Hybrid As-Operated Digital Twin of an Aircraft Brake

Vinayak Chandran[1], Vikram Pasumarti[2], Stefano Sinisi[2], Roberta Cumbo[2], Bill May[3], and Alessandro Ulisse[2]

[1]*Collins Aerospace, North Gate Office, Bengaluru, Karnataka, 560064, India*
vinayakvijaya.chandran@collins.com

[2]*Collins Aerospace, Piazza dell'Indipendenza, 23, 00185 Rome, Italy*
vikram.pasumarti@collins.com
stefano.sinisi@collins.com
roberta.cumbo@collins.com
alessandro.ulisse@collins.com

[3]*Collins Aerospace, 101 Waco St, Troy, Ohio, 45373, USA*
bill.may@collins.com

ABSTRACT

This paper proposes a Digital Twin-based approach to predict friction wear of an aircraft braking system between required brake overhauls, based on individual aircraft operating conditions. The classification of Condition-Based Maintenance (CBM) as a Type III approach in SAE ARP6887 introduces significant challenges for purely statistical or data-driven algorithms. Models must be component-specific and sensitive to varying operational conditions, leading to domain shift, data sparsity, and limited generalization across fleets. Ensuring robustness, interpretability, and certifiability under these constraints are open research challenge. Conventional data analytics approaches are of limited use in scenarios characterized by a lack of run-to-failure data. On the other hand, the usage of model-based approaches can be limited due to complex physics modeling. Hybrid technologies are a more recent research area in the field of CBM, integrating prior physical knowledge with Machine Learning (ML) solutions. In this context, the As-Operated Digital Twin paradigm has emerged as a promising framework, representing the virtual counterpart of physical assets to reflect actual in-service condition and usage history. It can provide several benefits in the field of CBM, as it provides real-time insights into the health and degradation status of the monitored component and reduces the common challenges of a lack of run-to-failure data and a lack of collocated sensors with respect to the source of degradation. This article proposes a Hybrid As-Operated Digital Twin architecture, relying on the Archard equation to model the degradation phenomenon. The unknown degradation parameters of Archard's law are usually modeled with empirical equations based on laboratory data, limiting the validity of the model to the laboratory domain. To improve model generalization and reduce the number of experiments to fit the degradation parameters, a physics-informed Recurrent Neural Network (RNN) model is proposed. The proposed Hybrid As-Operated Digital Twin includes the physics-informed RNN model and the heat-sink thermal model of the brake system. The robustness of the model is studied, employing Monte Carlo methods for uncertainty quantification. The discussed methodology is demonstrated on a sample aircraft brake model with laboratory-domain data.

1. INTRODUCTION

Aircraft maintenance is largely reactive or scheduled by generic intervals, leading to high costs, potential flight delays, and safety risks due to a lack of real-time, operational insight into aircraft systems. The aerospace sector faces: unexpected failures, causing operational delays or safety risks; high maintenance costs, and limitations of current time-based maintenance. Scheduled maintenance may not detect fast degradations or abnormal operating conditions, leading to either delayed interventions or unnecessary part replacements. Condition-based Maintenance (CBM) identifies potential failures before they occur, enabling a proactive approach to maintenance scheduling and thereby reducing downtime and optimizing future flight legs based on actual aircraft conditions rather than a fixed schedule. Moreover, CBM has a positive impact on the sustainability of the aerospace industry, as efficiently maintained aircraft consume less fuel and produce fewer discarded parts by optimizing time in operation.

The prognosis of a mechanical component is normally performed by analyzing historical in-field and Maintenance, Repair, and Overhaul (MRO) data or by using advanced

---

analytical solutions. The development of a Digital Twin (Thelen, Zhang, Fink, Lu, Ghosh, Youn, Todd, Mahadevan, Hu, and Hu, 2022) for predictive maintenance can improve the current prognostic approach by using real-time data and simulations to forecast failures and optimize maintenance schedules, enabling the bidirectional interaction between virtual and physical models. This article proposes the implementation of the As-Operated Digital Twin (Michael, Pfeiffer, Rumpe, & Wortmann, (2022)), (Hartwell, Montana, Jacobs, Kadirkamanathan, Ameri & Mills, (2024)), which can be associated to one or more failure causes and can correlate these failures, e.g. wear, oxidation, with operational data, evaluating their impact at component or subsystem level. In general, system degradation leading to a specific failure can be modeled using known physical laws, whose parameters are identified through empirical functions calibrated on offline test data. However, reduced accuracy in degradation prediction may be observed due to the influence of unknown phenomena occurring under real-world operating conditions. As an alternative to purely physics-based models, Machine Learning (ML) approaches are adopted to overcome these limitations. However, their applicability may be limited by the need for large amounts of data. Physics-informed ML or hybrid modeling is an emerging research area that aims to integrate prior physical knowledge with ML techniques. A hybrid model (Li, Zhang, Li, & Si, 2024) can *1)* provide real-time insights into the health and degradation status of the monitored component; and *2)* reduce common challenges such as the lack of run-to-failure field data and the lack of sensors at the site of faults.

This paper proposes a Hybrid As-Operated Digital Twin architecture to predict the friction wear of an aircraft braking system between required brake overhauls based on individual aircraft operating conditions. Brake degradation is predicted using the Archard equation (Archard, 1953), a first-principles law for estimating friction wear. The unknown degradation parameters of the Archard law are usually calibrated using laboratory test data, limiting the validity of the model to the laboratory domain. To improve model generalization and reduce the number of experiments required to fit the degradation parameters, a physics-informed Recurrent Neural Network (RNN) model is proposed. A RNN architecture incorporating physics-based constraints has been proposed in the literature (Yucesan & Viana, 2020) for modeling cumulative damage due to fatigue loading. The work demonstrated that such a physics-informed solution is particularly effective in scenarios where degradation evolves monotonically, the system state is only partially observed, and model calibration relies on sparse measurements.

In this work, a similar physics-informed RNN framework is employed to address the problem of friction wear. Moreover, the robustness of the Hybrid Digital Twin is evaluated using Monte Carlo methods (Gal & Ghahramani, 2016; Helton & Davis, 2003) for uncertainty quantification. The proposed methodology is demonstrated on a sample aircraft brake model with laboratory-domain data for training and performance evaluation.

2. BRAKE SYSTEM CHALLENGES

Predicting the friction wear of an aircraft braking system in operation presents several known challenges. Significant variation in friction wear is observed between brakes of the same design operating on the same aircraft platform. Observational studies of friction wear data have identified sources of variation grouped by three main categories: *1)* aircraft configuration influences (variant, engine, etc.); *2)* environmental influences (temperature, humidity, etc.); and *3)* operational influences (thrust reverser level, number of brake applications, etc.). Whereas correlation analyses have deepened the understanding of these effects on the resulting friction wear, practical attempts at building a comprehensive predictive model using data-driven approaches remain challenging.

Friction wear testing in the laboratory is more controlled than in field operations; however, significant sources of variation remain and influence the resulting data. Examples of such sources include variations in test loading and environmental conditions due to long and complex test procedures (Blau, 2017).

3. ARCHARD EQUATION FOR FRICTION WEAR ESTIMATION

The prediction of friction-induced wear in aircraft braking systems requires establishing a relationship between microscopic degradation mechanisms and the macroscopic conditions governing the braking interface. In this context, the Archard law (Archard, 1953) helps bridge this scale separation by relating material loss to system-level quantities such as normal load and sliding motion. In its standard form, the law defines cumulative material loss in terms of the wear volume V, given by:

$$V = k\frac{Fd}{H} \tag{1}$$

where k is the dimensionless Archard coefficient, F is the applied normal load, d is the sliding distance, and H is material hardness. For application to braking events, Eq. (1) is rewritten in terms of frictional power to define the wear degradation over a time window Δt as:

$$\Delta h = \int_0^{\Delta t} \frac{k}{HA} P(t)dt = \int_0^{\Delta t} \overline{W} P(t)dt \tag{2}$$

Here, Δh denotes the wear depth accumulated over the interval Δt, obtained by distributing the worn volume over an effective contact area A, and $P(t)$ is the dissipated frictional power. In Eq. (2), the wear coefficient $\overline{W}$ is defined, which is function of the brake temperature T, varying in time, as described by Olesiak, Pyryev, & Yevtushenko (1997).

Eq. (2) can be written as:

$$\Delta h = \int_0^{\Delta t} \bar{W}(t)P(t)dt = \int_0^{\Delta t} \bar{W}(t)\tau\omega(t)dt \tag{3}$$

where τ denotes the shaft torque, ω is the angular velocity and t defines the time dependency. For each braking event, τ is assumed to remain constant, while $\omega(t)$ is approximated as varying linearly over the braking duration Δt according to:

$$\omega(t) = \omega_i \left(1 - \alpha \frac{t}{\Delta t}\right) \tag{4}$$

where ω_i and ω_f denote the initial and final angular velocities, respectively, and the dimensionless parameter α is defined as:

$$\alpha = \frac{\omega_i - \omega_f}{\omega_i} \tag{5}$$

The total accumulated wear depth h_{Total} is obtained by summing the contributions from individual braking events:

$$h_{Total} = h_0 + \sum_{i=1}^{N} \Delta h_i \tag{6}$$

where h_0 is the initial wear measurement and Δh_i denotes the wear increment associated with the i^{th} braking stop. N denotes the total number of braking events within the operational window. However, while the Archard wear law provides a useful first-order approximation, it has several limitations in real-world braking applications:

- Transitions between wear regimes can occur with relatively modest changes in contact pressure, sliding speed, or interfacial conditions, leading to qualitatively different wear responses rather than smooth parameter variations (Meng & Ludema, 1995). Such transitions introduce abrupt changes in wear behavior that cannot be fully captured by smooth empirical functions of temperature alone, as in Olesiak, Pyryev, & Yevtushenko (1997).
- Tribology literature further shows that terms like $\bar{W}$ are often used as effective parameters to lump unresolved interfacial phenomena, such as third-body effects, contact mechanics variability, surface interactions, and temperature-dependent effects, into simplified wear laws (Godet, 1984). Representing these complex and coupled interactions through a predefined parametric expression constrains the model ability to adapt to variations in operating and environmental conditions beyond those captured during calibration.

These considerations suggest that, while the Archard accumulation framework provides a physically interpretable structure for wear estimation, a more flexible and state-dependent representation of the effective wear coefficient may be required to capture the variability observed across both controlled laboratory testing and in-service operating conditions.

4. HYBRID MODELING FOR WEAR PREDICTION

In cumulative degradation models, the accumulation law is typically well established, whereas the driving quantities governing degradation are often not directly observable or are only partially specified through empirical representations. This, combined with the challenges identified in Sections 2 and 3, hinders the accurate definition, calibration, or validation of such modes. This is particularly true in the absence of large and sufficiently rich data and motivates the adoption of a hybrid modeling approach.

In this contribution, the physics-informed RNN approach proposed in Yucesan & Viana, 2020 is adapted to address the problem of friction wear, by constraining the RNN parameters with the Archard law (Fig. 1). While the accumulation law is retained, the empirical temperature-dependent representation of $\bar{W}$ defined in Olesiak, Pyryev, & Yevtushenko (1997) is described by a Neural Network (NN) model, which estimates the effective wear coefficient as a function of the temperature provided by the brake thermal model and the previously accumulated wear state.

Temperature is included due to its well-established influence on frictional wear, while the accumulated wear state represents the evolving surface condition of the braking interface. As material is removed, changes in contact geometry and interfacial interactions influence subsequent wear rates. The predicted coefficient is embedded directly within the Archard wear integration to compute the corresponding wear increment, ensuring consistency with the accumulation structure. In this manner, learning is restricted to the representation of the degradation-driving coefficient, while cumulative wear evolution remains governed by first-principles modeling. To preserve physical consistency, the network output is constrained to be non-negative, ensuring monotonic wear increments.

4.1. Model Training and Wear Prediction Results

Landing gear systems undergo repeated use during operations, characterized by stops and cycles. A stop denotes an individual braking event in which the brakes are applied to decelerate the aircraft, either to a complete rest or to a

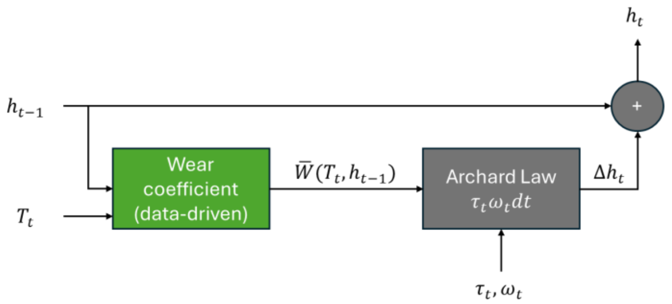

Figure 1. Schematic representation of the proposed physics-informed RNN architecture for friction wear.

reduced velocity. A cycle comprises a sequence of stops corresponding to a complete gate-to-gate aircraft operation. The data used in this study were obtained from controlled experimental tests designed to investigate the wear rate of a specific landing gear system. The data set consists of more than 600 consecutive operating cycles composed of a predefined sequence of stops. Wear measurements are obtained using a wear pin, whose readings increase monotonically with material loss, thereby providing a direct indication of accumulated wear. Wear pin measurements are recorded at the end of each cycle, so training and evaluation are performed at the cycle level. The learned wear coefficient is used within the Archard formulation to compute the corresponding cycle-level wear increment, and cumulative wear is updated sequentially. For the purposes of model development and evaluation, 80% of the available data are used for training, while the remaining data are reserved for validation and testing. Furthermore, the model is also evaluated on a test set of 175 consecutive operating cycles corresponding to another brake design.

To increase the effective number of training sequences and expose the model to different initial wear states, a windowing approach with a predefined length w_l is applied to the time-series data, i.e., the sequence of stops and cycles. w_l is defined in terms of operating cycles. Choosing the size of the training window size is a balancing act: a large window can capture long-term wear patterns but leads to increased training time; whereas a shorter window length may not capture sufficient temporal dependencies. In this work, we distinguish between the training and forecasting (prediction) window sizes, denoted as w_l^t and w_l^p, respectively. w_l^p can assume a different value than w_l^t; however, it is observed that the estimation accuracy decreases when $w_l^p \gg w_l^t$, since long-term dependencies are not sufficiently learned.

The analyses presented in this section are performed with $w_l^t = 15\, cycles$. The input features of the model, as indicated in Section 3, are τ, ω, and T, defined as time- series inputs, while the target is the accumulated wear at the end of each training window. The training process incorporates a scheduler to reduce the learning rate based on the validation performance, as well as an early stopping functionality which stops the training when validation loss fails to improve. The selected loss function is Mean Squared Error (MSE). Fig. 2a and Fig. 2c show respectively the relative percentage error, i.e. $\frac{|pred-meas|}{meas} * \mathbf{100}$, and the absolute error, i.e. $|\boldsymbol{pred - meas}|$, evaluated for each cycle and for different w_l^p values. Since the accumulated wear measurement is monotonically increasing, comparable absolute errors lead to larger percentage errors in the initial cycles than the later ones. The results of Fig. 2a and Fig. 2c, summarized in Table 1, show that the prediction accuracy decreases when $w_l^p \gg w_l^t$. The generalization of the model has also been tested by evaluating it on another dataset corresponding to a different brake design (Figure 2b and Fig. 2d). The two brakes differ in heat-sink properties, rotor count, associated aircraft etc. and estimation accuracy drops significantly on this design (Table 2). The acceptable accuracy threshold is defined as an MAE of 0.11, corresponding to a 5% error on the cumulative wear.

The results demonstrate that the proposed model architecture can accurately model the underlying dynamics of wear evolution for a given brake design. By replacing the fixed empirical wear formulation with a learned wear coefficient defined by temperature and wear state, the model was able to capture the operational and experimental variability observed in laboratory data. As expected, lower accuracy is observed when the model is evaluated on a different brake design. Model generalization will be explored in a future work by investigating transfer learning techniques. A reduced accuracy is also observed when $w_l^p \gg w_l^t$, indicating that long-term dependencies must be adequately represented during training to maintain stable forecasts. The following section focuses on the uncertainty quantification to explore robustness of the trained model.

Table 1. Mean Square Error (MSE) and Mean Absolute Error (MAE) of the wear prediction over all cycles and the last prediction for the training/validation brake design.

Prediction Window Size	Over all Cycles		On the Last Prediction	
	MSE	MAE	MSE	MAE
10 Cycles	2.6e-05	0.0042	8.2e-05	0.0091
50 Cycles	4.1e-04	0.0176	2.1e-03	0.0455

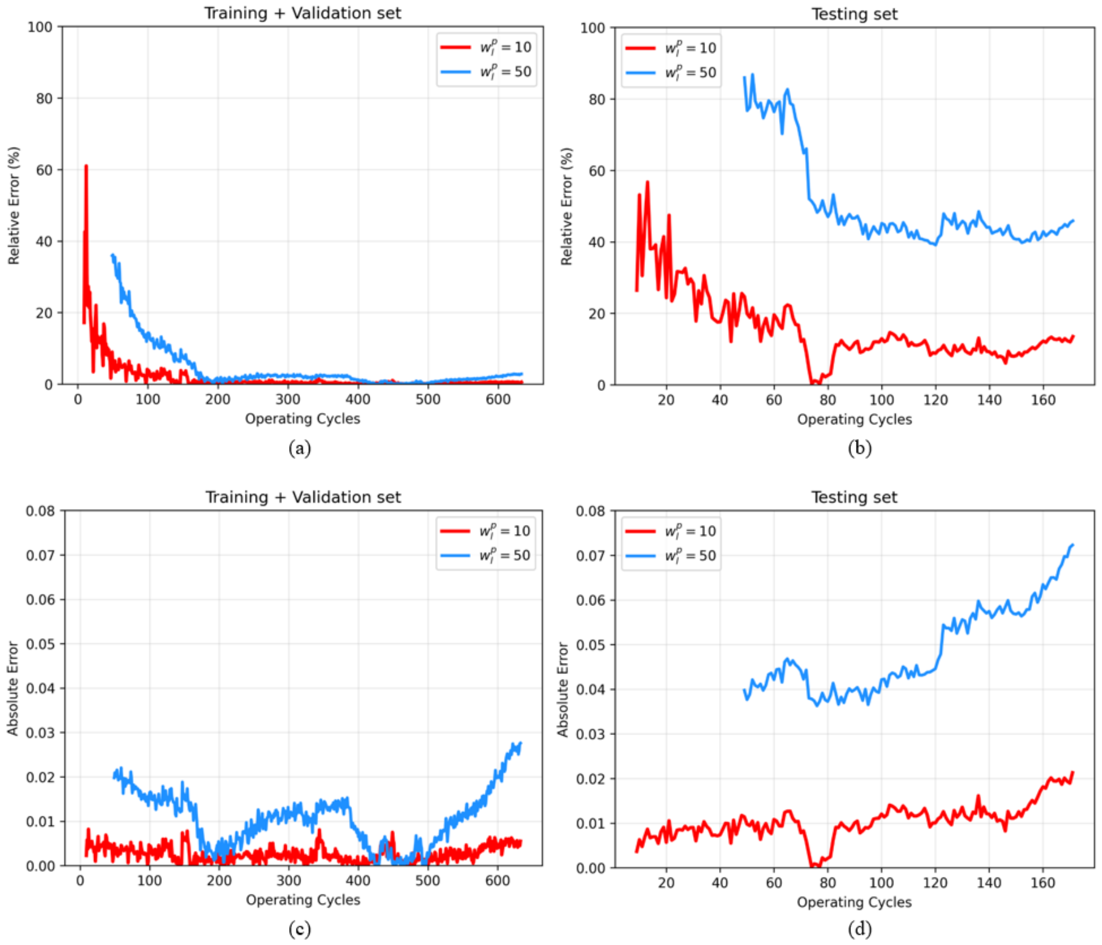

(a)

(b)

(c)

(d)

Figure 2. Per-Cycle relative percentage error and absolute error comparison for $\boldsymbol{w_l^p}$=10 cycles and $\boldsymbol{w_l^p}$=50 cycles for (a) (c) Training + Validation set and (b) (d) Testing set.

Table 2. Mean Square Error (MSE) and Mean Absolute Error (MAE) of the wear prediction over all cycles and the last prediction on different brake design.

Prediction Window Size	Over all Cycles		On the Last Prediction	
	MSE	MAE	MSE	MAE
10 Cycles	3.3e-04	0.0171	1.2e-03	0.0351
50 Cycles	6.6e-03	0.0802	1.4e-02	0.1192

4.2. Uncertainty Quantification

Degradation processes are influenced by variability in operating conditions, limited observability of system states, and lack of model fidelity. For these reasons, in practical prognostics applications, providing only a deterministic estimation of degradation is often insufficient, as uncertainty plays a critical role in maintenance planning, risk assessment, and decision-making. As highlighted in the prognostics and health management literature (Peng, Dong, & Zuo, 2010), uncertainty quantification is therefore a necessary component of predictive models to ensure robust and interpretable prognostic outcomes.

In this work, both epistemic uncertainties, associated with the data-driven component of the Hybrid As-Operated Digital Twin, and aleatory uncertainty, arising from variability in the model inputs, are considered. Model-related uncertainty is addressed using Monte Carlo dropout during inference. Originally proposed as a regularization technique, dropout has also been shown to provide an effective approximation of Bayesian uncertainty in neural networks when applied at prediction time (Gal & Ghahramani, 2016). By randomly deactivating a subset of network parameters at prediction time, multiple forward passes generate an ensemble of wear predictions that reflects uncertainty in the learned model parameters. A dropout rate of 0.0005 is used for this analysis, and the distribution of wear prediction error obtained from repeated evaluations is shown in Figure 3.

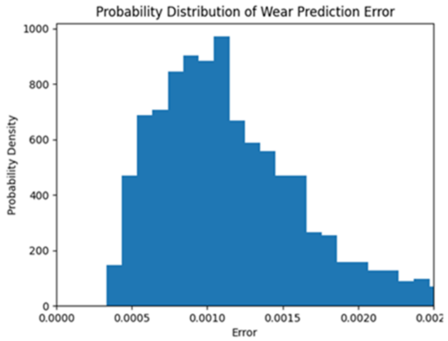

Figure 3. Results of Monte Carlo dropout–based uncertainty analysis – Distribution of cumulative wear prediction error obtained from multiple stochastic forward passes with a dropout rate of 0.0005

In real operating environments, physical parameters are subject to measurement noise and operational variability, which can propagate through the Digital Twin and affect wear estimates. Input-related uncertainty is assessed through Monte Carlo simulation (Helton & Davis, 2003; Saltelli et al., 2008) by introducing stochastic perturbations to τ and ω, the primary inputs defined in Eq. (3). Three noise levels are considered for each input, defined by increasing standard deviations corresponding to low, medium, and high variability scenarios. For each noise level, multiple simulations are performed under independently perturbed inputs, and the corresponding wear prediction errors are recorded. The distribution of cumulative wear prediction errors corresponding to the high-torque uncertainty case is shown in Figure 4, along with the corresponding uncertainty levels for both τ and ω and the associated mean absolute error values.

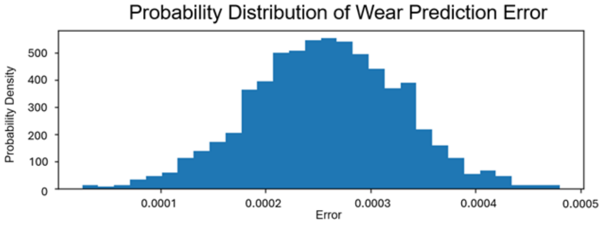

		STD [N/m]	Mean Absolute Error
Torque	Low	500	1.5e-4
	Medium	2500	1.9e-4
	High	5000	2.7e-4

		STD [rad/s]	Mean Absolute Error
Angular Velocity	Low	2	1.4e-4
	Medium	10	2.4e-4
	High	20	6.5e-4

Figure 4. Distribution of cumulative wear prediction error under input uncertainty for the high-torque case. The uncertainty levels for torque and angular velocity and the corresponding mean absolute error values are indicated.

The results indicate that, for the scenarios considered, the proposed degradation model exhibits limited sensitivity to both model-related and input-related uncertainties within the investigated variability levels. These findings provide an initial assessment of prediction variability and support the reliability of the proposed Hybrid As-Operated Digital Twin architecture within the examined operating regime.

5. CONCLUSIONS

This work presented a Hybrid As-Operated Digital Twin architecture for predicting friction wear in aircraft braking systems. Although the Archard formulation provides a physically interpretable framework for wear estimation, the difficulty of calibrating the wear coefficient without extensive data limits its applicability in real-world scenario. This motivates hybrid modeling strategies that integrate first-principles law with data-driven learning.

In the proposed architecture, the cumulative wear is modeled as a Hybrid RNN with the wear coefficient represented by a NN that takes brake temperature as a dynamic input and the previously predicted wear state as its recurrent input. This formulation enables the model to learn complex, possibly nonlinear and state-dependent wear behavior while preserving the physically grounded accumulation law of the Archard equation.

The model was trained and evaluated using laboratory friction-test data, which provided a controlled yet sufficiently complex environment given the sensitivity of wear to thermal and operational variations. Results showed that the proposed Hybrid As-Operated Digital Twin effectively captured the underlying wear dynamics for the investigated brake design. As expected, prediction accuracy decreased when the model was applied to a different brake design and when the forecasting window significantly exceeded the temporal patterns seen during training.

Given the sensitivity of degradation models to operational and environmental variability, uncertainty quantification is essential to ensure robust prognostic performance. The model uncertainty was approximated through Bayesian inference using Monte Carlo dropout at prediction, while input uncertainty was assessed via stochastic perturbations of torque and angular velocity under low, medium, and high variability scenarios. Within the ranges investigated, cumulative wear predictions exhibited limited sensitivity to both model and input variability, supporting the stability and robustness of the framework in the operating range.

Future work will focus on extending the inference capability of the proposed architecture to better reflect the complexity of in-service operational data. Improving cross-design generalization across different brake architectures will be a key objective, potentially through transfer learning strategies or domain adaptation. Finally, while Monte Carlo-based techniques offer a practical first step toward predictive uncertainty estimation, uncertainty quantification for Hybrid models remains an active research area.

REFERENCES

Archard, J. F. (1953). Contact and rubbing of flat surfaces. Journal of Applied Physics, 24, 981–988. doi:10.1063/1.1721448

Blau, P. J. (2017). Lessons learned from the test-to-test variability of different types of wear data. Wear, 376–377, 1830–1840.

Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. Proceedings of the 33rd International Conference on Machine Learning, 1050–1059.

Godet, M. (1984). The third-body approach: A mechanical view of wear. Wear, 100(1–3), 437–452.

Hartwell, A., Montana, F., Jacobs, W., Kadirkamanathan, V., Ameri, N., & Mills, A. R. (2024). *Distributed digital twins for health monitoring: resource constrained aero-engine fleet management*. The Aeronautical Journal, 128(1325), 1556-1575.

Helton, J. C., & Davis, F. J. (2003). Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. Reliability Engineering & System Safety, 81(1), 23–69.

Li, H., Zhang, Z., Li, T., & Si, X. (2024). A review on physics-informed data-driven remaining useful life prediction: Challenges and opportunities. Mechanical Systems and Signal Processing, 209, 111120. doi:10.1016/j.ymssp.2023.111120

Meng, H. C., & Ludema, K. C. (1995). Wear models and predictive equations: Their form and content. Wear, 181–183, 443–457.

Michael J., Pfeiffer J., Rumpe B., & Wortmann A. (2022, May). *Integration challenges for digital twin systems-of-systems.* In Proceedings of the 10th IEEE/ACM International Workshop on Software Engineering for Systems-of-Systems and Software Ecosystems (pp. 9-12).

Olesiak, Z., Pyryev, Y., & Yevtushenko, A. (1997). Determination of temperature and wear during braking. Wear, 210, 120–126.

Peng, Y., Dong, M., & Zuo, M. J. (2010). Current status of machine prognostics in condition-based maintenance: A review. International Journal of Advanced Manufacturing Technology, 50(1–4), 297–313. doi:10.1007/s00170-009-2485-2

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). Global sensitivity analysis: The primer. John Wiley & Sons. doi:10.1002/9780470725184

Thelen, A., Zhang, X., Fink, O., Lu, Y., Ghosh, S., Youn, B. D., Todd, M. D., Mahadevan, S., Hu, C., & Hu, Z. (2022). A comprehensive review of digital twin—Part 1: Modeling and twinning enabling technologies. Structural and Multidisciplinary Optimization, 65(12), Article 354. doi:10.1007/s00158-022-03425-4

Yucesan, Y. A., & Viana, F. A. (2020). A physics-informed neural network for wind turbine main bearing fatigue. *International Journal of Prognostics and Health Management*, *11*(1).