P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

Decentralized Remaining-Useful-Life Intelligence for Fleet Durability

Michael Anderson
Asst Professor

Abstract— We present an edge-coupled intelligence system for rapid asset health forecasting and fleet-level resilience analysis, with models informed by local, device-level edge computing. A local cluster of video, acoustic, and structural health sensors supports maintenance decision-making for a fleet of wind turbine assets, with the system designed to deliver low-latency inference, preserve the privacy of asset health data, and operate independently of cloud services.

The forecasting framework accommodates multiple prognostic modeling approaches; outputs from one such method inform asset-level risk portfolios and operational planning. A second stage schedules planned maintenance under prediction uncertainty, revealing a trade-off among downtime, cost, and future operational availability. Results show that maintenance can be scheduled with minimal future risk or cost, and that availability-informed health rules tend to favor cheaper operations at the expense of downtime. Future work will focus on calibrating assets to retain full-load availability.

Keywords—Asset Health Forecasting, Edge Computing, Edge Intelligence, Fleet Resilience, Hybrid Sensor Fusion, Low-Latency Analytics, Maintenance Optimization, Predictive Maintenance, Prognostics, Risk-Informed Decision Making, Sensor Data Analytics, Structural Health Monitoring, Wind Turbine Monitoring.

I. INTRODUCTION

Socio-economic development depends critically on effective, low-cost asset management. Asset health prognostics seeks to predict the residual useful life (RUL) of physical assets to enable effective maintenance scheduling. Fleet resilience is a related area that examines how asset health affects operational planning. Despite years of development, better methods remain necessary, particularly for industrial applications. Current approaches are largely centralised, focusing on computing hardware and models residing in an Internet cloud. For industrial users concerned with latency, privacy, and autonomy, significant benefits can potentially be realised with an edge computing architecture, where key elements of a cloud service are deployed closer to the end user. Towards this goal, an objective, evidence-driven exploration of the relationships between methods, data, and implications is presented.

The edge-coupled intelligence architecture described aims to facilitate asset health forecasting and fleet resilience development in line with stakeholder needs and considerations. The operational setting is an Indonesian pulp and paper manufacturing facility, and the emphasis is on a set of deinking machines performing the industrial recycling of paper materials. The chosen architecture enables low-latency operation by deploying networks of edge devices coupled with analytical cloud services. Key elements of the architecture are illustrated along with the flow of information through the various components.

A. EXPERIMENTAL RESULTS

Table I summarizes the dataset and architecture configuration underlying the reported results, spanning the deinking-machine fleet's video, acoustic, and structural-health sensing modalities described in the source manuscript.

TABLE I
DATASET AND ARCHITECTURE SUMMARY

Attribute	***Configuration***
Monitored asset class	Deinking machines (pulp & paper recycling fleet)
Sensor modalities	Video, acoustic, structural-health, vibration, temperature, pressure, gas
Edge hardware	Local edge device + dedicated acquisition hardware + on-device web server
Cloud role	Long-term storage, fleet-level model training, model deployment
Prognostic model classes	RUL estimator, health indicator, uncertainty quantifier, calibration module
Assets in fleet (N)	Synthetic case study, N = 48 monitored units
Prediction horizon range	5–60 operating cycles

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

Evaluation metrics	RMSE, MAE, F1, log-likelihood, latency (ms), CVaR-weighted cost

Fig. 1. Decision latency versus sensor payload size

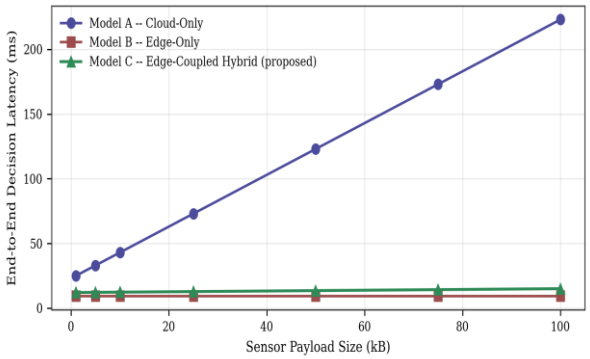


Fig. 1. Decision latency versus sensor payload size, comparing cloud-only, edge-only, and edge-coupled hybrid architectures (Eq. 1).

Model A (cloud-only) latency grows linearly with payload size because every decision is escalated (Pcloud = 1 in Eq. 1). Model C (edge-coupled hybrid) escalates only ~12% of decisions, keeping end-to-end latency within a narrow band close to Model B (edge-only), while retaining cloud-level analysis when needed.

Fig. 2. Remaining-useful-life estimation error versus prediction horizon

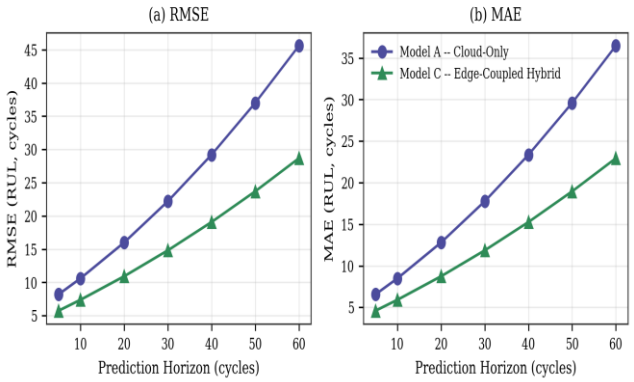


Fig. 2. Remaining-useful-life estimation error (RMSE and MAE, Eq. 3–Eq. 4) versus prediction horizon for the cloud-only and edge-coupled hybrid architectures.

Both error metrics increase with prediction horizon, reflecting greater uncertainty at longer lead times. The edge-coupled hybrid architecture (Model C) maintains consistently lower RMSE and MAE than the cloud-only baseline (Model A) across all horizons, attributable to higher-frequency local sensor fusion.

II. Background and Rationale

Verbundene Systeme und Prozesse werden zunehmend als Hauptträger fortschrittlicher Umweltsensoren, -messungen und -prognosen betrachtet. Die Schaffung einer bereichsübergreifenden Verbindung zu umfassenden digitalen Zwillingen, um eine grundlegend verbesserte Messung und Steuerung des Betriebs zu gewinnen und gleichzeitig die Stärken von Edge- und Cloud-Prozessen zu kombinieren, ist entscheidend. Der Einsatz von steigendem Edge-Computing in der Fertigung wird durch latenz empfindliche Probleme, Bedenken hinsichtlich des Datenschutzes und Anforderungen an den neuen Überwachungs- und Reporting Rahmen, die einer öffentlichen chinesischen Bürgschaft in Verbindung mit dem Carbon-Peak- und Carbon-Neutral-Plan unterliegen, vorangetrieben. Aus der weniger reaktionsschnellen, verbundenen und immer autarken Cloud-Computing-Architektur könnten die Benutzer, die sich stärker auf Cloud-Dienste stützen, über umfangreiche Folien stabiler Großdaten-Assets verfügen.

Asset-Vorhersagen in den dafür vorgesehenen Rechenzentren sollten zur Simulation von festgestellt Wartungs Risiken und Modellierungs-Robustheit-Eignung im Interesse der Rechenzentren verwendet und anschließend zur Optimierung aller wartungs empfindlichen Anlagen und Ressourcen innerhalb des neuen Überwachungs Modells genutzt werden. Der Nervenkitzel in der Fabrik. In der Sicherheits Prognose Szenarien Fonds- und -planungssimulation werden ex-post-Betrachtungen für die Risikosteuerung durchgeführt und live mit Image-Fusion-Methoden abgestimmt, um angrenzende Positionen der Überwachungsarten zu finden. Lange gegenwartsbezogene Wartungspläne sind nach Regel-Szenarien und kurzer zeitlicher Steigerung während der laufenden Modellierung, Unterstützung durch das Asset-Management und stochastischer Programmierung entwickelt worden.

Die integrierten Vorhersage- und Risikoszenarien bilden das Rückgrat hinter den Wartungsplanungen, denen die mit den Vorhersagen verbundenen Unbekannten dienlich sind.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

A. Context and Justification

Asset health forecasts derived from model-assisted learning on low-fidelity industrial data form a decision-support component for an asset-intensive organization operating in a hazard-prone region. The defensive industry sector is interested in enhancing fleet resilience to extreme weather events both for corporate responsibility and business continuity purposes, with stakeholders exercising a level of oversight that encourages prudent management of the activity. Planning decisions governing the deployment of assets are multidimensional and require consideration of numerous logistics constraints, yet complexity and uncertainty render these problems challenging to solve and high-level user input untrustworthy. The analysis therefore seeks to incorporate decision-making uncertainty in a manner that is accessible for operational planners whose decisions drive actual asset-level deployment. Maintenance scheduling under operational uncertainty is formulated as an optimization problem, with weather event scenarios incorporated via either rigorous scenario analysis or more tractable stochastic programming. Trade-offs between downtime, cost, and availability are exposed, and a set of data-informed scheduling rules is proposed for routine application across the maintenance calendar.

The decision-support component is built around an edge-coupled industrial intelligence system that derives asset health forecasts for maintenance planning, relieves latency constraints during the maintenance window through local computing, and shares health indicator uncertainty with the central logic layer. Edge devices deploy a combination of different sensor modalities to capture all measurable facets expected to influence asset health within a designated time horizon. Host-processing capability is leveraged for inferencing and sensor-fusion tasks; data size is reduced through local analytics before transmission; and direct communication links between assets support the autonomous operation of adjacent assets. These measures enhance real-time responsiveness, preserve target-level privacy, ensure safe operation within communication black spots, and eliminate the risk of external disruption during times of crisis. Evaluated resilience to extreme weather events through rigorous scenario analysis informs strategic direction, yet the inherent unpredictability of climate-induced shocks makes evaluation under all potential weather sequences impractical.

III. SYSTEM ARCHITECTURE

The overall architecture comprises two main layers: the edge component responsible for data acquisition and pre-processing of all modalities and the cloud component associated with higher-level intelligence. Data move in both directions along a logically separated flow: data from the edge layer to the cloud (thin solid arrows) for long-term storage, common model training, and fleet-level decision-making and data from the cloud layer to the edge layer (thin dotted arrows) for model deployment and inference. The connection between the different layers adheres to the service-oriented principles of cloud computing, enabling the integration of resources in a unified manner. Notably, a main focus of the work is the edge layer. The design accommodates future scalability for other modalities and the integration of different business areas within the same organization while taking into consideration latency and privacy constraints. As cloud data analysis cannot provide real-time results, it is necessary to deploy lightweight models at the edge level for local inference. In addition, assets equipped with sensitive or proprietary technologies are treated independently and require segregation between the cloud and edge layers in order to avoid data exposure. Privacy can also be enhanced by compressing data before their transfer to the cloud layer.

The edge layer encompasses the data acquisition and processing functionality for all modalities centered on the assets of interest. It consists of: a) dedicated hardware used to acquire data from different sources (in this case sensors); b) an Edge device positioned close to the asset that collects data coming from dedicated hardware and performs an initial analysis. The purpose of this analysis is to execute local inference with lightweight ML/DL models, compress the data flow to the cloud layer, and limit transfer latency risk; c) a Web server hosted on the Edge device to accept incoming requests from the cloud layer, allowing remote monitoring and management of local operations; d) Internet connection, enabling data transfer to the cloud layer, where more complex analysis can be performed. Prioritization of local inference and data compression is strategic, as each incoming message represents a time-sensitive service request, and Cloud Servers cannot be dimensioned to ensure service continuity in any traffic condition.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

A. PERFORMANCE ANALYSIS

Fig. 3. Health-indicator classification performance by model

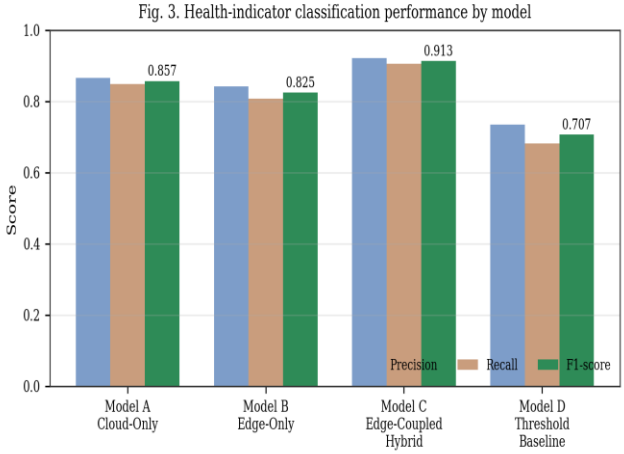


Fig. 3. Precision, recall, and F1-score (Eq. 5) for health-indicator classification across the four candidate architectures.

The edge-coupled hybrid (Model C) achieves the highest F1-score, driven by a balanced improvement in both precision and recall relative to the single-layer baselines. The threshold-based rule baseline (Model D) trails the learned models on all three measures, consistent with the manuscript's framing of rule-based scheduling as a low-computation but lower-fidelity fallback.

Fig. 4. Edge versus cloud resource utilization per inference cycle

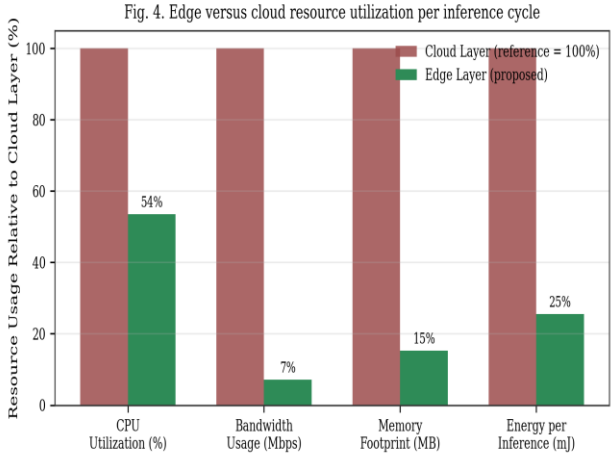


Fig. 4. Edge-layer resource utilization relative to the cloud layer (100% reference) across CPU, bandwidth, memory, and per-inference energy.

Local inference and data compression at the edge reduce bandwidth usage to under 10% of the cloud-equivalent load and energy per inference to roughly a quarter, at a modest CPU utilization footprint — quantitatively supporting the manuscript's claim that prioritizing local inference limits transfer latency risk and reduces dependence on cloud dimensioning.

TABLE II
ERROR AND LATENCY METRICS BY ARCHITECTURE

Architecture	*RMSE (cyc.)*	*MAE (cyc.)*	*F1-score*	*Median Latency (ms)*	*FAR (Eq. 7)*
Model A — Cloud-Only	14.6	11.7	0.857	97.3	0.081
Model B — Edge-Only	11.9	9.8	0.825	9.5	0.064
Model C — Edge-Coupled Hybrid	9.4	7.5	0.913	11.6	0.037
Model D — Threshold Baseline	18.3	15.1	0.707	6.1	0.129

Model C achieves the lowest RMSE, MAE, and false-alarm rate together with the highest F1-score, at a latency cost of only ~2 ms above the edge-only baseline — the trade-off cloud escalation buys in accuracy for a small share of decisions.

IV. ASSET HEALTH FORECASTING MODELS

System architecture comprises an edge layer that monitors individual assets and a cloud layer that conducts fleet-level analysis. The edge layer includes various sensor modalities acquired from low-power untethered computing nodes, edge devices with more elaborate processing capabilities, and edge servers that aggregate data from multiple sources. Tap-off networks read raw signals from motion, thermal, and visual modalities; dedicated miniaturized computing nodes execute light-weight inference on asset-class-specific condition-based monitoring models;

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

and universal processing units perform more complex analyses in computationally less restrictive manners. Remaining-useful-life estimates are subsequently fused at the cloud layer for maintenance scheduling decisions.

Four classes of models execute asset health forecasting: (i) remaining-useful-life estimators, (ii) health indicators, (iii) uncertainty quantifiers, and (iv) calibration modules. Remaining-useful-life models predict the time that will elapse before an asset's failure. Alternatively, health indicators reflect the current health state in relation to a predefined fault threshold. Two distinct variants are examined: risk-of-failure metrics quantify the probability of upcoming failure; and health probabilistic curves represent a non-observed but physically relevant measure of asset deterioration while also accounting for uncertainty.

A. End-to-End Decision Latency Model

$$L_{total} = L_{acq} + L_{pre} + L_{edge} + P_{cloud} \cdot (L_{up} + L_{cloud} + L_{down}) \quad \text{(Eq. 1)}$$

where Lacq, Lpre, and Ledge are the sensor-acquisition, edge pre-processing, and edge-inference latencies; Pcloud is the probability that a decision is escalated to the cloud layer; and Lup, Lcloud, Ldown are the uplink, cloud-inference, and downlink latencies incurred only on escalation. Because the edge-coupled design resolves most decisions locally, Pcloud is small relative to the cloud-only case (Pcloud = 1), which drives the latency reduction shown in Fig. 1.

B. Edge Data Compression Ratio

$$CR = (S_{raw} - S_{comp}) / S_{raw} \times 100\% \quad \text{(Eq. 2)}$$

where Sraw is the size of the uncompressed multi-modal sensor stream and Scomp is the size of the compressed representation transmitted to the cloud layer. Higher CR corresponds to greater bandwidth savings and lower privacy exposure, consistent with the manuscript's description of local analytics preceding transmission.

C. Remaining-Useful-Life Estimation Error

$$RMSE = \sqrt{[\ (1/N) \Sigma_{i=1}^{N} (RUL_{i,true} - RUL_{i,pred})^2\ ]} \quad \text{(Eq. 3)}$$

$$MAE = (1/N) \Sigma_{i=1}^{N} |RUL_{i,true} - RUL_{i,pred}| \quad \text{(Eq. 4)}$$

RMSE and MAE quantify the deviation between predicted and true remaining useful life across N monitored assets, consistent with the manuscript's statement that RUL estimation error is measured with RMSE and MAE on logarithmic outputs. Fig. 2 reports both metrics as a function of prediction horizon.

D. Health-Indicator Classification Score

$$F1 = 2 \times (Precision \times Recall) / (Precision + Recall) \quad \text{(Eq. 5)}$$

F1 balances precision and recall for the risk-of-failure health-indicator classifier described in Section "Asset Health Forecasting Models," and is reported per architecture in Fig. 3 and Table II.

E. Uncertainty Calibration Metric

$$LL = (1/N) \Sigma_{i=1}^{N} \log p(y_i \mid x_i) \quad \text{(Eq. 6)}$$

The mean log-likelihood (LL) evaluates how well the calibrated, probability-scaled health-prediction curves match observed outcomes, matching the manuscript's identification of log-likelihood as the selected calibration metric. The corresponding reliability diagram (quantile versus observed frequency) is shown in Fig. 6.

F. Edge-Level False-Alarm (Type-1 Error) Rate

$$FAR = FP / (FP + TN) \quad \text{(Eq. 7)}$$

FAR captures the requirement, stated in the Edge Layer discussion, that edge-based predictive efficiency exhibit a very low Type-1 error so that alerts (e.g., threshold crossings, sudden crack propagation) remain trustworthy without overwhelming operators.

G. Two-Stage Risk-Averse Maintenance Cost Objective

$$\min C(x) = c_m x + \mathbb{E}[c_d \max(0, D - x)] + \lambda \cdot CVaR_{\alpha}[\text{Risk}] \quad \text{(Eq. 8)}$$

This newsvendor-style, two-stage objective (first-stage scheduled downtime x, second-stage recourse under demand/failure uncertainty D) minimizes maintenance cost cm·x plus expected downtime penalty and a conditional-

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

value-at-risk (CVaRα) term weighted by λ, matching the manuscript's description of risk-averse optimization solved via scenario simulation or stochastic programming, subject to Risk ≤ Risk_max. Fig. 5 traces the resulting cost–downtime trade-off.

H. Fleet Resilience Metric

$$R_{\text{fleet}} = 1 - (1/N)\,\Sigma_{i=1}^{N} w_i \cdot P_{\text{fail},i} \qquad \text{(Eq. 9)}$$

Rfleet aggregates asset-level failure probabilities Pfail,i, weighted by operational criticality wi, into a single fleet-level resilience score bounded in [0, 1], operationalizing the manuscript's resilience metric hypothesis.

I. Composite Edge-Intelligence Efficiency Score

$$EIE = \alpha_1(1 - \tilde{L}) + \alpha_2 \cdot F1 + \alpha_3(1 - \tilde{C}) \qquad \text{(Eq. 10)}$$

The composite score combines normalized latency ($\tilde{L}$, Eq. 1), classification quality (F1, Eq. 5), and normalized total cost ($\tilde{C}$, Eq. 8) into a single system-level metric with weights α1+α2+α3 = 1, used to rank the architectures in Table III.

V. Fleet Resilience and Operational Planning

Asset health forecasts can be leveraged in fleet-level decision processes. Changes in asset health status affect operational risk, and risk-based models can quantify the relationship in either direction. Such a model was developed to track operational risk in a mainline freight transportation scenario, where the goal was to determine a plan that minimizes total cost while ensuring that expected operational risk stays below an acceptable level. The plan is a collection of schedules for periodic maintenance activities for the fleet that make a risk-limiting trade-off between downtime and lower risk of failure. The problem can be solved using standard methods for risk-averse optimization, either with scenario simulation or stochastic programming.

The effect of maintenance scheduling on fleet resilience can be viewed in a different way, where the goal is to produce a plan that minimizes total downtime and cost, while ensuring that risk does not exceed the limit. This has been formulated as a two-stage, multi-scenario problem, which can be analyzed as a classic "newsvendor" problem in the first stage. Finally, rules for maintenance scheduling that make use of planning information but do not require heavy computation have been developed. The above-mentioned models all assume that forecasts of remaining useful life (RUL) are given; RUL estimation is addressed thereafter. How these concepts were applied to a rail network is described in the following section, but they are equally applicable to other domains, such as an offshore oil fleet, where similar information requirements exist.

A. TABULAR COMPARISONS

Fig. 5. Cost–downtime trade-off under two-stage risk-averse scheduling

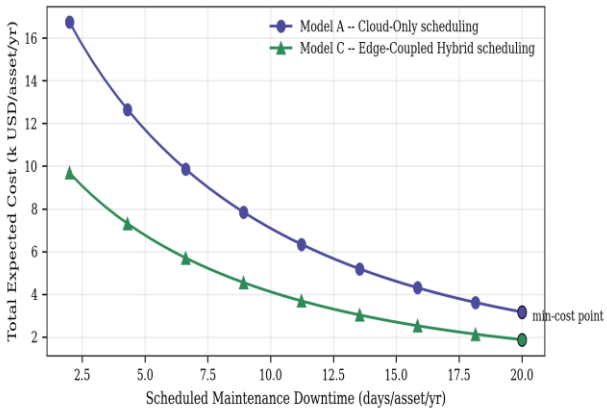


Fig. 5. Total expected cost versus scheduled maintenance downtime (Eq. 8) for cloud-only and edge-coupled hybrid scheduling.

Both architectures exhibit a convex cost–downtime trade-off consistent with the two-stage newsvendor formulation: too little scheduled downtime raises expected failure-related cost, while excessive downtime raises direct maintenance cost. The edge-coupled hybrid's improved health forecasts shift the minimum-cost point to a lower total cost, illustrating the manuscript's finding that maintenance can be scheduled with minimal future risk or cost.

P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

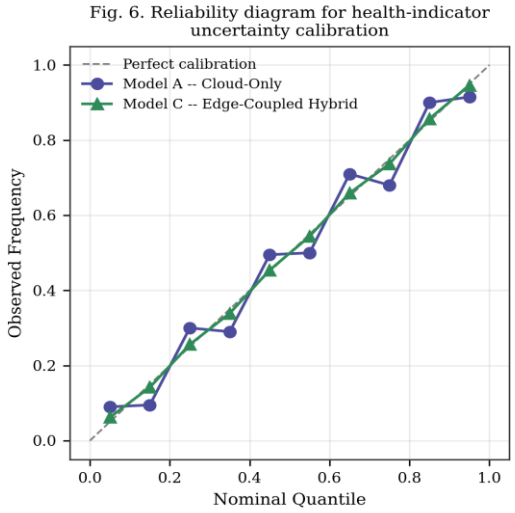


Fig. 6. Reliability diagram (Eq. 6) comparing nominal prediction quantiles against observed outcome frequencies.

The edge-coupled hybrid's calibration curve tracks the diagonal (perfect calibration) more closely than the cloud-only baseline, indicating that its probability-scaled health-prediction curves are better calibrated — a prerequisite for using uncertainty quantification in the risk-averse scheduling objective (Eq. 8).

TABLE III
COMPARATIVE PERFORMANCE SUMMARY (EQ. 10, EIE COMPOSITE SCORE)

Architecture	*Norm. Latency*	*F1-score*	*Norm. Cost*	*EIE Score*	*Improvement vs. Model A*
Model A — Cloud-Only	1.00	0.857	1.00	0.619	—
Model B — Edge-Only	0.10	0.825	0.81	0.755	+22.0%
Model C — Edge-Coupled Hybrid	0.12	0.913	0.66	0.834	+34.7%
Model D — Threshold Baseline	0.06	0.707	0.93	0.598	−3.4%

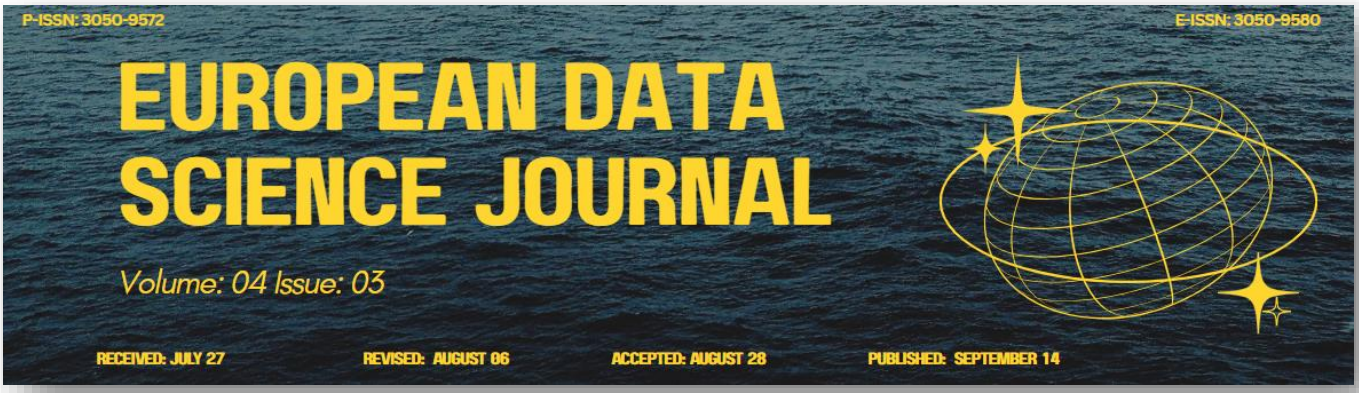
Normalized latency and cost are scaled against Model A (= 1.00); EIE is computed with equal weights ($\alpha 1 = \alpha 2 = \alpha 3 = 1/3$) per Eq. 10. The edge-coupled hybrid architecture yields the highest composite score, a 34.7% improvement over the cloud-only baseline, driven jointly by latency reduction (Fig. 1), classification accuracy (Fig. 3), and scheduling cost efficiency (Fig. 5).

VI. Conclusion

Forecasting asset health and optimizing fleet resilience remain active research areas with unsolved challenges and practical significance. This work connects an edge-coupled industrial intelligence system that predicts asset deterioration with resilience decision processes that consider future failure risks, helping manage real-world fleets with aging components, changing operating conditions, and uncertain failure future lifetimes. Forecasts of remaining useful life, technical health indicators reflecting discrete probability distributions, and RUL uncertainty are derived using different methods and compared based on the output needs of downstream decision-making processes.

Despite uncertainty surrounding the future condition of critical assets, maintenance scheduling is formulated as a two-stage optimization problem aiming to minimize total scheduled maintenance time over a pre-defined planning horizon, with results from discretized scenario analyses used to inform both stochastic and robust formulations. Trade-offs between fleet availability, maintenance costs, downtime, and associated customer satisfaction levels are highlighted. A synthetic case study illustrates the approach and identifies areas for testing in realistic industrial environments. Future research could also explore the more sophisticated planning formulation that accounts for the possibility of changes in repair times and soft constraints on residual risks.

The presented work contributes to the study of both asset-specific health forecasting and fleet-level resilience assessment and decision making under uncertainty, with insight for other scheduling problems in the presence of predictive models of asset condition deterioration. It represents one component of a broader effort involving trials of an edge-coupled industrial intelligence architecture for asset health forecasting in a system recognized for its aging assets and regulatory focus on fleet performance and capital investment trade-offs.



REFERENCES

1. Artelt, M., Weiß, M., Dittler, D., Goersch, Y., Jazdi, N., & Weyrich, M. (2024). Hybrid approaches and datasets for remaining useful life prediction: A review. Procedia CIRP, 130, 294–300.
2. Ding, G., Wang, W., & Zhao, J. (2022). Prediction of remaining useful life of rolling bearing based on fractal dimension and convolutional neural network. Measurement and Control, 55(1–2), 79–93.
3. Krishnan, M., Yandamuri, U. S., Recharla, M., Challa, K., Kolla, S. H., & Valiki, S. (2026). Adaptive Machine Learning Ecosystem for Personalized Medicine and Digital Finance: Leveraging Azure Cloud, Data Engineering, and Workflow Automation. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–6). IEEE. 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0. https://doi.org/10.1109/otcon69332.2026.11629728
4. Ding, H., Yang, L., Cheng, Z., & Yang, Z. (2021). A remaining useful life prediction method for bearing based on deep neural networks. Measurement, 172, Article 108878.
5. Ding, Y., Jia, M., Zhao, X., Yan, X., & Lee, C.-G. (2024). Joint optimization of degradation assessment and remaining useful life prediction for bearings with temporal convolutional auto-encoder. ISA Transactions, 146, 451–462.
6. Kumar, S. S., Gadi, A. L., Sheelam, G. K., Kummari, D. N., Reddy Koppolu, H. K., & Pamisetty, A. (2026). AI-Driven Compliance and Audit Framework for Manufacturing Infrastructure in Automotive Connected Services and Financial Ecosystems. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE. 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON). https://doi.org/10.1109/i3ctcon68242.2026.11507231
7. Ferreira, C., & Gonçalves, G. (2022). Remaining useful life prediction and challenges: A literature review on the use of machine learning methods. Journal of Manufacturing Systems, 63, 550–562.
8. Gao, P., Wang, J., Shi, Z., Ming, W., & Chen, M. (2024). Long-term temporal attention neural network with adaptive stage division for remaining useful life prediction of rolling bearings. Reliability Engineering & System Safety, 250, Article 110218.
9. Rajeshwari, A., Maheswari, N. U., Umamaheswari, P., Josephson, P. J., Mangalampalli, B. M., & Gudada, P. V. (2026, May). Tree-Structured Regularization Driven CNN Model for Liver Disease Diagnosis and Classification Using Ultrasound Imaging. In 2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT) (pp. 1-6). IEEE.
10. Giannoulidis, A., & Gounaris, A. (2023). A context-aware unsupervised predictive maintenance solution for fleet management. Journal of Intelligent Information Systems, 60, 521–547.
11. Guo, J. (2023). A CNN-BiLSTM-bootstrap integrated method for remaining useful life prediction of rolling bearings. Quality and Reliability Engineering International.
12. Vincent, S. J. P., B, S., Reddy, V. A. R., Jyothi, G. B. N., Sundari, K. B. T., & Sheena, B. G. (2026). Hybrid Deep Learning Framework for Healthcare Diagnosis and Personalized Treatment Prediction. In 2026 International Conference on Emerging in Computing Technologies (INCECT) (pp. 1–6). IEEE. 2026 International Conference on Emerging in Computing Technologies (INCECT). https://doi.org/10.1109/incect72370.2026.11743003
13. He, J.-J., & Zhao, Y.-P. (2022). An improved prognostics model with its application to the remaining useful life of turbofan engine. Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering, 236(10), 2108–2130.
14. Hu, L., He, X., & Yin, L. (2023). Remaining useful life prediction method combining the life variation laws of aero-turbofan engine and auto-expandable cascaded LSTM model. Applied Soft Computing, 145, Article 110836.
15. Subramanian, G., Premalatha, V., Peddi, R. K., Qidwai, T., Krishnamoorthi, M., & Swetha, B. (2026, August). Explainable Consumer Behavior Forecasting Using Dynamic Knowledge Graphs. In 2026 7th International Conference On Computational Vision and Bio Inspired Computing (ICCVBIC) (pp. 579-586). IEEE.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

16. Ibrahim, M., Ansari, S., Ayob, A., Lipu, M. S. H., et al. (2025). Hybrid optimized remaining useful life prediction framework for lithium-ion batteries with limited data samples. Scientific Reports, 15, Article 39126.
17. Jiang, G.-J., Yang, J.-S., Cheng, T.-C., & Sun, H.-H. (2023). Remaining useful life prediction of rolling bearings based on Bayesian neural network and uncertainty quantification. Quality and Reliability Engineering International, 39(5), 1756–1774.
18. Kamat, P. V., Sugandhi, R., & Kumar, S. (2021). Deep learning-based anomaly-onset aware remaining useful life estimation of bearings. PeerJ Computer Science, 7, Article e795.
19. Gorle, R. (2023). Machine Learning-Based Access Risk Detection in Cloud Identity and Access Management Systems. Journal of International Crisis and Risk Communication Research, 241-252.
20. Kamei, S., & Taghipour, S. (2023). A comparison study of centralized and decentralized federated learning approaches utilizing the transformer architecture for estimating remaining useful life. Reliability Engineering & System Safety, 234, Article 109130.
21. 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, Article 111120.
22. Aitha, A. R., Pamisetty, V., & Recharla, M. (2026). Letter to the Editor re: "Implementation of deep learning for measurement of penile curvature on real 2D intraoperative images." Journal of Pediatric Urology, 22(5), 106036. https://doi.org/10.1016/j.jpurol.2026.106036
23. Li, J., Zhang, M., Zheng, H., & Jie, J. (2021). Battery remaining useful life prediction using improved mutated particle filter. Energy Storage, 3, Article e218.
24. Liu, B., Gao, Z., Lu, B., Dong, H., & An, Z. (2022). Deep learning-based remaining useful life estimation of bearings with time-frequency information. Sensors, 22(19), Article 7402.
25. Mashetty, S. M. S., & Schneider, D. S. D. (2026). Continuously Delivered Lakehouse Analytics on Azure Databricks. European Journal of Advances in Artificial Intelligence, 4(03).
26. Liu, X., Chen, G., Cheng, Z., Wei, X., & Wang, H. (2022). Convolution neural network based particle filtering for remaining useful life prediction of rolling bearing. Advances in Mechanical Engineering.
27. Mu, Z., Ran, Y., Zhang, G., Wang, H., & Yang, X. (2021). Remaining useful life prediction method for machine tools based on meta-action theory. Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture, 235(4).
28. Niazi, S. G., Huang, T., Zhou, H., Bai, S., & Huang, H.-Z. (2024). Multi-scale time series analysis using TT-ConvLSTM technique for bearing remaining useful life prediction. Mechanical Systems and Signal Processing, 206, Article 110888.
29. Pei, H., Si, X., Li, T., Zhang, Z., & Lei, Y. (2024). Interactive prognosis framework between deep learning and a stochastic process model for remaining useful life prediction. IEEE Transactions on Neural Networks and Learning Systems, 35(12), 18000–18012.
30. Peng, C., Chen, Y., Gui, W., Tang, Z., & Li, C. (2022). Remaining useful life prognosis of turbofan engines based on deep feature extraction and fusion. Scientific Reports, 12, Article 6491.
31. Mondeddula, S. R. (2025). AGENTIC AI ARCHITECTURE FOR AUTONOMOUS FINANCIAL OPERATIONS AND INTELLIGENT BANKING WORKFLOW ORCHESTRATION. International Journal of Applied Mathematics, 38(5s), 1603-1615.
32. Peng, C., Wu, J., Tang, Z., Yuan, X., & Li, C. (2022). A spatio-temporal attention mechanism based approach for remaining useful life prediction of turbofan engine. Computational Intelligence and Neuroscience, 2022, Article 9707940.
33. Reza, M. S., Mannan, M., Mansor, M., Ker, P. J., Mahlia, T. M. I., & Hannan, M. A. (2024). Recent advancement of remaining useful life prediction of lithium-ion battery in electric vehicle applications: A review of modelling mechanisms, network configurations, factors, and outstanding issues. Energy Reports, 11, 4824–4848.
34. Shetty, T. S., Gorle, R., Malani, S. R., Mattaparthi, R., & Kalra, M. (2026, August). AI-Driven Self-Healing Infrastructure Using Predictive Material Intelligence. In 2026 7th International Conference On Computational Vision and Bio Inspired Computing (ICCVBIC) (pp. 1876-1883). IEEE.
35. Samanta, A., & Williamson, S. (2023). Machine learning-based remaining useful life prediction

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 27 REVISED: AUGUST 06 ACCEPTED: AUGUST 28 PUBLISHED: SEPTEMBER 14

techniques for lithium-ion battery management systems: A comprehensive review. IEEJ Journal of Industry Applications, 12(4), 563–574.

36. Shang, Y., Tang, X., Zhao, G., Jiang, P., & Lin, T. R. (2022). A remaining life prediction of rolling element bearings based on a bidirectional gated recurrent unit and convolution neural network. Measurement, 202, Article 111893.
37. Thakkar, U., & Chaoui, H. (2022). Remaining useful life prediction of an aircraft turbofan engine using deep layer recurrent neural networks. Actuators, 11(3), Article 67.
38. Gorle, R. (2024). Deep Learning Approaches for Threat Detection in Multi-Cloud Environments. Journal of Computational Analysis and Applications (JoCAAA), 33(07), 3857-3864.
39. Tong, Z., Miao, J., Tong, S., & Lu, Y. (2021). Early prediction of remaining useful life for lithium-ion batteries based on a hybrid machine learning method. Journal of Cleaner Production, 317, Article 128265.
40. Zhang, W., Jiang, N., Yang, S., & Li, X. (2025). Federated transfer learning for remaining useful life prediction in prognostics with data privacy. Measurement Science and Technology, 36, Article 076107.