Proceedings of the 35th European Safety and Reliability & the 33rd Society for Risk Analysis Europe Conference
Edited by Eirik Bjorheim Abrahamsen, Terje Aven, Frederic Bouder, Roger Flage, Marja Ylönen

doi: 10.3850/978-981-94-3281-3_ESREL-SRA-E2025-P8303-cd

Practical Application of Predictive Maintenance - Solving the Challenges

Tiedo Tinga

Dynamics based Maintenance, University of Twente & Faculty of Military Sciences, Netherlands Defence Academy, The Netherlands. E-mail: t.tinga@utwente.nl

Industry is eager to harvest the potential of predictive maintenance (PdM). At the same time, only a fraction of the methods proposed in the past decade is actually applied in practice. This paper will identify the most prominent barriers for practical application of PdM, which are mainly related to the quality, relevance and availability of data. After that, three solution directions will be presented. The first solution aims to properly match ambition level with the available data and knowledge. The second set of solutions circumvents the lack of data by using physical models in addition to data. The third set of solutions addresses alternative ways of collecting data, including experimental test benches, experiments on fielded systems, and using numerical (simulation) models to generate data. Finally, the required standardized registration of (failure) data will be addressed.

Keywords: Predictive Maintenance, data quality, diagnostics, prognostics, physics-of-failure, hybrid methods.

1. Introduction

With increasing demands for availability and reliability of critical assets, industry has a strong need for advanced maintenance concepts. In the present era of Industry 4.0, where sensor data is abundant, there is a lot of potential in data-driven maintenance policies. The benefits can either be obtained in better diagnostics or in improved prognostics. Better diagnostics allows to quickly or even automatically detect or diagnose failures in (complex) systems (Rijsdijk et al. 2024). Improved prognostics potentially provides even larger benefits, allowing to predict future failures in a timely manner. The associated maintenance policy is called Predictive Maintenance (PdM), aiming to plan maintenance tasks just-in-time, thus effectively preventing failures as well as over-maintenance. It also allows efficient preparation of required man power and spare parts. Therefore, it can be stated that industry is eager to harvest the potential of PdM.

At the same time, many methods and models have been proposed in the past decade, see e.g. the reviews on AI-based methods (Khan and Yairi 2018), system-level prognostics (Tamssaouet et al. 2022), PdM application (Zonta et al. 2020) and maintenance optimization (de Jonge and Scarf 2020; Pinciroli et al. 2023). However, practical application of these prognostic methods is still rather limited (Grubic et al. 2011; Tiddens et al. 2022; Akkermans et al. 2024). Practitioners apparently encounter barriers that prevent them to apply the available methods in practice. The present paper will address this gap between what is theoretically possible and practically feasible. In this, a case study and lessons-learned approach is followed, rather than solid theory-building research design.

The paper is organized as follows. In section 2 the various ambition levels in smart maintenance are introduced. Section 3 then identifies the two prominent barriers for practical application of PdM. Section 4, 5, and 6 will discuss how these two barriers can be tackled, supported by case studies and examples. Finally, section 7 contains the conclusions of the paper.

2. Ambition Levels in Smart Maintenance

It should be realized that organizations can have very different goals and objectives with their maintenance policies. For inexpensive systems in a commercial environment, low-cost maintenance will have highest priority, and a tradetional scheduled maintenance policy might be optimal. But for highly complex critical assets, the focus may be on just-in-time maintenance that ensures both effective prevention of failures (availability) and efficient replacement without spoiling remaining life time (cost benefits). These organizational objectives can be classified in four different ambition levels for the associated smart maintenance policies:

(i) Detection of failures: *is something wrong?*
(ii) Diagnosing failures: *what is wrong?*
(iii) Health assessment: *how wrong is it?*
(iv) Prognosis: *when is it expected to go wrong?*

Each of these ambition levels requires different types and amounts of data when developing or training the associated models. In practice, data availability might be limited, implying that certain ambition levels are infeasible.

The lowest level focuses on the automatic *detection* of failures (e.g. from sensor readings), and is typically based on an anomaly detection algorithm. This level has the lowest data requirements, as it only needs unlabelled data, i.e. without any information on whether the data represents a healthy or a failed system. The algorithm will detect any deviation from the common behaviour of the system. In practice, this merely requires a limited amount of sensor data (e.g. temperature, vibration level) to learn the normal behaviour, after which the algorithm will detect any deviation as an anomaly.

At the next ambition level of *diagnosis*, the aim is to not only detect failures, but also identify which of several potential failure modes has actually happened. This is achieved by application of a classification algorithm. Training of such an algorithm is denoted *supervised learning*, as it requires labelled data as input: it must be specified to which failure mode (i.e. class) a subset of the (training) data belongs. In practice this means that just collecting sensor data is not enough: from only healthy system data the algorithm will not learn how any of the failure modes is represented in the data. Therefore, sufficient traces of data representative for each failure mode must be collected. This firstly requires these failures to occur in practice, and secondly requires proper registration and reporting of the observed failures.

The next higher ambition level is the *health assessment*, focusing on determining the severity of the occurring degradation. This allows the comparison of the current health to a predetermined threshold value, leading to a condition-based maintenance policy. The data requirements are again more demanding than for the previous ambition levels. Now there is a need for measurements of the health (or condition) of the system. As explained in Keizers et al. (2024), these can be either direct measures of the system degradation, like crack length or corrosion depth, or indirect measures like vibration level or temperature rise. Direct measures allow immediate maintenance decisions, but indirect measures, common for many condition monitoring techniques, also need a relation between these measures and the actual degradation to quantify the threshold (at which maintenance is required). In practical situations this kind of health data can only be retrieved when dedicated sensors for condition monitoring are deployed.

The highest ambition level is the *prognosis*, which aims to predict when a future failure is expected. This is denoted as calculating the remaining useful life (RUL) of the system. Such predictions are either based on revealing the trajectories in time series of (sensor) data preceding the system failures, or on finding the quantitative relation between some input parameter (e.g. rpm, environment) and the resulting degradation rate, e.g. by regression. Note that the *reliability engineering* approach, based on deriving mean time to failure and failure rates for large populations of parts, is not considered a prognostic approach here, as it is unable to accurately predict the time to failure for an individual part under specific conditions.

Prognostic models must be trained from a large number of time series of sensor data. On top of that, these time series should represent *run-to-failure* trajectories, i.e. the data set must include the actual failure of the system. Especially the latter requirement makes derivation of prognostic algorithms in a practical setting very challenging: maintenance of critical systems is intended to prevent failures, so the presence of run-to-failure data is *by definition* always (very) limited.

Based on this overview of smart maintenance ambition levels, the main barriers will be discussed in the next section.

3. Barriers for the Application of PdM

Based on lessons-learned from research projects with industrial partners, the author has identified the main barriers for practical implementation of PdM. These barriers are all associated to the amount and quality of data that are available in the considered situation, and therefore also to the mentioned ambition levels.

The first barrier, as identified by Tiddens et al. (2022), appeared to be a (unknown) mismatch between the ambition level of an organization and the available (amount and quality of) data and knowledge. It was discovered that organizations are often not aware of the ambition level they have, but especially lack insight in what that ambition level implies for their data and knowledge requirements. Without this insight, companies start developing PdM methods for their systems, but often get stuck due to a lack of (relevant) data. Section 4 therefore proposes a decision support tool assisting in matching the ambition level with data availability. This allows selecting the right approach from the beginning, which significantly increases the success rate.

The second barrier that was identified is the actual lack of *useful* and *relevant* data in industrial practice. With the strong focus on data-driven approaches in PdM, the requirements for the amount and quality of data are rather demanding. Time series of sensor data are nowadays readily available, but, as discussed in section 2, without additional information these are only suitable for the lowest ambition level, i.e. anomaly detection. When aiming for a higher ambition level, i.e. diagnosis, health assessment or prognosis, the following limitations in data availability are encountered:

No labelling: although many companies do register faults and failures in their Computerized Maintenance Management Systems (CMMS), a precise and consistent description of the type of failure is often missing. This is typically the responsibility of the person operating the system, who is lacking the knowledge to pinpoint specific failures. The consequence is that the proper labels for the associated sensor data of the machine are missing, and training of classification algorithms becomes very hard;

No condition measurements: health assessment and prognostic algorithms can only be derived when a considerable set of degradation monitoring data is available. However, this requires the presence of dedicated sensors for condition monitoring, which is still not very common in industrial practice. The majority of sensors present in systems serve the purpose of monitoring (for safety reasons) and control, e.g. SCADA systems. These sensors could be used in some cases for smart maintenance purposes, but most of the time provide irrelevant data, especially if labels are lacking (see previous point).

No threshold value: when degradation monitoring data is available, taking maintenance decisions also requires to have a threshold value for the monitored quantity. In case of direct condition measurements (e.g. corrosion depth), this threshold can typically be derived from functional or structural integrity criteria. However, most of the commonly used condition monitoring techniques, like vibration analysis and oil analysis, are based on indirect measurements. In those cases either an experience-based (trial-and-error) threshold value must be used, or a quantitative diagnostic relation must be derived (Keizers et al. 2025b).

No run to failure data: for the prediction of failures, the availability of training data containing actual failures is essential. The patterns in sensor data associated to (or preceding) failures can only be discovered when sufficient examples are available. However, as was mentioned before, the strong focus of maintenance on preventing failures yields by definition a very limited amount of failure data. Especially for critical assets, like aircraft, industrial plants and military systems, current maintenance policies are conservative, ensuring to effectively reduce the number of actual failures to a minimum. This is very challenging for the development of prognostic methods, and also the main reason that ~ 95% of all methods proposed in academic literature have been developed and tested with artificial benchmark datasets, e.g. (de Pater et al. 2022). An example is the NASA CMAPSS dataset generated by a simulation code of an aircraft engine (Saxena et al. 2008). As was mentioned before, failure rates, that can be found in reliability handbooks or databases like NPRD/EPRD or OREDA (Anonym. 2025), do not solve this problem either, as they cover complete populations of systems rather than individual parts.

No operating history: after the development of a prognostic algorithm, it should be tested and validated to check its accuracy. This again requires a number of actual failures to be present (see previous point), but in addition also requires knowledge on the operational history of the considered system. As the time to failure of a system typically depends on the operating conditions (e.g. environment, speed, power setting), these factors will also be included in the algorithm. Validation of the method then requires that the conditions associated to a specific failure are fully known. The latter is very challenging in an industrial environment. While the proper registration of failures (see previous issues on run-to-failure and labelling) is already challenging, getting access to the full operational history of a system (which may cover a period of several years) is almost impossible in most organizations.

After identifying these barriers for practical application of PdM, the next three sections will discuss some directions to tackle these barriers:

(i) Finding the most suitable approach given the (limited) data;
(ii) Circumventing the limited amount of data by including physical models;
(iii) Extending the amount of (relevant) data;

4. Solution 1 - Matching Ambition and Data

The first solution tackles the first barrier, and aims to support practitioners in finding a proper match between ambition level and available data and knowledge. Whereas many companies directly aim for Predictive Maintenance, section 2 revealed that this is the highest and most challenging level in smart maintenance. It is therefore advisable to first consider the lower levels of detection and diagnosis, as these are much easier to attain. When after this check PdM is (still) the required ambition, then several maintenance techniques are available in literature, each with their benefits and (data) requirements.

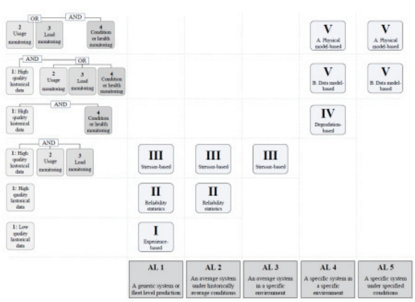

Fig. 1. Relating Maintenance Techniques to ambition level and data requirements (Tiddens et al. 2023).

To support the process of selecting the right technique, Tiddens et al. (2023) proposed a framework (Fig. 1) classifying these maintenance techniques (MT) in five categories, and linking them to both an ambition level (AL) and a set of data – knowledge requirements. The ambition levels deviate slightly from those in section 2, as they only focus on prognostics. The benefit of this framework is that a company can determine its ambition level, and find the associated MT with is data requirements. If there is a mismatch between AL and data, then either the AL should be lowered, or collection of the required (but still missing) data must be organized. Being aware of such a mismatch beforehand prevents a lot of useless efforts.

Answers		**EB**	**RS**	**CE**	**SB**	**DA**	**PB**
Is the selected approach required to provide lifetime prediction…	*Answer*	**Result score**					
(A1) … for individual systems? (No for fleet average values).	Yes	0	0	2	2	2	2
(A2) … for generalized cases?	No	2	2	1	1	1	1
(A3) … for varying operational conditions?	Yes	0	0	1	2	2	2
(A4) … for operational conditions that were not observed previously?	No	2	2	2	2	2	2
(B1) … with only limited accuracy (i.e. a rough estimate)?	Yes	2	2	0	1	0	0
(B2) … including insight into which parameters play an important role in the prediction?	Yes	0	0	0	2	2	2
(C1) Is an expert with practical experience available?	No	Ø	0	0	0	0	0
(C2) Is a specialist/ analyst available?	No	0	0	0	0	0	Ø
(C3) Is a data scientist available?	Yes	0	1	1	1	2	0
(D1) Are the (physical) degradation mechanism and the parameter(s) describing its evolution known ?	Yes	0	0	2	0	0	2
(D2) Is a stressor-base model available?	Yes	0	0	0	2	0	0
(D3) Is a physical model available?	No	1	1	1	0	1	Ø
(E1) Is run-to-failure sensor data available?	Yes	1	2	1	2	1	1
(E2) Is load/usage data available? Such data should cover the entire analyzed system.	Yes	0	0	1	1	2	1
(E3) Is historical failure data available?	Yes	1	1	2	2	2	2
(E4) Is a threshold value available?	No	0	0	Ø	0	0	0
(E5) Is real-time or periodic condition monitoring information available?	No	0	0	Ø	0	0	0
(E6) Are systems monitored individually (usage, loads, situation)?	Yes	0	0	2	0	2	2
(E7) Future operational conditions similar to past?	Yes	2	2	2	1	2	2
Total (T_S)		Ø	0.62	Ø	0.76	0.81	Ø
Suitability (S_S)		0.55	0.55	0.60	1.00	0.82	0.82
Feasibility (F_S)		Ø	0.70	Ø	0.60	0.80	Ø

Fig. 2. Calculation of the Suitability and Feasibility scores of various MTs (Alves da Silveira et al. 2023).

Recently, this framework has been extended to a more quantitative tool (Alves da Silveira et al. 2023). The tool (Fig. 2) quantifies the *suitability* and *feasibility* of a certain MT, based on answers (and predefined scores) provided by the user. The result is a total score for each MT, with the highest scoring method being the advised MT.

The included MTs are Experience-based (EB), Reliability statistics (RS), Condition extrapolation (CE), Stressor-based (SB), Data analytics (DA) and Physics-based (PB).

5. Solution 2 - Physics of Failure and Hybrid

The second set of solutions addresses the barrier of limited failure data and lack of labelling. Fully data-driven methods, like machine learning, require a lot of data to learn the patterns and relations that are present in the data. This is often not possible in industrial practice. However, more detailed knowledge of the considered system and its behavior can assist in deriving these relations with much less data. Physical models, based on the laws of nature, already contain the fundamental relations, e.g. between applied load and resulting degradation rate or time to failure. The only remaining challenge is then to find the values of the (material) parameters in these models, which might be slightly different for specific systems.

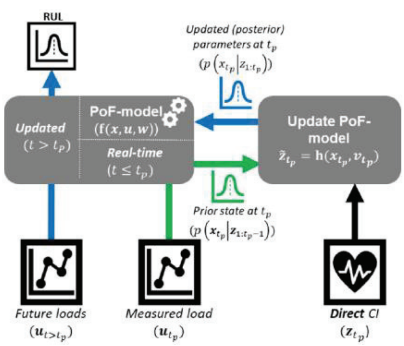

Fig. 3. Hybrid prognostic method combining a physical model with a particle filter (Keizers et al. 2025a).

This approach is followed in (Keizers et al. 2025a), where a hybrid method is developed for the prediction of corrosion damage (weight loss), combining a physical model with collected data. The adopted process is shown in Fig. 3. A physics-of-failure (PoF) is used to predict the RUL, based on the future loads (in this case the ambient temperature and humidity). However, the parameters in this PoF model are not precisely known for the specific system considered, so a Particle Filter (PF) is used to tune and update the model (parameters) with periodic measurements on the system condition (=condition indicator: CI). The limited amount of condition data would in itself not be sufficient to properly train a prognostic algorithm, especially under varying operating conditions. But combined with a physical model in a hybrid method called the Load-controlled Particle Filter (LCPF), it performs well.

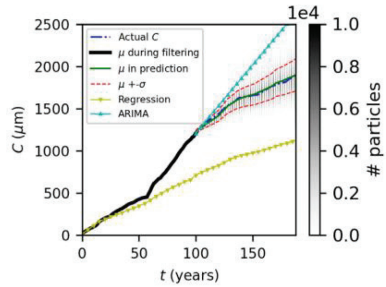

Fig. 4. Corrosion depth predictions for varying operating conditions with the LCPF compared to regression and ARIMA models (Keizers et al. 2025a).

Fig. 4 shows the predictions of the LCPF method starting from t = 100 (after tuning the model with measurements up till that moment). The results are compared to traditional regression and ARIMA (moving average) methods, revealing the much better performance of the hybrid LCPF method for this varying loads situation.

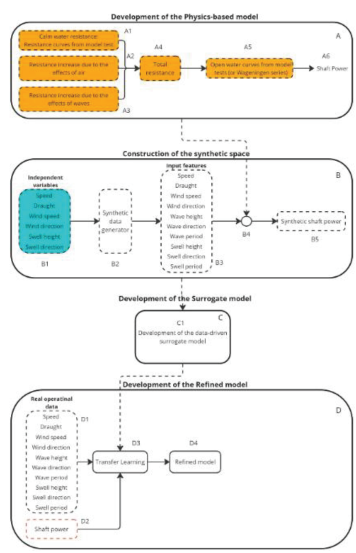

Fig. 5. Representation of Transfer Learning framework (Mavroudis and Tinga 2025).

Another hybrid approach (Fig. 5.), focusing on predicting the hull resistance of a ship operating in various sea conditions, is proposed in (Mavroudis and Tinga 2025), This allows to detect deviations that could indicate hull fouling, which can be used to trigger maintenance tasks. As the amount of data collected is insufficient to train a fully data-driven method, a physical model is constructed to calculate the resistance from first principles. Then Transfer Learning is used to tune the model, using the small amount of real data, to a specific ship.

6. Solution 3 – Generation of Failure Data

The third set of solutions addresses alternative ways of collecting data when field data is too limited or too low quality. Two solution directions exist (Fig. 6): 1) disclosing the failure path that remains covered when parts are replaced preventively; 2) assess the condition of the replaced part. Four approaches to achieve this are explored: executing lab-scale experiments on test benches, executing real-life experiments in fielded systems, improve labeling of field data by experts and using numerical simulations. The first two follow option 1 in Fig. 6, the final two option 2, as will be explained below. The required standardized registration of (failure) data will be addressed in 6.5.

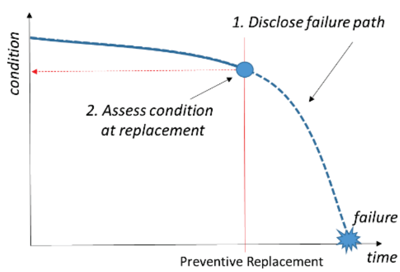

Fig. 6. Two solution directions for data improvement.

6.1. *Experimental test benches*

As was discussed before, the number of failures encountered in normal operation is typically very limited due to the executed PM tasks. An additional complication is that the types of failures and associated operational history cannot be controlled for fielded systems. The occurrence of failures fully depends on how the operator uses the system. This implies that a representative sample of failure types that covers the full operational range of the system will hardly ever be obtained from field data within reasonable time. To tackle this issue, fully controlled lab-scale experiments with real (sub)systems can be executed to generate the required data (option 1 in Fig. 6). The big advantage of such experiments is the full control over failure types and operating conditions. Additionally, the lab environment guarantees rather noise-free data, as no other systems are operating close-by.

Fig. 7. Experimental set-up of e-motor driven centrifugal pumps equipped with multiple sensors.

Bruinsma et al. (2024) used the centrifugal pump set-up shown in Fig. 7 to generate data for a large range of faults. Vibration sensors were used to collect data for both a healthy pump, and for various bearing faults, misalignment, unbalance, cavitation, etc. In contrast to field data, this data is perfectly labelled, as faults and operating conditions could be accurately controlled. This NLN-EMP dataset has been published (Bruinsma et al. 2024) to assist others in developing detection and diagnostic algorithms. The dataset is currently used to develop an automated failure identification method, that wil be applied in future naval ships.

6.2. *Real-life experiments on fielded systems*

An alternative way of generating failure data is to use fielded systems as experimental set-ups. Normally this would yield limited failure data, as the PM tasks typically prevent failures. However, in some situations it is possible to postpone or even skip PM tasks. As a result, the components will be used for their full service life (until failure) and the actual failure behaviour, is disclosed (solution 1 in Fig. 6). This will accelerate the data failure data collection, sup-porting decisions on extending PM intervals (which now may appear to be too conservative). Secondly, if sensors are deployed, the patterns in the sensor data associated to the actual failures can be obtained. Based on this insight, future failures can be timely detected or even predicted. Such an experiment cannot be applied to any system, as for critical systems the consequences of failure might be too severe. However, for non-critical or redundant sub-systems, this approach might well be feasible. Measures should then be taken to minimize the failure effects, e.g. by storing spare parts nearby. In addition, executing this experiment with a fraction of a larger fleet of parts would be a good option. If the systems with the highest age (or operating hours) are selected, denoted the '*front runners*', insights are obtained before the rest of the fleet reaches that age, which can be utilized (for interval extension) for the rest of the fleet.

In the Dutch MoD this experiment is now executed with the fuel injectors on a fleet of diesel engines. Each engine contains multiple injectors, whose failure therefore does not lead to a non-functioning engine. Moreover, injectors can easily be replaced and spare injectors are stored nearby, thus minimizing the down time. Also additional sensors have been installed on these engines to fully utilize the experiment for collecting failure-related data and gaining insights in relevant indicators.

6.3. *Improved labelling of field data by experts*

Although failures are typically rare due to PM, it is crucial to properly label the few failures that do occur in practice. Expert judgement of a skilled person can be used to execute a root cause analysis, specifying what failure mechanism occurred, and to register accumulated operating hours and conditions at failure.

Additionally, for preventively replaced parts, careful inspection by an expert could provide very useful labels and insights: quantification of the actual condition of the replaced component (option 2 in Fig. 6) offers crucial feedback on whether that replacement was just-in-time or way-to-early. These insights can assist in adjusting PM intervals or threshold values for condition monitoring.

6.4. *Numerical simulations*

If real data cannot easily be collected, simulation models can sometimes be used as alternative. Any simulation code that is based on the underlying physics (e.g. CMAPSS for aero-engines), can be used for that purpose. Defects or degradation can be incorporated in the models, and the associated data can quickly be generated. The obvious drawback is that simplifications or incorrect representations in the model lead to deviating data. However, as long as the dominant effects are included, detailed tuning to a real system can be achieved by updating the model with real data (see section 5).

Rijsdijk et al. (2024) demonstrate how numerical models can enhance diagnostics when only a limited number of sensors is present. Current work extends these relatively simple models to a simulation model of a chilled water system, simulating several realistic faults. As the simulation quantifies variables (temperature, pressure) at any location, an infinite number of virtual sensors is present. This allows to analyse how many (and which) sensors are minimally required to properly diagnose the system.

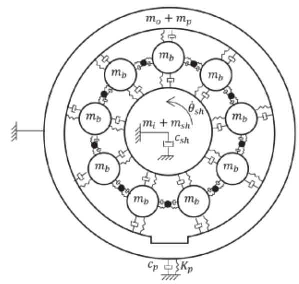

Fig. 8. Model of a bearing with outer race defect as represented by a bond graph (Keizers et al. 2025b).

Finally, Keizers et al. (2024) propose to use a bond graph model for a bearing containing an outer race defect, see Fig. 8. This model allows to derive a direct condition measurement (defect size) from the associated indirect vibration measurement, which is required for accurate prognostics (option 2 in Fig. 6).

6.5. *Standardized registration of (failure) data*

From the previous subsections it is clear that the generation and collection of data, both from the lab and the field, is a significant facilitator for smart maintenance algorithms. Around the world, several benchmark datasets (CWRU, FEMTO-ST, NLN-EMP) and many experiment-tal set-ups are available. However, these all use their own data format and especially the reporting of the meta-data (details of faults, operating conditions) is very inconsistent and incomplete. This makes it very hard for algorithm developers to use the data. To overcome this challenge, the *Dutch Prognostics Lab* (Tinga 2024) was initiated by the author, aiming to define a standard way of reporting failure data and associated meta-data and stimulating the sharing of datasets amongst algorithm developers.

Currently, a first version of a standardized template is nearing completion, based on the ISA metadata framework from the field of biology, and adapted to store diagnostic and prognostic test data and descriptors. Software code is now developed to translate inputs to the ISA standard, and store this in a database.

7. Conclusion

This paper has identified some of the major barriers for the application of PdM in industrial practice, which are predominantly related to the lack of relevant data. Three solution directions have been presented: accepting the data position and finding the best fitting method, circumventing the lack of data by using physical models and generating additional data. Following these directions is expected to yield a wider application of PdM in industrial practice.

Acknowledgment

The discussions with and ideas of S.J. Bruinsma, L.S. Keizers, R. Loendersloot, S. Mavroudis, A.A. Meghoe, C. Rijsdijk, N. Silveira, W.W. Tiddens, and M.J.R. van de Wijnckel are gratefully acknowledged.

References

Akkermans, H., Basten, R., Zhu, Q., and Van Wassenhove, L. (2024). Transition paths for condition-based maintenance-driven smart services. *Journal of Operations Management*.

Alves da Silveira, N.N., Meghoe, A.A., and Tinga, T. (2023). Integration of multiple failure mechanisms in a life assessment method for centrifugal pump impellers. *Advances in mechanical engineering* 15 (6).

Anonym. 2025. "Reliability Data Sources." NTNU. https://www.ntnu.edu/ross/info/data.

Bruinsma, S., Geertsma, R.D., Loendersloot, R., and Tinga, T. (2024). Motor Current and Vibration Monitoring Dataset for various Faults in an E-motor-driven Centrifugal Pump. *Data in Brief* 52: 1-16.

de Jonge, B., and Scarf, P.A. (2020). A review on maintenance optimization. *European Journal of Operational Research* 285 (3): 805-824.

de Pater, I., Reijns, A., and Mitici, M. (2022). Alarm-based predictive maintenance scheduling for aircraft engines with imperfect Remaining Useful Life prognostics. *Reliability Engineering & System Safety* 221.

Grubic, T., Redding, L., Baines, T., and Julien, D. (2011). The adoption and use of diagnostic and prognostic technology within UK-based manufacturers. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture* 225 (8): 1457-1470.

Keizers, L.S., Loendersloot, R., and Tinga, T. (2024). Towards a Hybrid Framework for Prognostics with Limited Run-to-Failure Data. European conference of the PHM society, Prague, Czech Republic.

Keizers, L.S., Loendersloot, R., and Tinga, T. (2025a). Bayesian filtering based prognostic framework incorporating varying loads. *Mechanical Systems and Signal Processing* 224.

Keizers, L.S., Loendersloot, R., and Tinga, T. (2025b). Utilizing bond graphs for quantitative fault diagnostics to enhance Bayesian filtering-based failure prognostics. *to be submitted*: 1-15.

Khan, S., and Yairi, T. (2018). A review on the application of deep learning in system health management. *Mechanical Systems and Signal Processing* 107: 241-265.

Mavroudis, S., and Tinga, T. (2025). Application of transfer learning on physics-based models to enhance vessel shaft power predictions. under review.

Pinciroli, L., Baraldi, P., and Zio, E. (2023). Maintenance optimization in industry 4.0. *Reliability Engineering & System Safety* 234.

Rijsdijk, C., Wijnckel, M.J.R.v.d., and Tinga, T. (2024). A maturity framework for data driven maintenance. European Conference of the Prognostics and Health Management Society, Prague, Czech Republic.

Saxena, A., Goebel, K., Simon, D., and Eklund, N. 2008. "Damage propagation modeling for aircraft engine prognostics." International Conference on Prognostics and Health Management.

Tamssaouet, F., Nguyen, K.T.P., Medjaher, K., and Orchard, M.E. (2022). System-level failure prognostics: Literature review and main challenges. *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability* 237 (3): 524-545.

Tiddens, W.W., Braaksma, J., and Tinga, T. (2022). Exploring predictive maintenance applications in industry. *Journal of Quality in Maintenance Engineering* 28 (1): 68-85.

Tiddens, W.W., Braaksma, J., and Tinga, T. (2023). Decision Framework for Predictive Maintenance Method Selection. *Applied Sciences* 13 (3).

Tinga, T. 2024. "Dutch Prognostics Lab." University of Twente. https://www.utwente.nl/en/et/ms3/research-chairs/dbm/projects/currentprojects/dutchprognosticslab/.

Zonta, T., da Costa, C.A., da Rosa Righi, R., de Lima, M.J., Silveira da Trindade, E., and Li, G.P. (2020). Predictive maintenance in the Industry 4.0: A systematic literature review. *Computers & Industrial Engineering* 150 (106889): 1-17.