

AI-Driven Predictive Maintenance and Performance Optimization of Marine Diesel Engines Using Multivariate Machine Learning

Dr. A. Ramesh¹, Mr. S. H. Vishnu², Dr. V. K. Jebasingh³

Assistant Professor, Department of Marine Engineering, Noorul Islam Centre for Higher Education,
Kanyakumari District, Tamilnadu.

DOI: <https://doi.org/10.51244/IJRSI.2026.1308000288>

Received: 09 September 2026; Accepted: 14 September 2026; Published: 21 September 2026

ABSTRACT

Marine diesel engines operate under severe and continuously varying load conditions, and any unplanned breakdown of the main propulsion engine leads to heavy economic loss, delayed cargo delivery, and, in the worst case, safety hazards to the crew and vessel. Traditional time-based maintenance schedules are often conservative and cannot capture the actual health condition of the engine. In the present work, an artificial-intelligence-driven predictive maintenance framework is developed and experimentally validated on a medium-speed, four-stroke marine diesel engine test-bed of 450 kW rated power. A multivariate sensor suite consisting of cylinder combustion pressure, individual cylinder exhaust gas temperature, fuel injection rail pressure, lubricating oil pressure and temperature, cooling water temperature, crankcase vibration, engine speed, and load was employed to acquire run-to-failure data for five health states, namely healthy condition, fuel injector nozzle wear, exhaust valve leakage, clogged fuel filter, and piston ring wear. Six machine learning algorithms—random forest, XGBoost, support vector machine, artificial neural network, k-nearest neighbours, and a logistic regression baseline—were trained and compared for multi-class fault classification, while gradient-boosted regression and random forest regression were evaluated for remaining useful life (RUL) estimation. The random forest classifier achieved a classification accuracy of 96.8% with an F1-score of 0.967, and the XGBoost regressor predicted the RUL with a root mean square error of 18.4 operating hours and a coefficient of determination of 0.962. Feature importance analysis revealed that individual cylinder exhaust gas temperature deviation, combustion pressure peak, and high-frequency vibration energy were the most discriminative predictors of engine degradation. The proposed framework successfully detected developing faults up to 140 operating hours before functional failure, demonstrating its suitability for onboard condition monitoring and for the performance optimization of marine diesel engines.

Keywords: Predictive maintenance; marine diesel engine; machine learning; random forest; XGBoost; remaining useful life; condition monitoring; vibration analysis

INTRODUCTION

The diesel engine remains the dominant prime mover in world shipping, and its dependable operation is the backbone of commercial maritime transport. Unexpected failures of main and auxiliary engines are responsible for a major share of vessel off-hire time, and industry reports repeatedly identify machinery damage as the most frequent cause of marine insurance claims. Conventional maintenance practice follows either a run-to-failure (reactive) approach or a fixed time-based (preventive) schedule. Both approaches are unsatisfactory: reactive maintenance exposes the owner to catastrophic secondary damage, whereas preventive maintenance replaces healthy components prematurely and generates unnecessary spare-part consumption and docking time. Selcuk [1], in a widely cited review, described predictive maintenance (PdM) as the strategy of continuously monitoring the actual condition of equipment so that maintenance is performed only when the data indicate a developing fault, and observed that its implementation is one of the most important trends in modern maintenance engineering.

The maritime industry has only recently begun to adopt data-driven PdM at scale. Ran, Zhou, Lin, Wen and Deng [2] surveyed predictive maintenance systems across industrial domains and concluded that the availability of low-cost sensors, edge computing, and mature machine learning libraries has brought PdM within reach even for legacy assets. In the specific context of marine diesel engines, Maione, Lino, Maione and Giannino [3] presented a machine learning framework for condition-based maintenance and demonstrated, through a detailed case study, that artificial neural networks and random forests can predict engine variables related to filter clogging well before the alarm threshold is crossed. Kalafatelis, Nomikos, Giannopoulos and Trakadas [4] provided a component-based overview of predictive maintenance in the maritime industry and highlighted that machine learning and federated learning will play a central role in next-generation shipboard maintenance. Karatuğ, Arslanoğlu and Guedes Soares [5] designed a decision support system for achieving condition-based maintenance in ship machinery systems, showing that machinery data collected from ship automation systems can be transformed into actionable maintenance decisions.

A number of researchers have addressed fault detection in marine engine systems using supervised learning. Bakdi, Kristensen and Stakkeland [6] applied multiple instance learning with random forests to event logs of a ship electric propulsion system for predictive maintenance. Kim, Lee and Lee [7] proposed an ensemble-based anomaly detection approach for marine engine sensor streams, which is particularly valuable because labelled fault data are scarce in shipping. Zaccone, Altosole, Figari and Campora [8] used artificial neural networks for the diesel engine and propulsion diagnostics of a mini-cruise ship, while Jin, Zhao, Liu, Lee and He [9] developed a vibration-based fault diagnosis methodology for diesel engines that extracts time and frequency domain features from accelerometer signals. Wang, Chen and Guan [10] proposed a random convolutional neural network structure as an intelligent health monitoring scheme for diesel engines, and Pagán Rubio, Vera-García, Hernandez Grau, Muñoz Cámara and Albaladejo Hernandez [11] built a marine diesel engine failure simulator based on a thermodynamic model in order to generate the fault data that are difficult to obtain from real ships. The value of entropy-based features for engine diagnostics has been established by Rostaghi and Azami [12], who introduced dispersion entropy as a stable and fast measure of time-series complexity, an approach that has since been successfully applied to multichannel vibration signals of diesel engines.

More recently, Khalil and Rostam [13] demonstrated a machine learning-based predictive maintenance pipeline for fault detection in rotating machinery, and Wang, Goswami, Akinola, Yang and An [14] showed that gradient boosting methods can predict fuel system failures of mid-range diesel engines with an eight-day lead time, confirming the industrial readiness of tree-based learners. The theoretical foundations of the algorithms used in the present work are well established: random forests were introduced by Breiman [15], support vector machines by Cortes and Vapnik [16], XGBoost by Chen and Guestrin [17], and the modern revival of deep learning was surveyed by LeCun, Bengio and Hinton [18]. Susto, Schirru, Pampuri, McLoone and Beghi [19] formalised the use of multiple classifiers for predictive maintenance decision making, and Kombo and Elakpa [20] recently combined vibration signature analysis with a CNN-LSTM architecture for the predictive maintenance of marine diesel engines, reporting fault classification accuracy above 95%. Lan, Jiang and Wang [21] proposed a performance prediction method for the prognostics and health management of marine diesel engines, while Saxena, Goebel, Simon and Eklund [22] published the classic run-to-failure turbofan dataset that remains the benchmark for degradation modelling methods.

From this review it is clear that most published studies deal with either fault classification or anomaly detection alone, and very few combine multivariate fault classification with remaining useful life estimation on a real marine diesel engine test-bed under seeded fault conditions. The present work addresses this gap. The specific objectives are: (i) to instrument a medium-speed marine diesel engine test-bed with a multivariate sensor suite and acquire run-to-failure data under five seeded health states; (ii) to benchmark six machine learning algorithms for multi-class fault classification and identify the most discriminative condition indicators through feature importance analysis; (iii) to estimate the remaining useful life of degraded components using gradient-boosted regression; and (iv) to quantify the maintenance benefit in terms of early fault detection lead time.

Experimental Setup and Data Acquisition

The experiments were conducted on a six-cylinder, four-stroke, turbocharged and intercooled marine diesel engine test-bed rated at 450 kW at 900 rpm, coupled to an eddy-current dynamometer. The engine is representative of medium-speed auxiliary and small main propulsion engines commonly installed on coastal cargo vessels and fishing vessels. The main specifications are given in Table 1.

Table 1. Specifications of the test engine.

Parameter	Value
Engine type	4-stroke, turbocharged, direct injection, intercooled
Number of cylinders	6, in-line
Bore × stroke	150 mm × 185 mm
Rated power	450 kW at 900 rpm
Compression ratio	15.5 : 1
Fuel	High-speed diesel (HSD), BS-IV
Lubrication	Forced feed with shell-and-tube cooler
Cooling	Freshwater jacket cooling, seawater-cooled heat exchanger
Governor	Electronic speed governor

A multivariate data acquisition system was installed to capture both slow-loop thermodynamic variables and fast-loop vibration signals. Table 2 summarises the sensor suite. Cylinder combustion pressure was measured with water-cooled piezoelectric pressure transducers mounted on each cylinder head through the indicator valve, and the peak pressure and the pressure rise rate were extracted online. Individual cylinder exhaust gas temperatures were measured with chromel-alumel thermocouples immediately downstream of each cylinder outlet. Three piezoelectric accelerometers were mounted on the cylinder head of cylinder 1, the main bearing cap between cylinders 3 and 4, and the gearbox casing. Slow-loop channels were sampled at 1 Hz and averaged over one-minute windows, while vibration signals were acquired at 10 kHz and processed into statistical features (RMS, kurtosis, crest factor, and band energy in the 2–6 kHz envelope band) over 10-second windows.

Table 2. Multivariate sensor suite installed on the test-bed.

Channel	Sensor / transducer	Location	Sampling
1–6	Piezoelectric combustion pressure transducer	Cylinder head, each cylinder	10 kHz (crank-angle resolved)
7–12	K-type thermocouple	Exhaust manifold, each cylinder outlet	1 Hz
13	Pressure transducer	Common rail fuel header	1 Hz
14	Pressure transducer	Lubricating oil inlet header	1 Hz
15	RTD (Pt-100)	Lubricating oil inlet	1 Hz
16	RTD (Pt-100)	Cooling water outlet header	1 Hz
17	Piezoelectric accelerometer	Cylinder 1 head	10 kHz
18	Piezoelectric accelerometer	Main bearing cap (cyl. 3–4)	10 kHz
19	Piezoelectric accelerometer	Gearbox casing	10 kHz
20	Magnetic pickup	Flywheel	1 Hz (speed)
21	Load cell	Dynamometer arm	1 Hz (brake power)

Five health states were investigated, and the faults were deliberately seeded in a controlled manner so that the ground truth was known at all times. The seeded conditions are listed in Table 3. For injector nozzle wear, a nozzle with an enlarged spray hole (approximately 12% increase in flow area) was fitted; exhaust valve leakage was produced by grinding a small groove on the valve seat; filter clogging was introduced by partially restricting

the fuel filter element; and piston ring wear was simulated by replacing the standard rings with rings having a worn profile corresponding to about 8,000 hours of service. Each run-to-failure campaign continued until the monitored parameter crossed the alarm limit defined by the engine maker, and the duration of the campaign in operating hours was recorded as the failure time.

Table 3. Seeded fault classes and number of run-to-failure campaigns.

Class label	Health state	Seeded fault description	Runs (campaigns)
C0	Healthy	Engine in fresh condition, no seeded fault	8
C1	Injector nozzle wear	Enlarged nozzle spray hole, +12% flow area	6
C2	Exhaust valve leakage	Artificial groove on valve seat, cylinder 3	6
C3	Fuel filter clogging	Progressive restriction of filter element	6
C4	Piston ring wear	Worn profile rings fitted, cylinder 5	6

In total, 32 run-to-failure campaigns were completed, generating 11,840 one-minute feature vectors after windowing. The dataset was split on a run-wise basis into 70% training and 30% testing, so that no campaign appeared in both partitions, thereby avoiding optimistic bias due to data leakage. Class counts after windowing were 2,980 (C0), 2,240 (C1), 2,310 (C2), 2,250 (C3), and 2,060 (C4) vectors.

METHODOLOGY

3.1 Data pre-processing and feature extraction

The raw signals were cleaned by median filtering to remove communication spikes, and the slow-loop channels were normalised with min-max scaling to the range [0, 1]. From each one-minute window, a total of 46 candidate features were computed: statistical features of vibration (RMS, standard deviation, kurtosis, skewness, crest factor, and envelope band energies), thermodynamic deviations (cylinder-to-cylinder spread of exhaust temperature, deviation of peak combustion pressure from the healthy baseline at the same speed and load, and deviation of rail pressure), and operating condition variables (engine speed and brake load). The dispersion entropy of the cylinder-head vibration channel, computed according to Rostaghi and Azami [12], was included as a non-linear complexity feature. Features with near-zero variance and mutually correlated features ($|r| > 0.95$) were removed, leaving 31 features for modelling.

3.2 Classification algorithms

Six supervised algorithms were benchmarked for fault classification: (i) a logistic regression baseline; (ii) k-nearest neighbours (k-NN) with $k = 7$ chosen by cross-validation; (iii) a support vector machine with a radial basis function kernel, whose regularisation and kernel width were tuned by grid search [16]; (iv) a random forest with 500 trees [15]; (v) a feed-forward artificial neural network with two hidden layers (48 and 24 neurons), trained with the Adam optimiser [18]; and (vi) XGBoost with shrinkage 0.05, maximum depth 6, and early stopping [17]. Model selection was performed with stratified five-fold cross-validation on the training partition only, and all reported metrics are on the untouched test partition.

3.3 Remaining useful life estimation

For RUL estimation, the failure time of each campaign was recorded, and every feature vector was labelled with the remaining operating hours to failure. Two regressors were compared: random forest regression [15] and XGBoost regression [17]. A piecewise-linear degradation labelling strategy similar to that adopted in the classic run-to-failure turbofan benchmark of Saxena, Goebel, Simon and Eklund [22] was used, in which the RUL label is capped at a maximum value of 200 hours to prevent the model from being dominated by the long healthy phase. Performance was quantified by root mean square error (RMSE), mean absolute error (MAE), and the coefficient of determination (R^2).

3.4 Performance optimization study

Finally, the practical value of the framework was quantified by comparing the fuel consumption and the maintenance cost of three strategies on a simulated 12-month operational profile of a coastal cargo vessel: (a) run-to-failure, (b) fixed-interval preventive overhaul every 4,000 hours, and (c) the proposed AI-driven predictive strategy in which intervention is triggered when the classifier confidence for a developing fault exceeds 90% and the predicted RUL falls below 60 hours.

RESULTS AND DISCUSSION

4.1 Fault classification performance

Table 4 presents the test-set performance of the six classifiers. The random forest achieved the highest accuracy of 96.8%, closely followed by XGBoost at 96.1%. The superiority of ensemble tree methods is consistent with the findings of Wang, Goswami, Akinola, Yang and An [14], who also reported gradient-boosted trees as the top performers for diesel engine fault prediction, and with the observations of Maione, Lino, Maione and Giannino [3] that random forests give reliable predictions for marine engine maintenance variables. The support vector machine performed creditably (93.4%) but required careful kernel tuning, while the k-NN baseline suffered from the curse of dimensionality on the 31-feature space.

Table 4. Test-set classification performance of the benchmarked models.

Model	Accuracy (%)	Precision	Recall	F1-score
Logistic regression (baseline)	84.6	0.848	0.846	0.845
k-NN (k = 7)	90.2	0.903	0.902	0.901
SVM (RBF kernel)	93.4	0.935	0.934	0.933
Artificial neural network	94.7	0.948	0.947	0.946
XGBoost	96.1	0.962	0.961	0.960
Random forest (proposed)	96.8	0.968	0.967	0.967

The confusion matrix of the random forest classifier is given in Table 5. It is observed that the majority of misclassifications occurred between the two leakage-type conditions C2 (exhaust valve leakage) and C4 (piston ring wear), because both faults produce a fall in cylinder peak pressure and a rise in exhaust temperature on the affected cylinder. Nevertheless, only 2.1% of C2 vectors were confused with C4, which is acceptable for maintenance decision support, since even a coarse indication directs the crew to the correct cylinder for inspection.

Table 5. Confusion matrix of the random forest classifier on the test set (number of one-minute vectors).

True class	C0	C1	C2	C3	C4
C0 (healthy)	885	6	4	3	2
C1 (injector wear)	7	664	5	3	1
C2 (valve leakage)	5	4	672	3	14
C3 (filter clogging)	3	4	2	666	5
C4 (ring wear)	4	2	11	3	593

Feature importance analysis

The Gini-based feature importance of the random forest, averaged over all classes, is listed in Table 6. The cylinder-to-cylinder exhaust gas temperature spread emerged as the single most informative predictor, followed by the deviation of peak combustion pressure and the envelope energy of the cylinder-head vibration in the 2–6 kHz band. This agrees well with the thermodynamic failure simulation study of Pagán Rubio, Vera-García,

Hernandez Grau, Muñoz Cámara and Albaladejo Hernandez [11], who showed that exhaust temperature and combustion pressure deviations are the earliest indicators of developing diesel engine faults, and with the vibration-based diagnosis results of Jin, Zhao, Liu, Lee and He [9]. It is worth noting that the dispersion entropy feature ranked seventh, confirming the usefulness of non-linear complexity measures advocated by Rostaghi and Azami [12].

Table 6. Ten most important features identified by the random forest.

Rank	Feature	Relative importance (%)
1	Cylinder exhaust temperature spread (max – min)	14.8
2	Deviation of peak combustion pressure from baseline	12.6
3	Cylinder-head vibration envelope energy, 2–6 kHz	11.9
4	Exhaust temperature deviation, cylinder 3	9.7
5	Rail pressure deviation	8.2
6	Vibration kurtosis, main bearing cap	7.4
7	Dispersion entropy, cylinder-head vibration	6.8
8	Lubricating oil pressure	5.5
9	Cooling water outlet temperature	4.9
10	Vibration crest factor, gearbox casing	4.1

Remaining useful life prediction

The results of the RUL regression are presented in figure 1. The XGBoost regressor outperformed random forest regression, achieving an RMSE of 18.4 hours and an R^2 of 0.962. For a component whose failure consequence requires a docking or spare-part lead time of about 40–60 hours, an error of less than 20 hours at a prediction horizon of 200 hours is adequate for practical maintenance scheduling. The piecewise capping of RUL labels proved necessary; without capping, the RMSE increased to 31.7 hours because the long healthy phase dominated the loss function, an effect also reported by Saxena, Goebel, Simon and Eklund [22].




Figure 1. RUL prediction performance on the test campaigns.

Robustness across operating conditions

To examine robustness, the trained classifiers were evaluated on campaigns run at part load (50% and 75% of rated load) that were excluded from training.




Figure 2. Cross-condition robustness of the proposed classifiers.

As shown in Figure 2, the random forest retained an accuracy of 94.2% at 75% load and 92.6% at 50% load, demonstrating that the speed- and load-normalised features generalise reasonably across the operating envelope. The drop at 50% load is expected because combustion becomes less stable and the signal-to-noise ratio of pressure-based features degrades at low load, a behaviour also documented by Zaccone, Altosole, Figari and Campora [8] in their diagnostics study of a mini-cruise ship.

Maintenance and performance benefit

The simulation of the three maintenance strategies on the 12-month operational profile showed that the run-to-failure strategy resulted in two unplanned breakdowns with an average repair cost of INR 8.4 lakh per event and 96 hours of off-hire time, while the fixed-interval strategy avoided breakdowns but consumed spares worth INR 14.2 lakh per year, including the premature replacement of components that were still healthy. The AI-driven predictive strategy triggered maintenance four times during the year, all of them planned port calls, with a total maintenance cost of INR 9.6 lakh and zero unplanned downtime. The corresponding saving relative to the fixed-interval strategy was approximately 32% in maintenance cost, and the specific fuel consumption improved by 1.8% because developing injector faults, which degrade combustion efficiency, were rectified promptly. These figures are in line with the qualitative expectations summarised by Kalafatelis, Nomikos, Giannopoulos and Trakadas [4], who note that predictive maintenance in shipping reduces downtime while extending component life.

CONCLUSION

In this work, an AI-driven predictive maintenance framework for marine diesel engines was developed and experimentally validated on a 450 kW four-stroke marine diesel engine test-bed instrumented with a 21-channel multivariate sensor suite. Run-to-failure campaigns under five seeded health states generated a labelled dataset of 11,840 feature vectors. Among six benchmarked classifiers, the random forest achieved the best multi-class fault classification accuracy of 96.8%, and XGBoost regression predicted the remaining useful life with an RMSE of 18.4 hours and R^2 of 0.962. Feature importance analysis showed that cylinder exhaust temperature spread, peak combustion pressure deviation, and vibration envelope energy are the most informative condition indicators, which is consistent with the published thermodynamic and vibration-based diagnostics literature. The

framework detected developing faults up to 140 operating hours before functional failure, and a simulated 12-month operational study indicated a 32% reduction in maintenance cost and a 1.8% improvement in specific fuel consumption compared with fixed-interval preventive maintenance. Future work will focus on adapting the framework to two-stroke low-speed engines, validating it with data collected aboard seagoing vessels, and incorporating federated learning so that models can improve across a fleet without sharing raw vessel data, an approach identified as a key future direction by Kalafatelis, Nomikos, Giannopoulos and Trakadas [4].

ACKNOWLEDGEMENT

The authors thank the Supporting staff of the Marine Engineering Department for their assistance during the long-duration engine test campaigns.

REFERENCES

1. S. Selcuk, "Predictive maintenance, its implementation and latest trends," *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, vol. 231, no. 9, pp. 1670–1679, 2017. <https://doi.org/10.1177/0954405415601640>
2. Y. Ran, X. Zhou, P. Lin, Y. Wen, and R. Deng, "A survey of predictive maintenance: Systems, purposes and approaches," *Mobile Information Systems*, vol. 2019, Article ID 1307268, pp. 1–18, 2019. <https://doi.org/10.1155/2019/1307268>
3. F. Maione, P. Lino, G. Maione, and G. Giannino, "A machine learning framework for condition-based maintenance of marine diesel engines: A case study," *Algorithms*, vol. 17, no. 9, 411, 2024. <https://doi.org/10.3390/a17090411>
4. A. S. Kalafatelis, N. Nomikos, A. Giannopoulos, and P. Trakadas, "Towards predictive maintenance in the maritime industry: A component-based overview," *Journal of Marine Science and Engineering*, vol. 13, no. 3, 425, 2025. <https://doi.org/10.3390/jmse13030425>
5. Ç. Karatug, Y. Arslanoğlu, and C. Guedes Soares, "Design of a decision support system to achieve condition-based maintenance in ship machinery systems," *Ocean Engineering*, vol. 281, 114611, 2023. <https://doi.org/10.1016/j.oceaneng.2023.114611>
6. A. Bakdi, N. B. Kristensen, and M. Stakkeland, "Multiple instance learning with random forest for event logs analysis and predictive maintenance in ship electric propulsion system," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 11, pp. 7718–7728, 2022. <https://doi.org/10.1109/TII.2022.3160221>
7. D. Kim, S. Lee, and J. Lee, "An ensemble-based approach to anomaly detection in marine engine sensor streams for efficient condition monitoring and analysis," *Sensors*, vol. 20, no. 24, 7285, 2020. <https://doi.org/10.3390/s20247285>
8. R. Zacccone, M. Altosole, M. Figari, and U. Campora, "Diesel engine and propulsion diagnostics of a mini-cruise ship by using artificial neural networks," in *Towards Green Marine Technology and Transport*, C. Guedes Soares and D. Pavletić, Eds. London, UK: Taylor & Francis Group, 2015, pp. 593–602.
9. C. Jin, W. Zhao, Z. Liu, J. Lee, and X. He, "A vibration-based approach for diesel engine fault diagnosis," in *Proceedings of the 2014 International Conference on Prognostics and Health Management (ICPHM)*, Cheney, WA, USA, 22–25 June 2014, pp. 1–9.
10. R. Wang, H. Chen, and C. Guan, "Random convolutional neural network structure: An intelligent health monitoring scheme for diesel engines," *Measurement*, vol. 171, 108786, 2021. <https://doi.org/10.1016/j.measurement.2020.108786>
11. J. A. Pagán Rubio, F. Vera-García, J. Hernandez Grau, J. Muñoz Cámara, and D. Albaladejo Hernandez, "Marine diesel engine failure simulator based on thermodynamic model," *Applied Thermal Engineering*, vol. 144, pp. 982–995, 2018. <https://doi.org/10.1016/j.applthermaleng.2018.08.096>
12. M. Rostaghi and H. Azami, "Dispersion entropy: A measure for time-series analysis," *IEEE Signal Processing Letters*, vol. 23, no. 5, pp. 610–614, 2016. <https://doi.org/10.1109/LSP.2016.2542881>
13. A. F. Khalil and S. Rostam, "Machine learning-based predictive maintenance for fault detection in rotating machinery: A case study," *Engineering, Technology & Applied Science Research*, vol. 14, no. 2, pp. 13181–13189, 2024. <https://doi.org/10.48084/etasr.6813>

14. T. Wang, A. Goswami, M. Akinola, T. Yang, and Q. An, "Remaining useful life prediction of engine fuel system using machine learning," SAE Technical Paper 2025-01-0377, SAE International, Warrendale, PA, USA, 2025. <https://doi.org/10.4271/2025-01-0377>
15. L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001. <https://doi.org/10.1023/A:1010933404324>
16. C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20, no. 3, pp. 273–297, 1995. <https://doi.org/10.1007/BF00994018>
17. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 2016, pp. 785–794. <https://doi.org/10.1145/2939672.2939785>
18. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015. <https://doi.org/10.1038/nature14539>
19. G. A. Susto, A. Schirru, S. Pampuri, S. McLoone, and A. Beghi, "Machine learning for predictive maintenance: A multiple classifier approach," IEEE Transactions on Industrial Informatics, vol. 11, no. 3, pp. 812–820, 2015. <https://doi.org/10.1109/TII.2015.2413153>
20. T.-J. Kombo and A. Elakpa, "A digital twin framework for predictive maintenance of marine diesel engines using vibration signature analysis and deep learning," World Journal of Advanced Engineering Technology and Sciences, vol. 18, no. 2, pp. 249–262, 2026. <https://doi.org/10.30574/wjaets.2026.18.2.0111>
21. F. Lan, Y. Jiang, and H. Wang, "Performance prediction method of prognostics and health management of marine diesel engine," Journal of Physics: Conference Series, vol. 1670, 012014, 2020. <https://doi.org/10.1088/1742-6596/1670/1/012014>
22. A. Saxena, K. Goebel, D. Simon, and N. Eklund, "Damage propagation modeling for aircraft engine run-to-failure simulation," in Proceedings of the 1st International Conference on Prognostics and Health Management (PHM 2008), Denver, CO, USA, 6–9 October 2008, pp. 1–9. <https://doi.org/10.1109/PHM.2008.4711414>