Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira, Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria
 Published by Research Publishing, Singapore.
doi: 10.3850/ESREL2026061419_esrel26-p26344-cd

Scenario-Based Optimization of RUL Models for Prescriptive Maintenance in Medical Imaging Equipment

Xi Wang, Pankaz Das, Bulent Alpay

GE Healthcare Technologies, The United States of America.
E-mail: xi.wang1@gehealthcare.com, pankaz.das@gehealthcare.com, bulent.alpay@gehealthcare.com

Medical imaging systems rely on high-precision components that degrade over time due to thermal stress, mechanical wear, and electronic drift. Unanticipated failures can disrupt clinical workflows and incur significant operational costs. This study focuses on the optimization of machine learning (ML) based models for Remaining Useful Life (RUL) prediction, aimed at enhancing the effectiveness of prescriptive maintenance strategies in medical imaging devices. The developed framework adopts failure-specific predictive models (e.g. CT detector) that trigger targeted actions to prevent or resolve issues. Models are optimized to align with distinct service strategic goals under multiple operational scenarios, such as minimizing part replacement costs, ensuring service-level agreement (SLA) compliance, maximizing uptime, or improving maintenance efficiency by reducing service cost.

The approach leverages multivariate time-series data from historical system logs and service records. Feature engineering includes high-order statistical transformations to enhance signal quality. Feature selection is guided through statistical hypothesis test and domain expertise to isolate failure-relevant features. The labelled failure-specific datasets feed into a supervised ML pipeline to predict RUL classes for specific degradation modes. Model performance is evaluated using precision-recall metrics across multiple failure modes. We focus onto two scenarios: (1) balancing uptime and maintenance efficiency, where we balance confidence of the RUL predictions with coverage (e.g. recall > 0.75, precision > 0.75) to obtain both service productivity and asset owner satisfaction; (2) minimizing part replacement costs for high-cost components, where high-confidence RUL predictions (e.g. precision > 0.95) are essential to enable automated part ordering via integration with inventory management systems, ensuring timely availability and reducing unplanned downtime. These two scenarios integrated with specific, actionable recommendations for maintenance tasks demonstrate the feasibility of our approach implementation in healthcare prescriptive maintenance area.

Keywords: Predictive Maintenance, Prescriptive Maintenance, Machine Learning, Medical Imaging Devices, System Log, Remaining Useful Life, Degradation Mode.

1. Introduction

In modern industrial and engineering systems, predictive maintenance (PdM) has emerged as a key enabler for maintaining safe, reliable, and cost-effective operations by shifting maintenance from reactive or schedule-based actions to health-informed, time-appropriate interventions. Within this paradigm, machinery prognostics leverages operational and environmental measurements to detect early anomalies, track degradation, and estimate remaining useful life (RUL), providing actionable foresight that directly supports maintenance planning and resource allocation. Data-driven, machine-learning-based prognostic models can learn degradation behavior from historical and in-service condition-monitoring data—typically organized through a pipeline of data acquisition, health indicator construction, health-stage characterization, and RUL prediction thereby offering a practical route to RUL estimation when first-principles modeling is limited or system behavior is complex. The resulting PdM delivers tangible benefits, including reduced unplanned downtime, fewer unnecessary maintenance actions, improved availability and safety, and lower life-cycle cost, while also strengthening logistics and decision support through timely, uncertainty-aware health information Hu et al. (2022); Lei et al. (2018); Podofillini et al. (2015); Sikorska et al. (2011); Zio (2022); Pecht and Kang (2018); Manchadi et al. (2023).

While PdM estimates when a fault may occur (e.g., RUL), it often cannot determine what action to take, when to schedule it, or how to execute it under real constraints such as limited technicians, spare parts lead times, production plans, and safety/sustainability requirements. This motivates the shift to prescriptive maintenance (PsM), which extends PdM by converting prognostic insights into actionable, optimized recommendations—repair/replace decisions, maintenance timing, and resource orchestration—often spanning beyond the asset to include production planning and logistics. PsM can be viewed as a "predict-and-prescribe" decision pipeline: it couples health forecasting with optimization/decision logic and uses feedback to continuously improve recommendations. This prescriptive decision layer based on rules, mathematical programming, simulation/discrete-event or agent-based modeling, or increasingly reinforcement learning enable PsM to choose feasible actions that optimize cost, availability, and operational impact Giacotto et al. (2025); Ansari et al. (2019); Orošnjak et al. (2025); Liu et al. (2019); Choubey et al. (2019); Meissner et al. (2021); Menezes et al. (2019); Matyas et al. (2017); Lee et al. (2011).

Recent progress shows PsM moving from concept to structured architectures (e.g., layered model Ansari et al. (2019): data management $\rightarrow$ predictive analytics $\rightarrow$ recommender/decision support $\rightarrow$ semantic reasoning) and to holistic end-to-end frameworks that emphasize distributed decisioning, knowledge management, and integration with enterprise systems Choubey et al. (2019). Measurable value was reported when prescriptions are deployed: Matyas et al. (2017) show real manufacturing gains including 30% maintenance cost reduction, 20% time saving in maintenance activities, and 12% equipment availability improvement, enabled by correlating production, quality, and condition data and bundling recommended actions. In Ansari et al. (2019), a proof-of-concept model reports 12–25% downtime reduction and improved planned/unplanned downtime ratios by combining analytics with rule-based prescriptions.

Along with PsM transitioning from a cutting-edge, niche application to a rapidly growing, high-value strategy in industrial sectors, in many industrial settings, maintenance and operational decisions must be optimized under heterogeneous objectives that vary with asset criticality, cost exposure, contractual obligations, and regulatory risk. A single predictive model often cannot deliver the best trade-off across these contexts. To address this, we propose an optimization approach that trains a portfolio of AI/ML models, each explicitly tuned to a distinct performance objective aligned with a specific operational constraint or risk profile. For example, one model may prioritize high precision for high-value or high-cost components to reduce unnecessary interventions, while another may emphasize high recall in scenarios where missed detections carry substantial contractual or regulatory penalties. A third model may focus on minimizing predicted lead time for components governed by warranty terms or service-level requirements. At inference time, the platform dynamically selects (or prioritizes) the appropriate model(s) based on the decision context and applicable objective, enabling more robust and actionable recommendations than a one-size-fits-all strategy. In addition, our model development framework is designed specific to subsystems/components. This context-aware, multi-objective model selection improves decision-making quality across diverse business environments and supports consistent optimization under varying operational constraints.

2. Methodology

In this section, we will demonstrate how to quantify different business scenarios through cost-benefit analysis and how our ML framework is designed to achieve PsM through optimization of each step in the workflow.

2.1. *Cost analysis*

To quantify the business scenarios, cost-benefit analysis is essential to set criteria for ML model development and assess the effectiveness of maintenance strategies. As business scenarios vary across industrial settings, factors to be included in a cost-benefit function also vary in practice, although general guidance and standards has been explored Feldman et al. (2009); Sun et al. (2012);

Pecht and Kang (2018); Wang and Pecht (2011); Banks et al. (2009); Meissner et al. (2021).

In our implementation of PsM in medical imaging devices, we have considered the factors including material cost, repair time-cost, liquidation rate, unplanned downtime, asset owner revenue loss from unplanned downtime, incurred cost of misplacing before end of life (EoL), and misclassification cost with PsM. By predicting potential failures, the costs or benefits gained from repairing components based on PsM versus the cost of leaving the machine run to failure can be quantified for the service provider and asset owner separately. The costs/benefits with PsM are directly related to the performances of predictive models from effective intervention (true positive: TP), unnecessary intervention (false positive: FP), and missed failure (false negative: FN) predictions.

Assuming that the components are still under warranty and that the material cost is borne solely by the service provider, the benefit to the service provider and asset owner can be estimated as the cost reduction achieved by adopting PsM instead of reactive maintenance (RM):

benefit *= RM cost - PsM cost*

RM cost$_{\textit{service provider}}$ *= (material cost + liquidation rate × repair time) × number of failures*

PsM cost$_{\textit{service provider}}$ *= (material cost + liquidation rate × improved repair time) × (TP + FP + FN) + incurred cost of misplacing before EoL × FP + misclassification cost × FN*

RM cost$_{\textit{asset owner}}$ *= unplanned downtime × asset owner revenue loss rate × number of failures*

PsM cost$_{\textit{asset owner}}$ *= unplanned downtime × asset owner revenue loss rate × FN*

Here, misclassification cost denotes the penalty associated with failing to detect a failure, while liquidation rate refers to the repair-related cost, including labor and travel. Because FN $<=$ number of actual failures, asset owners typically benefit from predictive maintenance. In contrast, from the service provider's perspective, the same predictive maintenance strategy—particularly when embedded in a PsM offering—can become economically unfavorable when model performance fails to satisfy operational criteria under realistic cost structures. Figure 1 illustrates the relationship between PsM benefits (for both asset owners and service providers) and material cost, holding other cost components constant, across three model-performance scenarios: low recall/high precision (0.50, 0.95), high recall/low precision (0.95, 0.50), and balanced recall/precision (0.75, 0.75). As shown, asset owners realize positive benefit in all three scenarios, with greater benefit under higher recall. Service providers, however, benefit more from high precision, reflecting the higher financial exposure to unnecessary interventions. Moreover, service provider benefit decline as material cost increases and can become negative beyond a threshold for all three scenarios. Under the assumptions used here, net benefit are maximized with higher recall when material costs are low, whereas precision becomes increasingly critical as material costs rise. These results suggest that model-performance criteria must be explicitly defined and enforced to ensure that PsM delivers positive economic value across stakeholders.

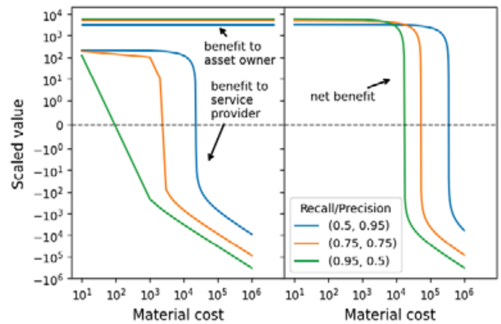

Fig. 1. Service provider benefit, asset owner benefit, and net benefit under PsM as a function of material cost for three model operating points: high-recall, high-precision, and balanced precision–recall.

2.2. *Model development framework*

We have developed a PsM framework for medical imaging systems that uses failure-specific, ML-based RUL models to predict component degradation (e.g., CT detector-related modes) and trigger targeted maintenance actions. Failure-specific

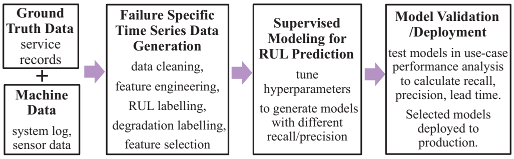

Fig. 2. Failure-specific RUL model development framework.

models are essential for PsM as the maintenance actions, cost analysis, and thus model success criteria vary for different component failures. The framework is demonstrated in Fig. 2, including data generation, failure-specific training data creation, supervised modeling for RUL prediction, and model validation and evaluation.

Customer relationship management (CRM) data contain service records (SRs) for imaging devices (reported faults, suspected failed components, timestamps, and corrective actions) and are often used as ground truth for model training. However, service-derived failure labels are noisy: interventions may reflect misdiagnosis, temporary workarounds, or overlapping symptoms across failure modes, and a single visit may include multiple actions (e.g., part replacement, calibration, software updates, reboots), obscuring the true resolution. Label fidelity is further affected by technician subjectivity, asset owner sensitivity to degradation, intermittent faults, and transient effects from operating or environmental conditions. To mitigate these issues, we infer degradation state using part-replacement evidence, an internal parts database, and domain-expert validation. For higher-confidence labels, we additionally restrict a subset of analyses to records with a single dominant corrective action, trading sample size for label precision.

Machine data consist of multivariate time-series signals derived from machine logs and/or device sensors. In this study, we primarily rely on log data because parametric sensor data are unavailable for many components, while we have developed models based on parametric data such as MR table where consistent sensor data are available. Machine log data consists of errors, warning, and operational data, and have been widely used for PdM, typically via supervised ML/DL when reliable ground-truth labels are available; however, label scarcity and label noise can substantially limit model performance Sipos et al. (2014); Huang et al. (2021). Because raw log entries are often free-text messages, we transform them into structured representations by aggregating the daily frequency of standardized message codes.

Failure-specific time series data generation. The structured multivariate time-series derived from system logs and SRs are integrated using the labeling logic summarized in Table 1 to assign a service action to each data point. Specifically, we label a time-series point as *actionable* if it occurs within a prescribed period (e.g. 21 days) prior to a recorded failure up to the failure date; all other retained points are labeled as *nominal*. Because the true transition from nominal to actionable can vary across failure events, we further reduce label noise by excluding potentially ambiguous regions, such as within 90 days prior to the start of the SR and the first 7 days following the corrective action to account for post-repair stabilization.

Table 1. Example of target class labelling.

Proposed lead time in days	Class
Before failure event date: 0 to 21	Actionable
Before failure event date: 90 to 180	Nominal
After failure is corrected: -60 to -8	Nominal

We also apply feature engineering to enhance the signal representation by deriving higher-order statistical features (e.g., mean, variance, and skewness) over rolling windows, capturing their temporal dynamics through cumulative sums and daily ratio changes, and subsequently applying principal component analysis (PCA) for further dimensionality reduction.

Failure-specific feature selection is essential for building component-targeted PsM models. System logs span a machine's major functionalities, resulting in a very high-dimensional feature space from nearly 3,000 unique log message types for each product line. To identify failure precursors for individual components, we employ a permutation-based hypothesis testing approach

Radivojac et al. (2004) that requires minimal domain-expert input. This method is broadly applicable across failure modes and consistently removes approximately 90–95% of non-informative features for diverse components (detector, collimator, etc.), substantially reducing dimensionality while preserving predictive signatures. With this feature selection technique, we have produced models beyond preset criteria for multiple components across CT product lines.

Supervised learning, classification or regression, learns failure-mode–specific degradation signatures from historical data and tests on unseen data. Hyperparameters and model parameters are tuned to optimize the output and also generate models with various performance profiles (recall, precision, lead time etc).

Model validation is performed using a use-case-driven evaluation framework that mirrors the service delivery setting. Using hold-out, we assess model performance in terms of recall, precision, and actionable lead time. Models that meet the target operating objectives are subsequently selected for deployment in production.

Optimization of decisions for PsM. The aforementioned model development framework is designed for training and deploying failure-mode-specific models. A portfolio of models configured with different performance objectives are aligned to varying operational constraints and risk profiles. Such objectives may include maximizing precision for high-value or high-cost items, maximizing recall for use cases associated with significant contractual or regulatory penalties, and minimizing predicted lead times for components under warranty or service-level obligations, as illustrated in Fig. 3. The selective use of models dynamically based on these heterogeneous performance criteria in an optimization platform enables improved decision-making across diverse business contexts.

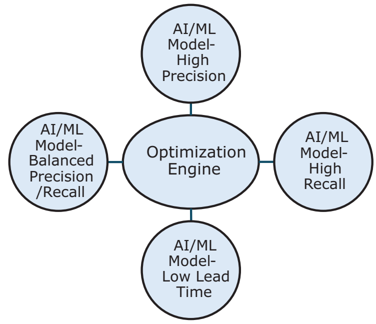

Fig. 3. Scenario-based optimization.

3. Case study & results

In this section, we will demonstrate the implementation of our ML framework to the CT detector models of a product line as an example for PsM in two scenarios, balancing uptime and maintenance efficiency using models with balanced recall and precision, minimizing part replacement costs for high-cost components using models with high-precision.

3.1. *Case study details*

This case study focuses on the GE HealthCare Revolution CT product line. Approximately 2,000 systems from all geographic regions are included, with about 500 SRs related to detector replacement. After feature selection, approximately 200 message codes are retained, from which roughly 5,000 features were derived. These features are then transformed to 200 principal components via principal component analysis (PCA), which are used as the final inputs for modeling. Data from 2020 to 2023 are used for model training, whereas unseen data from 2024 are used for testing.

3.2. *Assumptions*

There are a few assumptions along with the model development and implementation scenario analysis for CT detector, including

- A detector “failure event” is defined by a CRM-recorded detector module replacement (validated via parts database / expert review).
- Logs are standardized to message codes; daily aggregation of code frequency is sufficient to capture degradation signals.
- Actionable window = 0-21 days and nominal window = 90-180 days pre-failure (baseline), and other days within 90 days prior to the actionable window and 7 days after repair are excluded as ambiguous.

- Feature selection via permutation-based hypothesis testing yields generalizable precursors.
- Warranty assumption: service provider bears material cost (at least for the modeled regime).
- Proxy part costs (e.g., 1,000 vs. 100,000) represent low vs. high cost components, respectively.
- Precision/recall thresholds define scenario success (e.g., balanced > 0.75; high precision > 0.95).
- Precision–recall metrics sufficiently characterize deployment performance.

3.3. *Model optimization for various objectives*

To develop the detector-failure-specific RUL models, we extract ground truth information from SRs for detector module replacement events, select detector-failure-related features through statistical hypothesis testing complemented by domain expert judgment, train a Random Forest classifier, and tune key hyperparameters (e.g., tree complexity and class weighting) to control the model performance profiles. Representative models at different operating points including high recall, high precision, and balanced configurations are shown in Table 2.

Table 2. A portfolio of models generated by controlling hyperparameters and model parameters.

Hyperparameters	Recall	Precision
Set 1	0.42	0.97
Set 2	0.60	0.86
Set 3	0.75	0.75
Set 4	0.80	0.71
Set 5	0.85	0.68

The pool of models provides the capability for dynamic selection of models according to different business scenarios.

3.4. *Scenario based model selection*

In this section, we demonstrate scenario-based model selection using the detector RUL classifiers and their precision-recall operating points. We consider two representative business scenarios: (1) balancing uptime and maintenance efficiency,

Table 3. Cost analysis for scenario-based model selection. Values are scaled.

Material Cost	Recall, Precision	Benefit to Service Provider	Benefit to Asset Owner	Net Benefit
1e3	(0.42, 0.97)	60	750	810
1e3	(0.75, 0.75)	30	1340	1370
1e3	(0.85, 0.68)	10	1520	1530
1e5	(0.42, 0.97)	-60	750	690
1e5	(0.75, 0.75)	-2580	1340	-1240
1e5	(0.85, 0.68)	-4140	1520	-2620

where the objective is to achieve adequate coverage of impending failures (recall) while maintaining sufficient confidence in alerts (precision), thereby improving service productivity and asset owner satisfaction; and (2) minimizing part replacement costs for high-cost components, where the objective is to suppress FPs and ensure that maintenance actions are triggered only when the likelihood of true degradation is high. Here we use 1,000 and 100,000 as proxies for low- and high-cost components, respectively.

As shown in Table 3, under the low-cost scenario, models with balanced recall and precision (e.g., (0.75, 0.75) or (0.85, 0.68)) deliver greater benefit to the asset owner and comparable—though slightly reduced—benefit to the service provider relative to a high-precision, low-recall model. This outcome is consistent with the cost formulation in Section 2.1: asset owner benefits are dominated by reductions in unplanned downtime, which improves primarily as recall increases (i.e., fewer FNs). However, excessively prioritizing recall at the expense of precision (e.g., recall > 0.90 with substantially lower precision) can shift the service provider's economics from benefit to loss, because additional FPs drive unnecessary proactive interventions and associated logistics, labor, and replacement costs. Consequently, for low-cost parts, operating points that maintain a reasonable precision floor while achieving meaningful recall provide the best joint outcome by simultaneously improving uptime and maintenance efficiency.

For high-cost components, the optimal operating point shifts toward very high precision (e.g.,

> 0.95). In this regime, the marginal cost of a FP is large because it can trigger premature replacement of an expensive component and additional downstream costs (e.g., shipping, handling, and potential inventory holding). While higher recall remains desirable to avoid missed failures, the economic cost of unnecessary replacements and cost of replacing before EoL dominate as material cost increases; therefore, precision becomes the primary gating criterion for economic viability. High-confidence RUL predictions also enable reliable automation of part-ordering workflows through integration with inventory management systems, supporting just-in-time availability while reducing excess inventory exposure and mitigating unplanned downtime.

Overall, these results show that a single "best" model does not exist independent of context; instead, the proposed framework enables selection from a portfolio of models tuned to distinct objectives and constraints. Beyond the two canonical scenarios illustrated here, the same cost-driven selection logic can be applied to customized, emergent operational conditions and to planning horizons that incorporate both scheduled activities and asset owner-initiated service events.

4. Conclusion

In this work, we presented a scenario-based optimization framework that integrates failure-specific RUL prediction models with cost-driven PsM decision-making for medical imaging equipment. By combining multivariate time series data, supervised label refinement, and hypothesis test based feature selection, the framework generates a portfolio of ML models explicitly configured to meet heterogeneous operational objectives.

Through a detailed optimization results analysis, we demonstrated that the optimal model choice depends strongly on component material cost and stakeholder impact. For low-cost components, models with balanced recall and precision deliver the greatest combined value by improving equipment uptime and maintenance efficiency, benefitting both asset owners and service providers. In contrast, for high-cost components, high-precision models are essential to avoid costly unnecessary part replacements and to enable automated, just-in-time inventory workflows that minimize unplanned downtime.

The proposed framework supports dynamic selection among models optimized for different objectives—such as high recall, or high precision—allowing maintenance recommendations to be tailored to each operational scenario. This flexibility enables consistent value generation across diverse business conditions and improves the reliability and economic performance of prescriptive maintenance strategies.

Future extensions may incorporate reinforcement learning-based decision policies, richer representations of operational uncertainty, and broader integration with enterprise logistics and scheduling systems. Overall, the results confirm that scenario-aligned model optimization is a practical and effective approach for deploying PsM in complex, high-value healthcare assets.

References

Ansari, F., R. Glawar, and T. Nemeth (2019, May). PriMa: a prescriptive maintenance model for cyber-physical production systems. *International Journal of Computer Integrated Manufacturing 32*(4-5), 482–503.

Banks, J., K. Reichard, E. Crow, and K. Nickell (2009, March). How engineers can conduct cost-benefit analysis for PHM systems. *IEEE Aerospace and Electronic Systems Magazine 24*(3), 22–30.

Choubey, S., R. Benton, and T. Johnsten (2019, December). Prescriptive Equipment Maintenance: A Framework. In *2019 IEEE International Conference on Big Data (Big Data)*, Los Angeles, CA, USA, pp. 4366–4374. IEEE.

Feldman, K., T. Jazouli, and P. Sandborn (2009, June). A Methodology for Determining the Return on Investment Associated With Prognostics and Health Management. *IEEE Transactions on Reliability 58*(2), 305–316.

Giacotto, A., H. C. Marques, and A. Martinetti (2025, March). Prescriptive maintenance: a comprehensive review of current research and future directions. *Journal of Quality in Maintenance Engineering 31*(1), 129–173.

Hu, Y., X. Miao, Y. Si, E. Pan, and E. Zio (2022, January). Prognostics and health management: A review from the perspectives of design, development and decision. *Reliability Engineering & System Safety 217*, 108063.

Huang, C., A. Deep, S. Zhou, and D. Veeramani (2021, July). A deep learning approach for predicting critical events using event logs. *Quality and Reliability Engineering International 37*(5), 2214–2234.

Lee, J., M. Ghaffari, and S. Elmeligy (2011, April). Self-maintenance and engineering immune systems: Towards smarter machines and manufacturing systems. *Annual Reviews in Control 35*(1), 111–122.

Lei, Y., N. Li, L. Guo, N. Li, T. Yan, and J. Lin (2018, May). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing 104*, 799–834.

Liu, B., J. Lin, L. Zhang, and U. Kumar (2019). A Dynamic Prescriptive Maintenance Model Considering System Aging and Degradation. *IEEE Access 7*, 94931–94943.

Manchadi, O., F.-E. Ben-Bouazza, and B. Jioudi (2023). Predictive Maintenance in Healthcare System: A Survey. *IEEE Access 11*, 61313–61330.

Matyas, K., T. Nemeth, K. Kovacs, and R. Glawar (2017). A procedural approach for realizing prescriptive maintenance planning in manufacturing industries. *CIRP Annals 66*(1), 461–464.

Meissner, R., H. Meyer, and K. Wicke (2021, March). Concept and Economic Evaluation of Prescriptive Maintenance Strategies for an Automated Condition Monitoring System. *International Journal of Prognostics and Health Management 12*(3).

Meissner, R., A. Rahn, and K. Wicke (2021, October). Developing prescriptive maintenance strategies in the aviation industry based on a discrete-event simulation framework for post-prognostics decision making. *Reliability Engineering & System Safety 214*, 107812.

Menezes, B. C., J. D. Kelly, A. G. Leal, and G. C. Le Roux (2019). Predictive, Prescriptive and Detective Analytics for Smart Manufacturing in the Information Age. *IFAC-PapersOnLine 52*(1), 568–573.

Orošnjak, M., F. Saretzky, and S. Kedziora (2025, July). Prescriptive Maintenance: A Systematic Literature Review and Exploratory Meta-Synthesis. *Applied Sciences 15*(15), 8507.

Pecht, M. and M. Kang (2018). *Prognostics and health management of electronics: fundamentals, machine learning, and internet of things* (2nd ed ed.). Hoboken: Wiley-IEEE press.

Podofillini, L., B. Sudret, B. Stojadinovic, E. Zio, and W. Kröger (Eds.) (2015, September). *Safety and Reliability of Complex Engineered Systems: ESREL 2015*. CRC Press.

Radivojac, P., Z. Obradovic, A. K. Dunker, and S. Vucetic (2004). Feature Selection Filters Based on the Permutation Test. In *Machine Learning: ECML 2004*, Volume 3201, pp. 334–346. Springer Berlin Heidelberg.

Sikorska, J., M. Hodkiewicz, and L. Ma (2011, July). Prognostic modelling options for remaining useful life estimation by industry. *Mechanical Systems and Signal Processing 25*(5), 1803–1836.

Sipos, R., D. Fradkin, F. Moerchen, and Z. Wang (2014, August). Log-based predictive maintenance. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, New York New York USA, pp. 1867–1876. ACM.

Sun, B., S. Zeng, R. Kang, and M. G. Pecht (2012, June). Benefits and Challenges of System Prognostics. *IEEE Transactions on Reliability 61*(2), 323–335.

Wang, W. and M. Pecht (2011, July). Economic Analysis of Canary-Based Prognostics and Health Management. *IEEE Transactions on Industrial Electronics 58*(7), 3077–3089.

Zio, E. (2022, February). Prognostics and Health Management (PHM): Where are we and where do we (need to) go in theory and practice. *Reliability Engineering & System Safety 218*, 108119.