of simulated data in this work is intentional. Full run-to-failure brake wear data for mining vehicles are difficult to obtain because such tests are costly, time-consuming, commercially sensitive, and require measurement of the true degradation state throughout the brake life. Simulation allows the underlying brake pad volume, wear coefficient, operating conditions, and failure point to be precisely known, enabling fair and repeatable comparisons between predictive maintenance methods under controlled conditions.

Since simulation-based validation cannot provide the confidence afforded by experimental investigation, we considered it necessary to conduct comprehensive numerical testing to compensate, as far as possible, for the lack of experimental validation. We therefore devised a comprehensive evaluation process to test the prognostic models across nomi-

nal, anomalous, ablated, and model-mismatch conditions. In particular, the ablation and model-mismatch studies assess whether the observed digital-twin behaviour depends only on assumptions embedded in the nominal data-generation model, or whether the method remains robust when key aspects of the simulated physics are deliberately perturbed.

The datasets are produced using a calibrated multi-physics co-simulation framework, in which brake wear and vehicle dynamics evolve simultaneously. This ensures that the generated data capture the physical interaction between degradation and system response. The simulated data are treated as operational data collected from a vehicle equipped with a standard sensor suite, including brake temperature, wheel speed, brake line pressure, vehicle mass, and vehicle inclination. Braking events are identified when the brake line pressure exceeds a predefined threshold, and summary statistics are recorded for each event. Where direct measurements are not available, these quantities may be inferred from related signals, as demonstrated in the literature [19].

Using these datasets, three predictive maintenance approaches are constructed and evaluated: a data-driven model, a physics-based model, and a digital-twin model. Each approach uses the same underlying degradation data but differs in how the brake state is inferred and how future degradation is propagated.

The remainder of this section describes the brake degradation dataset construction process, followed by the formulation of the predictive maintenance models, and finally the investigation procedure used for evaluation of these predictive maintenance models.

2.1. Brake Degradation Dataset Construction

Brake degradation data were generated using a multi-physics simulation framework that captured the coupled evolution of brake wear and vehicle dynamics. The coupled dynamics–wear evolution is illustrated in Figure 1 and comprises five key steps which are repeatedly performed until a brake pad failure is reached.

Step 1: High-fidelity vehicle and brake simulation

Given the industrial relevance of brake wear monitoring in the mining industry, an Epiroc Scooptram ST14 SG load-haul-dumper (LHD) (Epiroc, Stockholm, Sweden) is utilised to demonstrate the digital-twin methodology. While this specific vehicle serves as the case study, the modelling procedure remains applicable to any vehicle. The vehicle dynamics are discretised into a multibody system of interconnected rigid bodies linked by kinematic joints and interacting through spring–damper force elements. The resulting model comprises 11 active degrees of freedom (DOFs): six for the spatial motion of the chassis, four for independent wheel rotations, and one for the articulated steering joint. Auxiliary DOFs associated with the bucket assembly are included in the model hierarchy but remain inactive for the scope of this study. The Simscape vehicle dynamics model is illustrated in Figure 2.

Brake dynamics are captured using a thermo-mechanical subsystem within Simscape (see Figure 3), which introduces a braking torque acting on the wheel rotational degrees of freedom. Although the underlying vehicle dynamics are parameterised from the ST14 platform, the brake architecture is adapted to meet the research objectives of this paper. Specifically, the physical ST14 typically employs a wet-braking system; however, for this study, it is replaced with a dry-friction rotor-disc system. Consequently, while the simulated performance is not an exact replication of the ST14 product specifications, the model remains phenomenologically correct in its coupling of thermal behaviour, friction variation, and wear evolution. By prioritising phenomenological accuracy over platform-specific replication, the study provides a robust environment for evaluating the digital-twin methodology while ensuring generalisability to other vehicle classes.




Figure 1. Five-step co-simulation framework used to generate brake degradation data. Step 1: High-fidelity simulation of vehicle braking behaviour. Step 2: Reduction of simulated braking behaviour to brake temperature and stopping-distance lookup maps. Step 3: Sampling of braking events from the active lookup map. Step 4: Brake wear calculation using Archard's wear law. Step 5: Update of brake thermal mass and transition to the next brake-volume map. The process is repeated until the brake failure threshold is reached.




Figure 2. Schematic layout of the multi-degree-of-freedom Simscape vehicle dynamics model.




Figure 3. Thermo-mechanical Simscape brake model used to calculate braking torque, heat generation, and brake temperature evolution.

In this framework, braking performance is primarily governed by the brake friction coefficient. As established in the literature [3,8,23–27], this coefficient is highly sensitive to the interface temperature. To capture this relationship, the temperature-dependent brake torque within the Simscape model is calculated as follows [28]:

$$\tau_b(T) = \mu(T) N_p \frac{\pi D_b^2}{4} r_m, \quad (1)$$

where $\mu(T)$ is the temperature-dependent friction coefficient, N_p is the number of brake pads, D_b is the actuator bore diameter, and r_m is the effective friction radius. In this work, we employ the temperature-dependent friction relationship outlined in Yevtushenko et al. [23]:

$$\mu(T) = \mu_0 \left(\mu_1 + \frac{\mu_2}{\mu_3(T - \mu_4)^2 + 1} \right). \quad (2)$$

The friction parameters to this formulation are outlined in Table 1. Note that the aforementioned relationship assumes temperature in degrees Celsius.

Table 1. Friction parameters as outlined in Yevtushenko et al. [23].

μ_0 [-]	μ_1 [-]	μ_2 [-]	μ_3 [°C ⁻¹]	μ_4 [°C]
0.45	0.01	1.07	0.0015	−250

The rate of heat generation during braking is given by

$$P_{\text{gen}} = \tau_b(T) \omega_{\text{slip}}, \quad (3)$$

where P_{gen} is the heat generation rate, and ω_{slip} is the relative angular velocity between the brake pad and the brake disc.

The resulting brake pad temperature evolution is governed by

$$\frac{dT}{dt} = \frac{P_{\text{gen}}}{C_{\text{th}}}, \quad (4)$$

where T is the brake temperature, and C_{th} is the thermal mass of the brake system. Note that for this work we assume a uniform temperature field. More extensive brake thermal models exist [8,9,24,25,27]. However, the focus of this work is to compare various predictive

maintenance models, and therefore, the simplified thermal distribution assumption is deemed sufficient for this preliminary study. The thermal mass is defined as

$$C_{th} = m_{brake} C_p, \quad (5)$$

where m_{brake} is the mass of the brake pads, and C_p is the specific heat capacity of the brake material. In this study, C_p is assumed constant across all operating conditions.

Importantly, C_{th} varies throughout the degradation process as brake pad material is worn away, resulting in a reduction in m_{brake} . This reduction in thermal mass leads to faster temperature increases during braking, which in turn reduces the effective friction coefficient (brake fade). Consequently, brake pad wear directly affects braking performance, establishing a key coupling between degradation (volume loss) and system dynamics (lower friction, longer stopping distance).

Step 2: Construction of high-fidelity maps and atlas

To enable efficient simulation of brake degradation over long time horizons, a brake performance atlas is constructed. The atlas serves as a high-speed surrogate model for the coupled Simscape dynamics–wear simulations in the form of a lookup table. It comprises a collection of high-fidelity lookup maps, where each lookup map corresponds to a discrete brake pad volume state. In this work, a new lookup map is constructed for each 1% reduction in brake pad volume.

For each lookup map, high-fidelity Simscape simulations are conducted over a grid of operating conditions and summarised into interpolants:

$$\mathcal{A}(E, \theta | V) \rightarrow \{T, d\}, \quad (6)$$

where E is the vehicle kinetic energy, θ is the terrain slope, V is the brake pad volume of interest, T is the average brake pad temperature, and d is the stopping distance. Therefore, this lookup map returns the expected vehicle behaviour (stopping distance and mean brake temperature) for a given set of vehicle inputs (kinetic energy and vehicle slope) and brake wear state (brake volume), without needing to simulate said event. Rather, the event can be interpolated between known high-fidelity simulation results.

The use of kinetic energy ($E = \frac{1}{2}mv_{vcl}^2$) in this lookup table is motivated by both physical and computational considerations. Kinetic energy captures the combined effect of vehicle mass and velocity on braking. This reduces the interpolation problem from three dimensions (m, v_{vcl}, θ) to two (E, θ), improving computational efficiency and robustness of the interpolation process.

To simulate real-world measurement uncertainties, additive white Gaussian noise is injected into each sensor channel. For each simulated sensor measurement s_i , the noisy observation, $s_{i, \text{noisy}}$, is defined as:

$$s_{i, \text{noisy}} = s_{i, \text{true}} + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_i^2), \quad (7)$$

where the standard deviation σ_i for each channel was selected to represent approximately 5% of its nominal operational magnitude. Specifically, the noise levels were defined with σ values of 5.0 K for temperature, 1.0 m for stopping distance, 2000 J for energy, and 0.01 rad for vehicle slope.

Each lookup map is constructed using twelve representative simulated braking scenarios, summarised in Table 2.

Table 2. High-fidelity braking scenarios used to construct each atlas map. Note that variations in vehicle mass occur due to variations in the loaded mass in the LHD bucket.

Run ID	Speed [km/h]	Vehicle Mass [kg]	Slope [°]
1	4.5	42,000	−10
2	15.0	42,000	−10
3	4.5	54,000	−10
4	15.0	54,000	−10
5	4.5	42,000	0
6	15.0	42,000	0
7	4.5	54,000	0
8	15.0	54,000	0
9	4.0	42,000	10
10	8.5	42,000	10
11	4.0	54,000	10
12	7.2	54,000	10

Step 3: Sampling operating conditions from atlas maps

For each braking event, vehicle operating conditions (speed, mass and slope) are sampled with equal probability from a range of possible values in the active lookup map. The lookup map is then queried to obtain the corresponding braking response, namely, mean brake temperature and stopping distance for the given sample. These sampled values represent realistic braking behaviour without requiring repeated high-fidelity simulation, dramatically speeding up brake wear simulation.

Step 4: Wear calculation using Archard law

For each sampled braking event, brake wear is computed using the Archard wear formulation. Due to the strong influence of temperature on braking behaviour, tribochemical effects have been shown to cause the effective hardness of brake pad materials to vary with temperature. Therefore, a modified, temperature-dependent Archard formulation is used in this work:

$$\dot{V}(T) = K \frac{Pv}{H(T)}, \quad (8)$$

where $\dot{V}(T)$ is the volumetric wear rate at brake pad temperature T , K is the dimensionless wear coefficient, P is the normal contact force between the brake pad and rotor, v is the relative sliding velocity between the brake pad and rotor, and $H(T)$ is the temperature-dependent hardness of the brake pad material.

The hardness term is calculated from the parameterised formulation given in Yevtushenko et al. [23]:

$$H(T) = H_0 \left(H_1 + \frac{H_2}{H_3(T - H_4)^2 + 1} + \frac{H_5}{H_6(T - H_7)^2 + 1} \right). \quad (9)$$

The hardness parameters to this formulation are outlined in Table 3. Note that the aforementioned relationship assumes temperature in degrees Celsius.

Table 3. Hardness parameters as outlined in Yevtushenko et al. [23].

H_0 [MPa]	H_1 [-]	H_2 [-]	H_3 [°C ^{−1}]	H_4 [°C]	H_5 [-]	H_6 [°C ^{−1}]	H_7 [°C]
137	−0.93	0.83	0.00234	546	2.02	0.002	−233

The wear increment is calculated based on the sampled braking response, and the brake pad volume is updated accordingly. Sampled values are only interpolated between

the known conditions of the current map, although future versions of the proposed data generation framework may consider doing a second linear interpolation across maps, to account for fractional brake pad volumes not captured by the atlas maps. This process is repeated until the brake pad volume reaches the next discrete volume level corresponding to another lookup map within the atlas.

Step 5: Map update and co-simulation loop

Once the brake pad volume crosses a threshold, the corresponding high-fidelity map is updated by modifying the brake thermal mass. Steps 1–5 are then repeated, ensuring that degradation and vehicle dynamics remain coupled throughout the simulation.

Training and Testing Dataset Compilation

Multiple co-simulation loops need to be completed and compiled to construct training and testing run-to-failure (RTF) datasets. In this work, ten RTF samples were used for training models, and two sets of ten RTF samples were used for testing the models. The first test set corresponded to RTF data with failure times similar to those of the training set. This is referred to as in-distribution (ID) testing data and represent normal historical wear rates. The second testing set corresponded to RTF data with failure times that were dramatically faster or slower compared to those in the training set. This is referred to as out-of-distribution (OOD) testing data and represent anomalously resilient (slower failure) brake pads or anomalously sensitive (faster failure) brake pads. The OOD dataset was constructed to test the predictive maintenance models' ability to accurately estimate failure times of brake pads that wear in some anomalous fashion.

To introduce variability between degradation trajectories, the Archard wear coefficient was sampled from a normal distribution,

$$K \sim \mathcal{N}(\mu_K, \sigma_K^2), \quad (10)$$

where $\mu_K = 2 \times 10^{-6}$ and $\sigma_K = 2 \times 10^{-7}$ were used in this work. Training and ID test datasets were generated using this distribution, while the OOD test dataset was constructed by shifting the mean by $\pm 6\sigma_K$. Five OOD samples were selected from six standard deviations below the training and ID mean (i.e., slower wear rates), and five OOD samples were selected from six standard deviations above the training and ID mean (i.e., higher wear rates).

The simulated degradation trajectories are illustrated in Figure 4a. A summary of the information contained in each of these trajectories is illustrated in Figure 4b. Note that the trajectories stop at around 18% remaining brake pad volume. During simulation, this was seen to be the point at which, under extreme braking, the brakes failed and could not bring the vehicle to a stop.

Notice that the degradation curves are nonlinear. Recall that wear rates are fixed during dataset generation. This suggests that although a linear Archard wear model is used with a fixed degradation rate, the overall wear rate of the brakes accelerates over the life of the brakes due to dynamics–wear coupling effects. The lower brake pad volume leads to higher mean braking temperature, which in turn leads to a lower mean brake friction coefficient, ultimately increasing the number of wheel rotations required to stop the vehicle. This leads to higher amounts of wear near the end of life of the brake pad. These results showcase that prognosis of brake pad wear is not as simple as propagating the wear model forward in time, as nonlinear dynamics–wear coupling effects may drastically alter the behaviour of brake wear.

A summary of the datasets used in this study is provided in Table 4. Each dataset consists of complete run-to-failure trajectories, where each row corresponds to a braking

event. For each braking event, the recorded variables include the braking event index, brake pad volume, mean brake pad temperature during braking, stopping distance, kinetic energy, terrain angle, and remaining useful life. The training dataset was used for model calibration, while the ID and OOD datasets were reserved exclusively for testing.




Figure 4. Comprehensive view of the Archard dataset simulations and individual trace structure. (a) Simulated run-to-failure data. Ten training, ten in-distribution and ten out-of-distribution run-to-failure runs are illustrated. (b) Representative event-level braking data contained in a single run-to-failure trajectory.

Table 4. Summary of the simulated run-to-failure datasets.

Dataset	Number of RTF Trajectories	Wear Coefficient Distribution	Purpose
Training	10	$K \sim \mathcal{N}(\mu_K, \sigma_K^2)$	Model calibration and training
ID testing	10	$K \sim \mathcal{N}(\mu_K, \sigma_K^2)$	Evaluation under nominal historical wear behaviour
OOD testing	10	$K \sim \mathcal{N}(\mu_K \pm 6\sigma_K, \sigma_K^2)$	Evaluation under anomalously slow and fast wear behaviour

2.2. Data-Driven Model Construction

The data-driven approach estimates the remaining useful life (RUL) of the brake pads directly from measured operating data using a neural network. In contrast to the physics-based and digital-twin approaches (described shortly), no explicit degradation law is imposed. Instead, the model learns a direct regression from braking-event measurements to the corresponding RUL.

Each run-to-failure trajectory is decomposed into individual braking events, and each event forms one supervised training sample. For braking event k within trajectory i , the operational input vector is defined as

$$\mathbf{x}_{i,k} = [N_{i,k}, d_{i,k}, E_{i,k}, \theta_{i,k}, T_{i,k}]^T, \quad (11)$$

where $N_{i,k}$ is the braking event index, $d_{i,k}$ is the stopping distance, $E_{i,k}$ is the vehicle kinetic energy at the onset of braking, $\theta_{i,k}$ is the vehicle slope, and $T_{i,k}$ is the mean brake temperature.

The corresponding regression target is the true remaining useful life,

$$y_{i,k} = \text{RUL}_{i,k}. \quad (12)$$

The training dataset consists of multiple run-to-failure trajectories, each with a different number of braking events. Let $i = 1, \dots, N_{\text{traj}}$ index the trajectories, and let n_i denote the number of events in trajectory i . The full training dataset is then defined as

$$\mathcal{D}_{\text{train}} = \bigcup_{i=1}^{N_{\text{traj}}} \{(\mathbf{x}_{i,k}, y_{i,k})\}_{k=1}^{n_i}, \quad (13)$$

where $N_{\text{traj}} = 10$ in this study.

The neural network therefore learns a mapping

$$\hat{y}_{i,k} = g_{\phi}(\mathbf{x}_{i,k}), \quad (14)$$

where $g_{\phi}(\cdot)$ denotes the neural network with parameters ϕ , and $\hat{y}_{i,k}$ is the predicted RUL for event k in trajectory i .

Before training, each input feature is normalised using the statistics of the training set:

$$\tilde{\mathbf{x}}_{i,k} = (\mathbf{x}_{i,k} - \boldsymbol{\mu}_x) \oslash \boldsymbol{\sigma}_x, \quad (15)$$

where $\boldsymbol{\mu}_x$ and $\boldsymbol{\sigma}_x$ are the feature-wise mean and standard deviation computed across all trajectories and events in $\mathcal{D}_{\text{train}}$, and $\oslash$ denotes element-wise division.

The regression model is implemented as a feedforward neural network with two hidden layers of 20 and 10 neurons, respectively, followed by a single output neuron for RUL prediction. The network parameters are trained by minimising the discrepancy between predicted and true RUL values across all training samples. The optimisation problem is defined as

$$\phi^* = \arg \min_{\phi} \sum_{i=1}^{N_{\text{traj}}} \sum_{k=1}^{n_i} \mathcal{L}(g_{\phi}(\mathbf{x}_{i,k}), y_{i,k}), \quad (16)$$

where $\mathcal{L}(\cdot)$ is the regression loss function, and ϕ^* denotes the optimal network parameters.

Figure 5 illustrates the training process. The model receives only event-level operational measurements and returns a direct RUL estimate, without producing an intermediate estimate of brake pad volume or degradation parameter.

A key consequence of this modelling choice is that the degradation mechanism is inferred implicitly from the data rather than represented explicitly. Under the continuous data acquisition assumption considered in this work, no direct brake pad volume measurement is available to the network. The model must therefore rely entirely on correlations between observable operating variables and the degradation state. Its performance is

consequently governed by the representativeness of the training data and its ability to generalise to degradation behaviour not previously observed.




Figure 5. Data-driven RUL prediction model. Each braking event is represented by measurable operating variables, which are passed to a neural network to obtain a direct RUL estimate. The prediction is compared to the true RUL, and the network parameters are updated iteratively during training.

Potential overfitting is addressed in several ways. First, the training, ID testing, and OOD testing datasets consist of separate run-to-failure trajectories, so no test trajectory is used during model fitting. Second, feature normalisation is performed using only the training data statistics, which are then applied unchanged to the test sets. Finally, the data-driven model is intentionally kept small, with two hidden layers of 20 and 10 neurons, to reduce the risk of memorising the limited training trajectories.

2.3. Physics-Based Model Construction

The physics-based approach models brake degradation by explicitly representing the underlying wear mechanism and propagating it into the future. In this work, the standard Archard wear law (the same form as in Equation (8) but without temperature dependence) was used to describe the evolution of brake pad volume as a function of operating conditions. The model behaviour is outlined in Figure 6.

For braking event k within trajectory i , the brake pad volume is denoted by $V_{i,k}$. The wear evolution is governed by

$$V_{i,k+1} = V_{i,k} - \Delta V_{i,k}, \quad (17)$$

$$\Delta V_{i,k} = f_{\text{Archard}}(\mathbf{u}_{i,k}; K), \quad (18)$$

where $\Delta V_{i,k}$ is the wear increment, $\mathbf{u}_{i,k}$ represents the operating conditions, and K is the wear coefficient. f_{Archard} describes the Archard wear formulation from Equation (8) but with temperature-dependency removed. The hardness value used in this formulation is calculated from Equation (9) at 20 °C.

The wear coefficient K is calibrated using the available run-to-failure training trajectories. Let K_i denote the degradation parameter associated with trajectory i . The calibrated parameter is obtained as the mean value across all training trajectories,

$$K_{\text{avg}} = \frac{1}{N_{\text{traj}}} \sum_{i=1}^{N_{\text{traj}}} K_i. \quad (19)$$

degradation state from operational data, and second, to propagate that state into the future while accounting for the fact that brake wear alters subsequent vehicle behaviour. This is achieved by combining the Archard wear formulation introduced previously with the brake performance atlas, yielding a model that is both physically grounded and computationally efficient. The summarised methodology is illustrated in Figure 7 and described in the following sections.




Figure 7. Digital twin structure used in this study. At each braking event, measured operating variables are used to infer the current brake pad volume through the atlas. This inferred state is used both to update the wear model parameter and to initialise forward degradation simulation. The degradation model is iteratively coupled with the atlas, allowing the vehicle response to evolve with brake wear.

2.4.1. Brake Degradation State Estimation (Diagnosis)

The brake performance atlas acts as a surrogate for the high-fidelity vehicle and brake simulation model. Recall that for a discrete set of brake pad volume states $V \in \mathcal{V}$, and for given operating inputs (v_{vcl}, m, θ) , the atlas returns the corresponding braking response, $\mathcal{A}(v_{vcl}, m, \theta; V) \rightarrow \{T, d\}$. Each atlas layer therefore represents the expected vehicle response at a specific brake wear state.

For braking event k within trajectory i , the digital twin receives measured quantities

$$\mathbf{u}_{i,k} = \{v_{i,k}, m_{i,k}, \theta_{i,k}, T_{i,k}, d_{i,k}\}, \quad (22)$$

where $T_{i,k}$ and $d_{i,k}$ denote the measured brake temperature and stopping distance, respectively. The first diagnosis task of the digital twin is to infer the current brake pad volume, denoted $\hat{V}_{i,k}$. This is achieved through reverse inference using the atlas. Since each atlas layer corresponds to a known brake volume, the twin determines the volume that best explains the observed braking temperature from the vehicle inputs:

$$\hat{V}_{i,k} = \arg \min_{V \in \mathcal{V}} \left[(T^{\text{atlas}}(v_{i,k}, m_{i,k}, \theta_{i,k}; V) - T_{i,k})^2 \right], \quad (23)$$

where $T^{\text{atlas}}(\cdot)$ is the atlas-predicted temperature. An illustration of this process is shown in Figure 8. In theory, both stopping distance and temperature may be used to perform reverse inference. However, after initial experimentation, it was seen that only temperature values monotonically increased with increased brake wear, and therefore, stopping distance

was not a reliable reverse inference parameter. Therefore, only temperature was used for the reverse inference task.




Figure 8. Reverse inference using the atlas. The measured brake temperature is compared against atlas predictions across all wear states, and the brake pad volume is inferred from the minimum-error match.

The second diagnosis task is the online calibration of the wear model. The wear evolution follows the Archard formulation in Equation (8), but unlike that model, the degradation parameter is continuously updated. This update is also formulated as an inverse problem.

At a given update event, the digital twin has access to an initial reference volume V_{start} , an inferred current volume $\hat{V}_{i,k}$, and the number of elapsed braking events $\Delta N_{i,k}$. The objective is to determine the wear parameter $K_{i,k}$ such that the wear model reproduces the observed degradation over $\Delta N_{i,k}$ events. This is expressed as

$$\hat{K}_{i,k} = \arg \min_K |N_{\text{sim}}(V_{\text{start}}, \hat{V}_{i,k}; K) - \Delta N_{i,k}|, \quad (24)$$

where $N_{\text{sim}}(\cdot)$ denotes the number of simulated braking events required to degrade from V_{start} to $\hat{V}_{i,k}$ under the wear model. This quantity is evaluated using the atlas-based degradation simulation.

In practice, this optimisation is solved numerically by iteratively simulating degradation trajectories for candidate values of K until the simulated degradation time matches the observed event count. This effectively inverts the wear model, allowing the digital twin to identify the degradation parameter that is consistent with the observed system behaviour.

2.4.2. Brake Degradation State Propagation (Prognosis)

Starting from the current estimated state $(\hat{V}_{i,k}, \hat{K}_{i,k})$, the twin simulates future braking events by repeatedly querying the atlas corresponding to the current wear state, computing wear increments, and updating the brake volume:

$$\hat{V}_{i,j+1}^+ = \hat{V}_{i,j}^+ - f_{\text{Archard}}^* \left(\mathcal{A}(v_{i,j}, m_{i,j}, \theta_{i,j}; \hat{V}_{i,j}^+), \hat{K}_{i,k} \right), \quad j \geq k. \quad (25)$$

Note that in this instance, f_{Archard}^* refers to the temperature-modified version outlined in Equation (8), with the mean braking temperature from the lookup map used to calculate a mean brake pad hardness, $\bar{H} = H(\bar{T})$.

In implementation, the forward simulation is performed in batches of 100 braking events at a time. For each batch, operating conditions are sampled from the admissible range of the currently active atlas layer, and the corresponding brake temperatures and stopping distances are obtained from the atlas. These responses are then passed through the Archard wear model to accumulate the total wear over the batch. After each batch, the brake volume is updated and compared against the atlas's discrete volume levels. If the updated volume crosses into the next degradation interval, the active atlas layer is switched before the next batch is simulated. In this way, the prognosis remains computationally efficient while still allowing the predicted vehicle response to evolve with brake wear. An illustration of this process is depicted in Figure 9.




Figure 9. Atlas-driven degradation prognosis loop used in the digital twin. Future braking events are simulated from the current estimated brake state, with the active atlas layer updated as the predicted brake volume decreases.

This process continues until a predefined failure threshold is reached, at which point the remaining useful life is determined as the number of braking events remaining.

The digital twin therefore introduces three key capabilities. First, it performs cycle-by-cycle estimation of brake pad volume through reverse inference. Second, it continuously updates the degradation parameter using operational data through an inverse modelling procedure. Third, it propagates degradation through a wear-dependent dynamics model, ensuring that future behaviour reflects the evolving condition of the brake system. Together, these elements enable the digital twin to capture the coupled evolution of brake wear and vehicle dynamics, which is not possible with purely data-driven or physics-based approaches.

2.5. Simulation-Based Predictive Maintenance Model Investigations

The investigations were constructed to evaluate how the predictive maintenance models performed under both nominal and anomalous degradation conditions. All models were trained on the same ten run-to-failure (RTF) training samples outlined in Section 2.1 and evaluated on the same ID and OOD test datasets to ensure fair comparison. Six complementary investigations were performed and are discussed in greater detail below. The first four investigations evaluated the three complete predictive maintenance models, namely, the data-driven model, the physics-based model, and the digital-twin model. The fifth investigation performed an ablation study to isolate the mechanisms responsible for digital-twin performance. The sixth investigation evaluated all models' robustness under deliberate model mismatch.

Investigation 1: Brake pad volume estimation

The first investigation evaluated brake pad volume estimation accuracy. This experiment was applied only to the physics-based and digital-twin models, since these were the only two approaches that maintained an explicit degradation state. The objective was to determine whether each model could estimate the current brake pad volume throughout the asset life, particularly under OOD degradation conditions. A representative accelerated OOD test sample was used to visualise the volume-estimation behaviour, while all runs were retained for the quantitative RUL analyses described in later investigations.

Investigation 2: Wear coefficient estimation

The second investigation evaluated online estimation of the Archard wear coefficient. This investigation was applied only to the digital twin, since the data-driven model did not contain an explicit wear parameter and the physics-based model used a fixed coefficient calibrated from historical data. The objective was to determine whether the digital twin could infer a latent degradation parameter from operational data and adapt when the true wear behaviour differed from the historical training distribution.

Investigation 3: Remaining useful life estimation

The third investigation evaluated remaining useful life (RUL) prediction accuracy for all three predictive maintenance models. Representative ID and OOD test samples were used to visualise each method's prediction behaviour. In addition, RUL prediction errors were summarised across all ten ID and ten OOD runs using the root-mean-square error (RMSE). This provided a pointwise accuracy measure, although the convergence analysis in investigation 4 was used as the primary practical evaluation criterion because it captured how early each model became useful for maintenance decision-making.

Investigation 4: RUL convergence-sensitivity meta-analysis

The fourth investigation evaluated the convergence behaviour of all three predictive maintenance models across all ten ID and ten OOD test samples. This investigation determined how early each model reached a practically useful RUL prediction for different acceptable error tolerances.

For a given testing sample trajectory i and braking event k , the instantaneous RUL prediction error is defined as

$$e_{i,k}^{\text{RUL}} = \left| \widehat{\text{RUL}}_{i,k} - \text{RUL}_{i,k} \right|, \quad (26)$$

where $\widehat{\text{RUL}}_{i,k}$ is the model-predicted remaining useful life, and $\text{RUL}_{i,k}$ is the true remaining useful life. If an acceptable tolerance ε is prescribed, the convergence event for a model is defined as the earliest braking event k^* for which

$$e_{i,k}^{\text{RUL}} \leq \varepsilon \quad \forall k \geq k^*. \quad (27)$$

This provides a practical measure of the earliest point at which a model becomes sufficiently accurate for maintenance decision-making at accuracy level ε for various planning horizon points. The convergence analysis was repeated for several tolerance values to assess model usefulness over different planning horizons. Mean convergence behaviour is reported across the ten runs, with shaded confidence regions corresponding to half a standard deviation.

Investigation 5: Digital twin ablation study

The fifth investigation isolated the contribution of the main digital-twin mechanisms. Instead of comparing the full digital twin only against other models, the full digital twin was compared against simplified variants of itself. This allowed the influence of dynamics–

wear coupling, brake-state estimation, and online wear-parameter updating to be assessed separately. The ablation variants are summarised in Table 5.

The static physics baseline corresponds to the physics-based model described previously, in which degradation is propagated using a fixed wear parameter and a fixed unworn brake response. The coupled propagation variant retains the fixed historical wear parameter but propagates degradation through the wear-dependent atlas so that the simulated vehicle response changes as brake volume decreases. The state-updated variant additionally performs periodic brake-volume inference from operational temperature measurements but retains a fixed wear parameter. Finally, the full digital twin performs both brake-volume inference and online wear-parameter updating, while also propagating future degradation through the wear-dependent atlas.

Table 5. Digital-twin ablation variants.

Variant	Online Parameter Update	State Estimation	Dynamics–Wear Coupling
Static physics baseline	No	No	No
Coupled propagation only	No	No	Yes
State-updated DT, fixed K	No	Yes	Yes
Full digital twin	Yes	Yes	Yes

The difference between the static physics baseline and the coupled propagation variant isolates the effect of dynamics–wear coupling. The difference between coupled propagation and the state-updated variant isolates the effect of brake-state estimation. The difference between the state-updated variant and the full digital twin isolates the effect of online wear-parameter updating. The same RUL convergence-sensitivity procedure from investigation 4 was then used to compare the four ablation variants under ID and OOD conditions.

Investigation 6: Model-mismatch robustness study

The sixth investigation evaluated robustness under deliberate model mismatch. The primary validation dataset was generated from the same nominal modelling framework used to construct the digital-twin atlas. To determine whether the results depended too strongly on shared modelling assumptions, additional datasets were generated using perturbed physics. During evaluation, however, the digital-twin and physics-based baseline retained the original unperturbed atlas and original internal wear-law structure. This created a deliberate mismatch between the data-generating model and the predictive models used during evaluation.

Four perturbation cases were considered: a friction–temperature perturbation, a hardness–temperature perturbation, a speed-dependent wear perturbation, and a combined perturbation, which simultaneously introduced the latter three effects. These cases are summarised in Table 6.

Table 6. Model-mismatch perturbation cases.

Perturbation Case	Perturbed Atlas	Perturbed Wear Law	Purpose
Friction perturbation	Yes	No	Braking-response mismatch
Hardness perturbation	No	Yes	Material degradation mismatch
Speed-dependent wear	No	Yes	Missing speed-dependent wear effect
Combined perturbation	Yes	Yes	Combined response and wear mismatch

The friction and hardness perturbations were constructed by scaling and shifting the nominal temperature-dependent relationship from Equations (2) and (9). In both cases, the perturbation was selected so that the perturbed relationship differed from the nominal relationship by approximately 20% near the lower and upper ends of the plotted temperature range. This provided a significant but controlled perturbation of the material behaviour. The nominal and perturbed friction and hardness relationships are shown in Appendix A.

The friction perturbation changed the high-fidelity braking response directly, and therefore, a new friction-perturbed atlas was generated for the data-generation model. The hardness perturbation affected the wear increment but did not directly alter the high-fidelity braking response at a fixed brake volume. Therefore, no new atlas was required for the hardness-only case. Instead, the perturbed hardness law was applied during degradation synthesis.

The speed-dependent wear perturbation was included to represent a structural wear-law mismatch. In the nominal degradation model, event severity is already represented through the sampled braking response and the resulting wear increment. However, real brake wear may contain additional speed-dependent effects that are not fully captured by the nominal event-level wear model. To test the effect of such an omitted mechanism, the data-generation process was modified so that the nominal wear increment was scaled by an additional speed-dependent multiplier.

For each sampled braking event, the perturbed wear increment is defined as

$$\Delta V_{\text{spd}} = \Delta V_{\text{nom}} \gamma_v, \quad (28)$$

where ΔV_{nom} is the nominal wear increment calculated from the temperature-dependent Archard formulation (Equation (8)), and γ_v is an additional speed-dependent multiplier. The calculation of γ_v is provided in Appendix B. In essence, lower-speed braking events reduce the nominal wear increment, while higher-speed braking events increase it, thereby creating a degradation pattern that is not represented in the internal wear model used by the predictive maintenance methods.

This perturbation changes only the degradation evolution during data generation. It does not require a new atlas, because the high-fidelity braking response at a fixed brake volume remains unchanged. During evaluation, the digital-twin and physics-based model still used the nominal wear formulation without the additional speed-dependent multiplier. The speed-dependent case therefore tested whether the models remained robust when the validation data contained an unmodelled speed-related wear mechanism.

For each perturbation case, new training, ID, and OOD datasets were generated. The models were then evaluated using the same RUL convergence-sensitivity procedure defined in investigation 4. The digital twin was allowed to estimate its initial wear parameter from the perturbed training data, representing calibration from historical fleet behaviour, but it was not given access to the perturbed atlas or perturbed wear-law structure during evaluation.

3. Results

This section presents the performance of the predictive maintenance models across the six investigations outlined previously. The first set of results considers representative trajectory-level behaviour, including brake volume estimation, wear-coefficient estimation, and remaining useful life (RUL) prediction. The second set of results summarises model behaviour across all test trajectories using RMSE and convergence-sensitivity analyses. Finally, the ablation and model-mismatch investigations are presented to support the later interpretation of digital-twin mechanisms and robustness.

The brake pad volume predictions for representative in-distribution (ID) and out-of-distribution (OOD) test samples (investigation 1) are illustrated in Figure 10. The selected OOD sample corresponds to a case with an elevated wear rate relative to the training dataset. Since the data-driven model does not maintain an explicit brake degradation state, only the physics-based and digital-twin models are included in this comparison.




Figure 10. Brake pad volume estimates for the physics-based and digital-twin predictive maintenance models. The top plot shows in-distribution test data with wear behaviour similar to the training dataset, while the bottom plot shows out-of-distribution test data with a significantly higher wear rate.

The online wear-coefficient estimates produced by the digital twin (investigation 2) are shown in Figure 11. This analysis was applied only to the digital-twin model, since the data-driven model did not contain an explicit wear parameter and the physics-based model used a fixed wear coefficient obtained from historical training data. The figure presents representative ID and OOD test trajectories.

The RUL predictions for the data-driven, physics-based, and digital-twin models (investigation 3) are presented in Figure 12. The results are shown for representative ID and OOD test samples, with the OOD sample corresponding to a significantly accelerated wear scenario. These trajectory-level results provide a visual comparison of each model's prediction behaviour over the brake life. To summarise RUL prediction accuracy across all test trajectories, RMSE values are reported in Table 7. The table presents the RMSE for each of the ten ID and ten OOD test runs, as well as the mean RMSE for each predictive maintenance model. RMSE values are reported in number of braking events.

Table 7. Root-mean-square error (RMSE) performance leaderboard for remaining useful life (RUL) predictions across ten individual runs for in-distribution (ID) and out-of-distribution (OOD) test sets. RMSE values are reported in braking events.

Model Name	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	Mean
In-Distribution (ID) Test Set											
Physics-Based Model	1175.00	4875.00	2575.00	8775.00	12,775.00	9025.00	11,125.00	12,075.00	3425.00	1925.00	6775.00
Data-Driven Model	2308.80	1261.61	1625.98	2660.43	4668.73	2730.95	3782.12	4265.18	1365.51	1897.22	2656.65
Digital-Twin Model	2825.82	6409.74	2627.00	1889.20	3270.93	2560.07	2149.38	5100.67	1536.28	3474.40	3184.35
Out-of-Distribution (OOD) Test Set											
Physics-Based Model	14,688.08	32,733.00	41,295.64	30,642.70	15,020.99	25,675.00	25,275.00	24,525.00	26,825.00	26,075.00	26,275.54
Data-Driven Model	11,542.92	24,698.06	32,066.39	22,823.37	11,474.44	12,392.06	11,653.21	11,277.23	12,804.44	12,136.86	16,286.90
Digital-Twin Model	3952.08	13,117.18	10,100.63	37,815.11	5573.15	5081.98	4994.25	4440.19	5224.15	6003.57	9630.23




Figure 11. Online Archard wear-coefficient estimates produced by the digital twin for representative in-distribution and out-of-distribution test trajectories.



Figure 12. Remaining useful life estimates for the data-driven, physics-based, and digital-twin models. The top plot shows an in-distribution test sample with typical degradation behaviour, while the bottom plot shows an out-of-distribution sample with accelerated wear.

While the preceding results report trajectory-level behaviour and pointwise RUL accuracy, the convergence-sensitivity analysis (investigation 4) evaluated model performance across all test samples. In that analysis, all ten ID and ten OOD trajectories were considered. Each model was evaluated based on the braking event at which its RUL prediction remained within a prescribed error tolerance for the remainder of the trajectory. The mean convergence behaviour over ten runs per test set is shown in Figure 13, with shaded regions representing half a standard deviation over these ten runs.




Figure 13. RUL convergence sensitivity for the data-driven, physics-based, and digital-twin models, with shaded bands indicating ± 0.5 standard deviation. The top plot shows results for in-distribution test samples, while the bottom plot shows results for out-of-distribution samples.

The digital-twin ablation results (investigation 5) are presented in Figure 14. The figure compares the convergence sensitivity of a static physics baseline, a coupled-propagation digital twin, a state-updated fixed- K digital, and the full proposed digital twin under ID and OOD test conditions.




Figure 14. Convergence-sensitivity results for the digital-twin ablation study. The figure compares static physics, coupled propagation, state-updated fixed- K , and full digital-twin variants under ID and OOD test conditions.

The model-mismatch robustness results (investigation 6) are presented in Figure 15. To improve compactness and enable direct comparison across perturbation cases, all convergence-sensitivity results are presented in a single figure. The left column shows the ID test results, while the right column shows the OOD test results. From top to bottom, the rows correspond to the friction perturbation, hardness perturbation, speed-dependent wear perturbation, and combined perturbation cases, respectively.




Figure 15. Convergence-sensitivity results for the model-mismatch robustness study. The rows show friction, hardness, speed-dependent wear, and combined perturbation cases, while the columns show ID and OOD test results. Shaded bands indicate ± 0.5 standard deviation across ten runs.

4. Discussion

This section interprets the results of the six investigations outlined in Section 2 in the context of the modelling assumptions underlying each predictive maintenance approach, with particular focus on their ability to capture degradation behaviour under both nominal and anomalous conditions.

4.1. Investigation 1: Brake Volume Estimation Performance

The brake volume estimation results for the physics-based and digital-twin models (Figure 10) highlight a fundamental limitation of static, physics-based degradation models: the inability to adapt to changing operating conditions.

For the in-distribution dataset, the physics-based model tracks the degradation trajectory with reasonable accuracy, as expected, since the test sample is representative of the training data used to calibrate the wear coefficient. However, under out-of-distribution conditions with elevated wear rates, the physics-based model significantly overestimates the remaining brake volume. This is a direct consequence of using a fixed degradation parameter that no longer represents the true system behaviour.

In contrast, the digital twin maintains accurate volume estimates for both in-distribution and out-of-distribution scenarios. This performance can be attributed to three key mechanisms:

1. State updating: The digital twin continuously updates its estimate of brake volume through reverse inference, ensuring consistency with observed system behaviour.

2. Parameter adaptation: The Archard wear coefficient is updated online, allowing the model to adjust to changing degradation rates.
3. Dynamics-wear coupling: Unlike the physics-based model, the digital twin captures the interaction between degradation and system dynamics, enabling accurate prediction of nonlinear wear behaviour from dynamics–wear coupling.

The ability to capture this nonlinear degradation behaviour represents a key advantage of the digital-twin approach.

4.2. Investigation 2: Wear Coefficient Estimation

Only the digital-twin model can estimate the underlying Archard wear coefficient from operational data. This is enabled by the coupling between system dynamics and degradation, which allows the wear parameter to be inferred through inverse modelling.

As shown in Figure 11, the digital twin converges towards the true wear coefficient for both in-distribution and out-of-distribution datasets. For in-distribution data, convergence occurs relatively early in the brake life, since the initial parameter estimate, derived from historical data, is already close to the true value. For out-of-distribution data, convergence occurs later but still sufficiently before failure to enable meaningful prognosis.

This result demonstrates that the digital twin can infer latent degradation parameters that are not directly measurable in practice. This reduces reliance on extensive experimental calibration and enables adaptation to changing wear conditions during operation.

4.3. Investigation 3: RUL Estimation Performance

The RUL estimation results follow directly from the observed volume estimation behaviour. The physics-based model, which underestimates the wear rate under out-of-distribution conditions, produces overly optimistic RUL predictions. In contrast, the digital twin provides accurate RUL estimates across both datasets.

For out-of-distribution scenarios, the digital twin initially exhibits optimistic RUL predictions due to model initialisation based on training data with longer lifetimes. However, the model rapidly corrects this bias as additional operational data become available, demonstrating its ability to converge to accurate predictions even when initial assumptions are incorrect.

The data-driven model exhibits an RUL prediction accuracy limitation. While it generally captures the overall trend of degradation, its predictions are significantly more variable. This variability arises from the fact that RUL is inferred directly from instantaneous sensor measurements, which may correspond to multiple plausible degradation states. As a result, the model produces noisy predictions and struggles to provide stable estimates, particularly under out-of-distribution conditions. Furthermore, the increased variability observed in out-of-distribution scenarios indicates limited generalisation capability, as the model has not learned the underlying degradation physics required to extrapolate beyond the training data.

When referring to the summary statistics for all runs in Table 7, it is clear that the digital-twin and data-driven methods perform similarly for in-distribution test datasets. However, as discussed above, the digital twin provides substantially lower RUL error than the data-driven and physics-based methods for the out-of-distribution test dataset. This implies that the digital twin is a more robust methodology, performing comparably to the best method under in-distribution test scenarios, but significantly outperforming other methods under out-of-distribution test scenarios.

4.4. Investigation 4: RUL Convergence Sensitivity

The convergence-sensitivity analysis in Figure 13 provides a more practical evaluation of model performance by assessing how soon reliable RUL predictions are obtained.

For all models, convergence occurs earlier and with lower variability under in-distribution conditions compared to out-of-distribution conditions. This highlights the increased difficulty of predicting degradation behaviour that deviates from historical trends.

Across the evaluated thresholds, the digital twin generally converges earlier than both the physics-based and data-driven approaches. It achieves accurate RUL predictions significantly earlier in the asset life. Even under out-of-distribution conditions, the digital twin converges to practically useful accuracy levels well before failure.

The physics-based model shows poor convergence performance under out-of-distribution conditions, reflecting its inability to adapt to changing degradation rates. The data-driven model exhibits poor convergence performance due to the stochastic nature of its predictions. However, when considering the data-driven method's overall trend-tracking results in Figure 12, smoothing strategies could improve its practical performance.

These results demonstrate that the primary advantage of the digital twin lies not only in improved final accuracy, but also in earlier and more reliable convergence to actionable predictions.

4.5. Investigation 5: Digital-Twin Ablation Study

The ablation study, shown in Figure 14, clarifies the individual contribution of the digital-twin components. The coupled propagation variant improves over the static physics baseline because it allows the predicted braking response to evolve with brake condition. This confirms that dynamics–wear coupling is important for brake wear prognosis. The state-updated fixed- K variant further improves interpretability by periodically correcting the estimated brake volume using operational temperature measurements.

However, Figure 14a shows that the full digital twin is not uniformly superior under ID conditions. When the historical wear parameter is already representative of the test distribution, the fixed- K variants can perform competitively because they avoid variance introduced by online parameter re-estimation. This result is important because it shows that online adaptation is not inherently beneficial in all cases. Rather, its value depends on whether the degradation behaviour differs from the historical population.

Under OOD conditions, shown in Figure 14b, the full digital twin outperforms the ablated variants because the fixed historical wear parameter is no longer representative. In this case, state estimation alone can correct the current brake volume but cannot fully correct the future degradation rate. Online wear-parameter updating enables the full digital twin to adapt its future projections to the observed degradation behaviour, explaining its earlier convergence in the OOD test set.

4.6. Investigation 6: Prognostic Robustness Under Model Mismatch

The model-mismatch study, shown in Figure 15, addresses the limitation that the primary validation data are simulated using a modelling framework related to the digital twin. By generating additional validation datasets with perturbed friction, hardness, and speed-dependent wear behaviour, the evaluation tests whether the predictive maintenance models remain reliable when the data-generating physics differ from the nominal model used internally during prediction.

The ID results in the left column of Figure 15 show that no model dominates all perturbation cases and tolerance levels. This is expected because the training and ID testing data are generated from the same perturbed distribution, allowing fixed-parameter or data-driven models to perform competitively when the perturbation is consistent between

training and testing. However, the OOD results in the right column show a clearer separation. Across the perturbation cases, the digital twin generally converges earlier than both the data-driven and physics-based models. This indicates that the combination of state estimation and online parameter updating provides robustness when both the wear behaviour and the validation physics deviate from the nominal model.

The digital-twin-based approach therefore provides the strongest robustness under anomalous degradation conditions, particularly when the test behaviour deviates from the historical training distribution. However, the ablation results in Figure 14 also show that simpler fixed-parameter variants may perform competitively under ID conditions when the historical wear parameter is already representative.

4.7. Implications of the Results

The results highlight several important implications for predictive maintenance modelling. First, the data-generation process demonstrates that coupled dynamics–wear modelling leads to inherently nonlinear degradation behaviour. This has significant implications for prognosis, as simple forward propagation of wear models with fixed parameters may lead to substantial errors when degradation alters future system response. The limited performance of the static physics-based implementation supports this conclusion, as it cannot account for changes in braking behaviour caused by evolving brake wear.

Second, the comparison between modelling approaches highlights the limitations of single-paradigm methods. Data-driven models can perform well when the test data remain close to the training distribution, but their extrapolation capability is limited when degradation behaviour becomes anomalous. Conversely, physics-based models provide interpretability, but fixed-parameter implementations cannot correct modelling errors or adapt to changing degradation behaviour during operation. Improving either approach in isolation would likely require either substantially more representative run-to-failure data or more detailed prior knowledge of the true degradation mechanism, both of which are often difficult to obtain in industrial environments.

Third, the ablation results in Figure 14 show that the value of digital-twin components depends on the operating regime. Under ID conditions, simpler coupled or state-updated variants may perform competitively when the historical wear parameter is already representative. However, under OOD conditions, the full digital twin benefits from online wear-parameter updating, which allows future degradation projections to adapt to observed degradation behaviour. This indicates that online adaptation is most valuable when the asset deviates from the historical population.

Finally, the model-mismatch results in Figure 15 strengthen the simulation-based validation by showing that the digital twin remains robust even when the validation data are generated using perturbed friction, hardness, speed-dependent wear, and combined mismatch cases. These results suggest that the proposed framework does not rely solely on the nominal data-generation assumptions. Rather, its robustness arises from combining operational state estimation, online parameter adaptation, and dynamics-aware degradation propagation.

5. Conclusions

This work addressed the problem of brake wear prognosis in cases where degradation and vehicle dynamics evolve in a coupled manner over the asset life. Existing brake wear approaches have largely focused on diagnosis or on prognosis using fixed degradation assumptions and therefore do not adequately represent the feedback between wear, thermal behaviour, friction loss, and subsequent braking performance. To address this limitation, a digital-twin framework was proposed in which brake pad volume was inferred from oper-

ational data, the underlying wear coefficient was updated online, and future degradation was propagated through a wear-dependent vehicle dynamics model.

Using a simulated mining vehicle case study, the proposed digital twin was compared against data-driven and physics-based predictive maintenance models. The results showed that the digital twin produced accurate brake volume estimates, inferred the underlying Archard wear coefficient from operational data, and delivered stable remaining useful life predictions. The advantage of the digital twin was most evident under out-of-distribution degradation conditions, where the test trajectories differed substantially from the historical training distribution.

An ablation study was conducted to isolate the contribution of the main digital-twin components. The results showed that dynamics–wear coupling and state estimation could provide strong performance when the test data remained close to the training distribution. However, the full digital twin, which combined state estimation with online wear-parameter updating and dynamics-aware propagation, provided the strongest robustness when the degradation rate was out of distribution. This finding indicates that online adaptation is not universally beneficial under all conditions but becomes critical when the asset behaviour deviates from historical expectations. In practice, deviations between the model and the physical system are inevitable due to factors like unmodelled dynamics, external disturbances, sensor noise, manufacturing variability etc. The digital-twin approach is therefore well suited to practical applications, as it can continually update its estimates in response to observed system behaviour.

A further model-mismatch study was introduced to investigate how the predictive maintenance models performed when their internal assumptions differed from the data-generating mechanism. This was achieved by generating additional datasets with perturbed friction behaviour, perturbed material hardness, speed-dependent wear behaviour, and their combined effects. The results showed that the digital twin achieved the most robust convergence under anomalous degradation conditions across these mismatch cases, although simpler methods remained competitive in some in-distribution settings.

The main contribution of this study is therefore not only the use of a digital twin for brake wear prognosis, but the demonstration that robust prognosis in coupled degradation problems requires three elements to act together: online state estimation, online degradation-parameter identification, and explicit propagation through a dynamics–wear coupled model. The proposed framework provides a structured approach for combining these elements in systems where accumulated damage changes future system response.

Although the study was developed for the brakes of a mining load-haul-dumper, the underlying methodology is more broadly applicable to degradation problems in which system response evolves with accumulated damage. The proposed framework is therefore not limited to braking systems but can be extended to other assets where strong coupling exists between physical degradation and system dynamics.

Future work should focus on experimental validation of the proposed framework, since simulation-based validation cannot fully establish quantitative performance under real operating conditions. This remains challenging, as full life-cycle brake degradation data are required to observe the coupled effects identified in this study. Accelerated testing strategies, such as evaluating system behaviour at discrete degradation levels corresponding to the atlas maps, may provide a practical pathway to bridge this gap.

Further model development could incorporate additional dependencies of friction and wear on operating conditions such as speed and braking pressure, as suggested in the existing literature. In addition, the choice of wear formulation remains an open question. Alternative approaches, such as entropy-based models [27], predict different end-of-life behaviour compared to mechanics-based formulations. Future studies should therefore

investigate which degradation mechanisms best represent real brake systems to further improve the fidelity of digital-twin-based prognosis.

Author Contributions: Conceptualisation, L.v.E., B.E., S.S. and S.H.; methodology, L.v.E.; validation, L.v.E. and J.M.; formal analysis, L.v.E.; investigation, L.v.E.; resources, J.M.; data curation, L.v.E.; writing—original draft preparation, L.v.E.; writing—review and editing, L.v.E., J.M., B.E., S.S. and S.H.; visualisation, L.v.E. and J.M.; supervision, S.S. and S.H.; project administration, S.H.; funding acquisition, S.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Mandela Mining Precinct under the South African Mining Extraction Research, Development & Innovation (SAMERDI) Mechanised Mining Systems (MMS) and Real-time Information Management Systems (RTIMS) programmes.

Data Availability Statement: The models used to generate the data for this work are the intellectual property of Epiroc and cannot be openly shared. However, the synthetic run-to-failure simulation datasets used in this study are made available. These datasets include the data used for model training, ID testing, OOD testing, ablation analysis, and model-mismatch evaluation.

Acknowledgments: We would like to acknowledge the key role that the Real-time Information Management Systems (RTIMS) programme played in prompting, guiding and overseeing large parts of this research through their call for proposals on Equipment Condition Monitoring. We would also like to acknowledge Epiroc for allowing a model of their load-haul-dumper to be used to perform the brake wear analysis.

Conflicts of Interest: The authors declare no conflicts of interest.

Nomenclature

Symbol	Meaning
A	Brake performance atlas or lookup-map operator.
C_p	Specific heat capacity of the brake material [J/(kg K)].
C_{th}	Thermal mass of the brake system [J/K].
D_b	Brake actuator bore diameter [m].
d	Stopping distance [m].
$d_{i,k}$	Measured or simulated measured stopping distance for event k in trajectory i [m].
dT/dt	Brake temperature rate of change [K/s].
E	Vehicle kinetic energy at the start of a braking event [J].
$\hat{E}_{i,k}$	Vehicle kinetic energy estimated for event k in trajectory i .
$e_{i,k}^{RUL}$	Absolute RUL prediction error for event k in trajectory i [events].
$f_{Archard}$	Archard wear-update function.
$f_{Archard}^*$	Temperature-modified Archard wear-update function used in the digital twin.
$g\phi(\cdot)$	Neural network mapping from event-level inputs to RUL.
$H(T)$	Temperature-dependent brake pad hardness [MPa].
$H_0, \dots, H_7$	Empirical coefficients defining the temperature-dependent hardness relationship in Equation (9).
$\bar{H}$	Mean brake pad hardness calculated from mean brake temperature [MPa].
i	Trajectory or run index.
j	Forecast event index.
k	Braking event index.
k^*	First event at which the RUL prediction remains within a prescribed tolerance.
K	Dimensionless Archard wear coefficient.
$\hat{K}$	Estimated Archard wear coefficient.
K_{avg}	Mean wear coefficient calibrated from training trajectories.

K_i	Wear coefficient associated with trajectory i .
$\hat{K}_{i,k}$	Wear coefficient estimated for event k in trajectory i .
m	Vehicle mass [kg].
m_{brake}	Brake pad mass [kg].
$m_{i,k}$	Vehicle mass for event k in trajectory i [kg].
n_i	Number of braking events in trajectory i .
$N_{i,k}$	Braking event index for event k in trajectory i .
N_{fail}	Predicted failure event index.
N_p	Number of brake pads.
N_{sim}	Simulated number of braking events required to reach a specified brake volume.
N_{traj}	Number of run-to-failure trajectories.
P	Normal contact force between the brake pad and rotor [N].
P_{gen}	Heat generation rate during braking [W].
r_m	Effective friction radius [m].
$\widehat{\text{RUL}}_{i,k}$	Predicted remaining useful life for event k in trajectory i [events].
$\text{RUL}_{i,k}$	True remaining useful life for event k in trajectory i [events].
$s_{i,\text{true}}$	True simulated sensor value in channel i .
$s_{i,\text{noisy}}$	Noisy simulated sensor value in channel i .
t	Time [s].
T	Measured/simulated brake temperature [K or °C, depending on context].
$\bar{T}$	Mean brake pad temperature used for hardness calculation [K or °C, depending on context].
T_{atlas}	Atlas-predicted brake temperature [K].
$T_{i,k}$	Measured or simulated measured brake temperature for event k in trajectory i [K].
$\mathbf{u}_{i,k}$	Operating or measured input vector for event k in trajectory i .
$\mathbf{u}_j$	Sampled future operating conditions for forecast event j .
$\mathbf{x}_{i,k}$	Data-driven model input vector for event k in trajectory i .
$\tilde{\mathbf{x}}_{i,k}$	Normalised data-driven model input vector.
v	Relative sliding velocity between brake pad and rotor [m/s].
v_{vcl}	Vehicle velocity [m/s].
$v_{i,k}$	Vehicle speed for event k in trajectory i [m/s].
v_{eq}	Equivalent braking speed inferred from kinetic energy [m/s].
v_{norm}	Normalising speed used in the speed-dependent wear perturbation [m/s].
V	Brake pad volume or volume fraction.
$V_{i,k}$	Brake pad volume or volume fraction for event k in trajectory i .
$\hat{V}_{i,k}$	Estimated brake pad volume or volume fraction for event k in trajectory i .
V_j^+	Forecast brake pad volume at future event j .
$\hat{V}_{i,j}^+$	Forecast estimated brake pad volume at future event j in trajectory i .
V_{start}	Reference brake volume used for online wear-parameter estimation.
$\dot{V}(T)$	Temperature-dependent volumetric wear rate.
$y_{i,k}$	True RUL target for the data-driven model.
$\hat{y}_{i,k}$	Data-driven model RUL prediction.
$\Delta N_{i,k}$	Number of elapsed braking events used for online parameter estimation.
$\Delta V_{i,k}$	Brake pad volume loss for event k in trajectory i .
ΔV_{nom}	Nominal wear increment from the temperature-dependent Archard formulation.
ΔV_{spd}	Speed-perturbed wear increment.
$\mathcal{D}_{\text{train}}$	Training dataset.
ϵ_i	Additive Gaussian sensor noise in channel i .
ϵ	Acceptable RUL error tolerance used in convergence analysis [events].
γ_v	Speed-dependent wear multiplier.
γ_{min}	Lower bound on the speed-dependent wear multiplier.
γ_{max}	Upper bound on the speed-dependent wear multiplier.

μ_x	Feature-wise mean of the training inputs.
$\mathcal{L}(\cdot)$	Regression loss function.
$\mathcal{N}(\cdot)$	Normal distribution.
σ_i	Standard deviation of sensor noise in channel i .
σ_K	Standard deviation of the wear-coefficient distribution.
σ_x	Feature-wise standard deviation of the training inputs.
τ_b	Brake torque [N m].
θ	Terrain slope angle [rad or deg, depending on context].
$\theta_{i,k}$	Terrain slope for event k in trajectory i [rad or deg, depending on context].
ϕ	Neural network parameter vector.
ϕ^*	Optimised neural network parameter vector.
$\mu(T)$	Temperature-dependent brake friction coefficient.
$\mu_0, \dots, \mu_4$	Empirical coefficients defining the temperature-dependent friction relationship in Equation (2).
μ_K	Mean of the wear-coefficient distribution.
$\mathcal{V}$	Set of discrete brake pad volume states represented in the atlas.
ω_{slip}	Relative angular velocity between the brake pad and disc [rad/s].
$(\cdot)^+$	Forecast quantity propagated into the future.
$\oslash$	Element-wise division.

Appendix A. Perturbed Material Relationships

The friction and hardness perturbations used in the model-mismatch study are shown in Figure A1. The perturbations were selected to provide a strong but controlled stress test of the predictive maintenance models. In both cases, the perturbed curve differs from the nominal curve by approximately 20% near the lower and upper ends of the plotted temperature range, while also crossing the nominal relationship within the intermediate temperature range.




Figure A1. Nominal and perturbed friction and hardness relationships used in the model-mismatch robustness study.

Appendix B. Speed-Dependent Wear Perturbation

The speed-dependent wear perturbation scales the nominal wear increment by a multiplier based on the equivalent braking speed. The multiplier is calculated as

$$\gamma_v = \min \left[\max \left(\frac{v_{\text{eq}}}{v_{\text{norm}}}, \gamma_{\text{min}} \right), \gamma_{\text{max}} \right], \quad (\text{A1})$$

where $\gamma_{\min}$ and $\gamma_{\max}$ define lower and upper bounds on the speed multiplier. In this work, these bounds were set to

$$\gamma_{\min} = 0.25, \quad \gamma_{\max} = 2.00. \quad (\text{A2})$$

The equivalent braking speed is inferred from the sampled kinetic energy as

$$v_{\text{eq}} = \sqrt{\frac{2E}{m}}, \quad (\text{A3})$$

where E is the sampled kinetic energy and m is the vehicle mass. The normalising speed was selected as

$$v_{\text{norm}} = 15 \text{ km/h} = 4.17 \text{ m/s}. \quad (\text{A4})$$

This value corresponds to the upper nominal braking speed used during atlas generation. Therefore, braking events near 15 km/h retain approximately the nominal wear increment, while lower-speed and higher-speed events produce reduced and increased wear increments, respectively.

References

1. Feo, M.L.; Torre, M.; Tratz, P.; Battistelli, F.; Tomassetti, L.; Petracchini, F.; Guerriero, E.; Paolini, V. Laboratory and on-road testing for brake wear particle emissions: A review. *Environ. Sci. Pollut. Res.* **2023**, *30*, 100282–100300. [\[CrossRef\]](#) [\[PubMed\]](#)
2. Duan, Y.; Cheng, P.; Zhang, Z.; Chen, K.; Fei, J. Detection methods, influencing factors, and formation mechanisms of brake particle emissions from automotive brakes: A review. *J. Hazard. Mater. Adv.* **2026**, *21*, 101013. [\[CrossRef\]](#)
3. Kalhapure, V.A.; Khairnar, D.H.P. Wear mechanism and modelling for automotive brakes with influence of pressure, temperature and sliding velocity: A review article. *Eur. J. Adv. Eng. Technol.* **2018**, *5*, 333–343. [\[CrossRef\]](#)
4. Voloacă, Ș.; Badea-Romero, A.; Badea-Romero, F.; Toma, M.F. Motor vehicle brake pad wear—A review. *Vehicles* **2025**, *7*, 52. [\[CrossRef\]](#)
5. Xiao, X.; Yin, Y.; Bao, J.; Lu, L.; Feng, X. Review on the friction and wear of brake materials. *Adv. Mech. Eng.* **2016**, *8*, 1–10. [\[CrossRef\]](#)
6. Kumar, N.; Bharti, A.; Goyal, H.; Patel, K. The evolution of brake friction materials: A review. *Mater. Phys. Mech.* **2021**, *47*, 796–815. [\[CrossRef\]](#) [\[PubMed\]](#)
7. Ilie, F.; Cristescu, A.C. Tribological behavior of friction materials of a disk-brake pad braking system affected by structural changes—A review. *Materials* **2022**, *15*, 4745. [\[CrossRef\]](#) [\[PubMed\]](#)
8. Yevtushenko, A.; Kuciej, M.; Topczewska, K. Some theoretical model for determining the temperature field of a multi-disk brake. *Adv. Mech. Eng.* **2020**, *12*, 1–15. [\[CrossRef\]](#)
9. Yevtushenko, A.; Topczewska, K.; Zamojski, P. The Effect of Functionally Graded Materials on Temperature during Frictional Heating: Under Uniform Sliding. *Materials* **2021**, *14*, 4285. [\[CrossRef\]](#) [\[PubMed\]](#)
10. Xie, G.; Huang, C.; Li, Z.; Xiang, Z.; Tang, B.; Wang, X.; He, D. A coupled thermal-wear computational method and its application in high-speed train braking systems. *Tribol. Int.* **2026**, *219*, 111847. [\[CrossRef\]](#)
11. Aleksendrić, D. Neural network prediction of brake friction materials wear. *Wear* **2010**, *268*, 117–125. [\[CrossRef\]](#)
12. Chen, J.; Yu, C.; Cheng, Q.; Guan, Y.; Zhang, Q.; Li, W.; Ouyang, F.; Wang, Z. Research on friction performance and wear rate prediction of high-speed train brake pads. *Wear* **2023**, *514–515*, 204564. [\[CrossRef\]](#)
13. Yu, C.; Li, W.; Guo, Y.; Sun, X.; Hong, F.; Sun, N.; Zhang, Q. Research on wear rate of train brake pads driven by small sample data. *Wear* **2024**, *536–537*, 205169. [\[CrossRef\]](#)
14. Zheng, K.; Le, Q.; Pan, L.; Huang, J. Friction and wear prediction of copper-free resin-based brake materials: A hybrid PSO-FPA-BP neural network approach. *Wear* **2026**, *589*, 206536. [\[CrossRef\]](#)
15. Li, Y.; Liu, B.; Zhang, Z.; Ouyang, W.; Shi, W.; Chang, Z.; Hu, Z.; Zhang, K.; Wei, B. Physics-informed machine learning for explainable prediction of friction and wear in carbon-ceramic composites. *Ceram. Int.* **2025**, *51*, 56400–56416. [\[CrossRef\]](#)
16. Ferreiro, S.; Arnaiz, A. Prognostics applied to aircraft line maintenance: Brake wear prediction based on Bayesian Networks. *IFAC Proc. Vol.* **2010**, *43*, 146–151. [\[CrossRef\]](#)
17. Guerin, F.; Barreau, M.; Cloupet, S.; Hersant, J.; Hambli, R. Bayesian estimation of degradation model defined by a Wiener process—Application on disc brake wear. *IFAC Proc. Vol.* **2010**, *43*, 74–79. [\[CrossRef\]](#)

18. Jensen, K.M.; Santos, I.F.; Corstens, H.J.P. Estimation of brake pad wear and remaining useful life from fused sensor system, statistical data processing, and passenger car longitudinal dynamics. *Wear* **2024**, 538–539, 205220. [[CrossRef](#)]
19. Jensen, K.M.; Santos, I.F.; Corstens, H.J.P. Prediction of brake pad wear and remaining useful life considering varying vehicle mass and an experimental holistic approach. *Wear* **2024**, 552–553, 205433. [[CrossRef](#)]
20. Wang, Y.; Yang, X.; Huang, Y.; Sun, C.; Shi, W. Thermo-stress-wear multiphysics coupling mechanism in drum brakes: A comparative study on material fatigue life through simulation and experiment. *Tribol. Int.* **2025**, 212, 110904. [[CrossRef](#)]
21. Zhang, H.; Goltsberg, R.; Etsion, I. Modeling adhesive wear in asperity and rough surface contacts: A review. *Materials* **2022**, 15, 6855. [[CrossRef](#)] [[PubMed](#)]
22. Rajesh, P.; Manikandan, N.; Ramshankar, C.; Vishwanathan, T.; Sathishkumar, C. Digital twin of an automotive brake pad for predictive maintenance. *Procedia Comput. Sci.* **2019**, 165, 18–24. [[CrossRef](#)]
23. Yevtushenko, A.; Topczewska, K.; Kuciej, M. Analytical determination of the brake temperature mode during repetitive short-term braking. *Materials* **2021**, 14, 1912. [[CrossRef](#)] [[PubMed](#)]
24. Bellini, C.; Di Cocco, V.; Iacoviello, D.; Iacoviello, F. Temperature influence on brake pad friction coefficient modelisation. *Materials* **2023**, 17, 189. [[CrossRef](#)] [[PubMed](#)]
25. Saiz, C.B.; Ingrassia, T.; Nigrelli, V.; Ricotta, V. Thermal stress analysis of different full and ventilated disc brakes. *Frat. Integrità Strutt.* **2015**, 9, 608–621. [[CrossRef](#)]
26. Marko, M.D.; Kyle, J.P.; Wang, Y.S.; Terrell, E.J. Tribological investigations of the load, temperature, and time dependence of wear in sliding contact. *PLoS ONE* **2017**, 12, e0175198. [[CrossRef](#)] [[PubMed](#)]
27. Koottaparambil, L.; Khonsari, M. A unified treatment of tribo-components degradation using thermodynamics framework: A review on adhesive wear. *Entropy* **2021**, 23, 1329. [[CrossRef](#)] [[PubMed](#)]
28. MathWorks. Disc Brake Block. 2025. Available online: <https://www.mathworks.com/help/sdl/ref/discbrake.html> (accessed on 30 March 2026).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.