www.ijarp.com ISSN 2456-9992 Page: 01-15

International Journal Advanced Research Publications

Review Article

Volume: 01
Issue: 03

PREDICTIVE MAINTENANCE IN THE ERA OF AI: TECHNOLOGIES, CHALLENGES, AND FUTURE DIRECTIONS

[*1]Mrs. Swati Chiplunkar, [2]Dr. Sunil Wankhade, [3]Mrs. Kajal Patel

[1,3]Thakur College of Engineering & Technology, Mumbai, India.

[2]Rajiv Gandhi Institute of Technology Mumbai, India.

Article Received: 11 November 2025, Article Revised: 01 December 2025, Published on: 21 December 2025

***Corresponding Author: Mrs. Swati Chiplunkar**

Thakur College of Engineering & Technology, Mumbai, India.

DOI: https://doi-doi.org/101555/ijarp.5504

ABSTRACT

Predictive maintenance (PdM) has completely changed the way industries are operated from a reactive and preventive mind set to real-time, data-driven, and condition-based decision-making for maintenance activities. As a result of the continuous capture of condition data leading to signal processing and artificial intelligence (AI) being utilized for maintenance, this can reduce downtime by enabling maintenance of equipment before catastrophic failure and simultaneously improve reliability and lower costs. This paper reviews the development of PdM technologies and practices over time, tracing a path from more traditional, adaptive statistical and signal analysis practices, advancing to and including more recent developments in machine learning, and deep learning algorithms. We compare the various approaches based on performance criteria such as accuracy, scalability, ease of understanding, and cost of implementation and describe the pros and cons of each approach when used with different industrial applications. In addition, we discuss an emerging scheme regarding explainable AI (XAI) and causal discovery which addresses the challenge of trust, and transparency with approaches employing black-box models that is critical in situations where safety is paramount. We also describe the usefulness of new paradigms, such as data categorization approaches, multimodal data fusion, edge and federated learning, and quantum computing, for enhancing scalability and real-time deployment. By collecting existing applications and examining positives and negatives, applicable principles, and main trade-offs to usability outcomes, we have noted areas of research that are lacking, such as understanding interpretability to the PdM system created by the data utilized, availability of data-based

information, and deployment, which allows us to create next steps for developing PdM systems that are accurate, interpretable, and scalable within the Industry 4.0 space.

INDEXTERMS: About four key words or phrases, in alphabetical order, separated by commas, with only the first index term capitalized. All terms following the initial term should be lowercase unless they are proper nouns, in which case they should have an initial cap.

INTRODUCTION

Predictive maintenance (PdM) is emerging as a common theme of today's industrial activities and changing the framework of maintenance from reactive and planned maintenance to data-driven, and condition-based maintenance 1 2. Given the rise in complexity and interconnectedness of industrial systems, the consequences of unplanned equipment failures are particularly costly to industry, with unplanned downtime costing the manufacturing industry billions annually 3,4. Thus, the transition from corrective maintenance of an asset, conducted after a failure, and preventive maintenance that performs maintenance tasks regardless of the asset's operating condition, to predictive maintenance, is grounded in better optimization of maintenance activities based on the real-time condition of equipment health, performance measures, and metrics. The concept of predictive maintenance is to give maintenance teams the ability to foresee equipment failure prior to its occurrence, potentially allowing maintenance practitioners to conduct a maintenance activity during a planned downtime, incur a lower maintenance cost, extend the life of the asset, and ultimately increase performance objectives5,6. Because predictive maintenance is developed from continuous monitoring of dynamic parameters of an asset, operational condition and state of the process and equipment conditions, and records of past maintenance activities, environmental and prognostic conditions, predictive maintenance is based on the development of specific description models related to operational performance behaviour and specific degradation models for the physical asset condition 7.

The performance of predictive maintenance methods is fundamentally relational to the quality and interpretability of the models 89. While classical statistical models might provide fundamental insights into equipment actions, they are unlikely to characterize the nonlinear, complex relationships that exist in an industrial setting 10. Industrial processes both include and surround us with a variety of data most prominently including time-series sensor data, categorical operational parameters, maintenance logs, or environmental parameters11. Moreover, the dynamic character of industrial processes creates temporal dependencies, and aging mechanisms which would ask for advanced modelling techniques in order to respond to those changing conditions.

The presence of multiple, simultaneous degradation processes (with varied failure mechanisms and timescales) adds to the complexity of predictive maintenance. Failures are often driven by complex interdependencies across multiple parameters which requires an understanding not only of the relevant failure predictive parameters, but how those relevant parameters interact with each other and exert influences over time. The ever-increasing volume and velocity of industrial data being realized from the widespread introduction of IoT sensors, and Industry 4.0 initiatives, opens-up both a challenge and an opportunity in relation to the implementation of predictive maintenance1213. Access to high resolution, multi-dimensional data can provide unprecedented insights into the performance of equipment, yet the complexity of the data will require an analytic approach that can find meaningful patterns while still supporting computational efficiency and practical application 14.

ARTIFICIAL INTELLIGENCE AS A TRANSFORMATIVE FORCE IN PREDICTIVE MAINTENANCE

The use of artificial intelligence (AI) for predictive maintenance demonstrates a very significant improvement leading to better, more accurate, reliable and actionable predictions of failure1516. Scheduling predictive maintenance is mainly based on the use of machine learning (ML) algorithms, which practitioners love, based on their ability to identify automatically complex nonlinear patterns in high-dimensional feature spaces and adapt to changing operating conditions; both of which have been one of the significant weaknesses of statistical algorithms17. In addition to traditional models, deep neural network architectures such as convolutional neural networks for spatial and spectral feature extraction from sensor signals, and recurrent neural networks for representing temporal dependencies, have been shown to increase performance capturing complex relationships between different operational parameters and indicators of equipment health18, 19]. Furthermore, deep learning architectures and ML algorithms have been verified for their ability to capture the nonlinear, time-dependent nature of degradation processes, and effectively accommodate the vast and diverse volumetric data types typically generated in industrial environments (e.g., vibration, acoustic, thermal, and image) [20]. Advanced algorithms, such as ensemble models, and advanced ML algorithms, not only leverage predictions from multiple ML algorithms, along with an uncertainty quantification, to increase the predictive reliability of diagnostic models, they also enable AI models to dynamically act live and learn continuously from streaming data, which leads to more adaptive and responsive predictive maintenance strategies.

Regardless of these advancements, there are major issues to face in adopting it on an industrial level. The sophistication of modern AI models, particularly deep learning models, can lead to systems that are "black boxes" which have models that give highly accurate predictions but limited explainability, making it difficult for people in the field to understand how recommendations for maintenance were reached [21], [22]. The lack of transparency presents challenges in safety critical industries for explaining decisions and actions [23], [24] to support operator trust in the operation of their systems. Safety critical industries, where people depend on the safety and performance of systems, require an understanding of how the model arrived at results to maintain trust and implement proper procedures (for projects with regulatory oversight) while making informed decisions. Addressing all necessary challenges to deployment of more modern AI models will not only require innovations in explainable AI but also consideration of deployment issues, i.e. computer cost, data access, how to combine the AI with human workflows, and providing proper training of human operators to use AI.

OVERVIEW OF PREDICTIVE MAINTENANCE TECHNOLOGIES

Predictive maintenance involves the use of data-driven approaches and algorithms to predict equipment failures and improve maintenance. Predictive maintenance's technologies involve data acquisition, data preparation, failure prediction, and maintenance optimization, and have evolved over the years through advances in sensors, computing, and analytics. The adoption of predictive maintenance solutions is contingent on the type of equipment, the modes of failure for the equipment, availability of data, computational demands, and interpretability. It takes careful design of effective frameworks to balance the strengths and weaknesses of these technologies to provide valid and actionable information within the constraints known for industrial resources. Time-domain analysis is one of the foundational methods for processing equipment performance data, focusing on how signals change over time. Time-domain analysis is beneficial for the identification of transient events, to analyse trends, and to capture the evolution of equipment health indicators over time [25].

Statistical time-domain features allow the foundation of a predictive maintenance system, as these features provide quantifiable and statistical measures of the signal characteristics, which can be indicative of the condition of the equipment. Root Mean Square (RMS) values are useful as a proxy for overall vibration energy and as a particularly sensitive indicator of the presence of developing faults in rotating machines [7]. Peak values and crest factors indicate the presence of impulsive events indicative of bearing failure or gear tooth damage. The measurements for skewness and kurtosis define the statistical nature of the signal amplitudes, with significant deviations from normal being indicative of faulty conditions [26].Trend Analysis Methods emphasize discovering differences over a prolonged time period in the behaviour of equipment that would reflect long-term degradation. Moving averages and exponential smoothing methods provide filtering of short-term variability while preserving the trends indicated of the health evolution of the equipment. Regression analysis and curve fitting methods are appropriate to quantify equipment degradation and extrapolate trending behaviours into future state predictions [27].

Autoregressive and Time Series Models offers more advanced methods for capturing the temporal dependency behaviour of monitoring data of the equipment. The AR, MA, and ARMA methods have been shown to represent the underlying stochastic processes of equipment behaviour and could help identify an anomaly as a deviation from expected temporal behaviour [28]. These different approaches provide a means of being more sensitive to the dynamic behaviour of equipment responses and may lead to prior identification of subtle changes in behaviour that precede failure.

Frequency-Domain Analysis Frequency-domain analysis converts time-domain signals into their frequency components, and may provide spectral features that are more directly connected to fault mechanisms than time-domain features. This forms the basis of many predictive maintenance strategies, especially for rotating machinery where various fault types create frequency-based signatures [29].

Fast Fourier Transform (FFT) analysis is the most commonly used method in predictive maintenance in that it converts time domain signals into their frequency based components, allowing the analyst to isolate frequency peaks associated with particular phenomena, such as shaft speed, meshing gear frequency, or bearing element passing frequency [30]. Power spectral density measures based on FFT analysis will provide the quantitative measure of energy pass-through in frequency bins and can be used for the detection of incipient faults through observation of changes in spectral features.

Advanced Spectral Analysis Methods allow us to go beyond merely conducting an FFT analysis, providing improved resolution and fault characterization. Welch's power spectral density estimation produces more reliable spectrum estimates for noisy industrial signals by segmenting a signal into shorter segments, estimating, and then averaging [31]. Parametric spectral estimation methods, such as autoregressive spectral analysis, can provide better frequency resolution when analysing limited length data. These are particularly useful when analysing transient phenomena or taking measurements over a short time scale.

Envelope Analysis and Demodulation Techniques are specialized frequency domain methods. They tend to show strengths when identifying bearing and gear faults, where fault-related impulses can be modulated onto a higher frequency signal from the bearing's resonance characteristic. This approach is often comprised of a band-pass filter for the high frequency components of our signal, followed by envelope detection and spectral analysis of the envelope signal [32].

Time-Frequency Analysis methods provide an essential link between time-domain and frequency-domain methods by providing information about the time occurrence of signal components alongside their frequency content. Time-frequency methods are particularly useful when signals are non-stationary and, therefore, their frequency content evolves over time. This is a common characteristic of industrial applications characterized by dynamic operating conditions or transient fault scenarios [33].

Short-Time Fourier Transform (STFT) is the most basic form of time-frequency analysis and involves performing the Fourier transform over a short and overlapping time window in order to observe the trend in frequency content over time [34]. STFT is most useful for analysing signals with slowly evolving frequency content and typically provides a reasonable compromise between temporal and frequency resolution. Nevertheless, STFT has a typical and prominent application in predictive maintenance scenarios when the monitored machinery operates at various speeds, or when the faults evolve with time.

Wavelet Transform Analysis Compared to the STFT, wavelet transform analysis has better time-frequency resolutions, especially when examining signals that contain both transient occurrences and steady-state components [35]. The continuous wavelet transform (CWT) is great for high-frequency component detection due to high-resolution time properties, and it can identify low-frequency components with good frequency resolution. The CWT is also particularly useful for detecting impulsive fault signatures, such as impacts generated by faulted bearings or damaged gear teeth. Discrete wavelet transforms (DWT) can break down signals efficiently into the frequency bands of interest via multiresolution analysis and help extract fault related features from complex signals in a range of frequency bands [36].

Advanced time-frequency methods include the Winger-Ville distribution, the Hilbert-Huang transform, and empirical mode decomposition for their individual capabilities and effectiveness at solving specific types of analysis challenges. The Hilbert-Huang transform combines empirical mode decomposition and Hilbert spectral analysis to provide a detailed analysis of nonlinear and non-stationary signals without the need for a priori basis functions [37]. Finally, these advanced methods have begun theoretical expansions and applications into the analysis of complex industrial signals when analysis results based on traditional techniques are insufficient.

The increasing interest in Explainable Artificial Intelligence (XAI) to support predictive maintenance is driven by operational and regulatory requirements [38]. In other words, increasingly, industrial operators and practitioners want to attain not just predictive accuracy, but insight into the prediction: why a particular failure is predicted, which parameters matter most regarding health or wellbeing of the equipment, and if operational conditions were to change, to what extent would it influence risk [39]. XAI is especially necessary in industries with a high public safety consequence, which are heavily regulated, as maintenance plans, inferences, or recommendations are required to be justified, auditable, and compliant with regulatory bodies [40]. Traditional XAI methods only partially facilitate explainability; feature importance analysis or local interpretation methods partially explain model representation and notion but do not typically or deeply provide insights about causality [41]. It is vital to know which variables and parameters are actively degrading the health or degrading the equipment, rather than simply knowing there is a correlation for events resulting in failures; and this information is essential for preventive maintenance to act on root causes rather than symptoms [42], [43].

To address these challenges, causal discovery has a growing presence as a promising approach for predictive maintenance. Causal discovery, by using approaches that model causal structures from observational data, overcomes the limitations of causal reasoning based on correlation, and also provides insight into the causes of equipment behaviour and failure [44], [45], [46]. In addition to being an alternative source of robustness and generalizability for predictive models, causal discovery will also provide more actionable information as to how maintenance decisions can connect back to the actual causes of the degradation of equipment [47]. Recent work has uncovered the potential for causal discovery and causal inference to be leveraged together alongside XAI in developing predictive maintenance systems that are valid with modern AI while using the explanations of causal inference [48], [49]. All in all, this synergistic approach provides an opportunity to develop maintenance frameworks that are valid, interpretable, trustworthy, and fundamentally linked to the causal mechanisms of equipment failures [50].

Comprehensive Method Comparison Analysis

1. Core Machine Learning Algorithms

Core ML algorithms exhibit trade-offs between performance, scalability, interpretability, and computational complexity. Random Forest [51] exhibits strong accuracy and can handle missing data, but takes up a lot of memory space and is not very interpretable. XGBoost [52] gives good prediction performance and scalability and handles missing data and cross-validation well, but needs delicate tuning of hyper-parameters and much more interpretation difficulty. Deep learning with CNNs [53], provides accuracy in image, signal, and pattern recognition but needs a large dataset, is computationally intensive, and like many deep learning approaches, can be considered "black-box" approaches. DBSCAN [54] provides interpretable and low-complexity clustering methods useful for anomaly detection while not requiring pre-defined clusters, however, DBSCAN can be quite sensitive to density and parameter selections. Overall, no single method dominates all criteria; deep learning suits high-dimensional, unstructured data, while Random Forest and XGBoost balance interpretability and scalability for structured data.

2. Predictive Maintenance Methods

Different predictive maintenance methods vary in accuracy, resource needs, and value trade-offs. A traditional statistical method, like an ARMA model, is relatively easy (days–weeks) and fast to implement, achieves medium accuracy (75%–85%), and has limited cost savings when put into practice. Machine Learning (ML) methods, on the other hand, achieve medium-high accuracy (85%–95%) and potential cost savings of 25%–40% by using multi-modal data, but require a longer timeframe to implement, larger data sets, and feature engineering [2]. Deep Learning (DL) methods yield the greatest accuracy (90%–98%) and would also create the highest possible cost savings (35%–50%) in large industrial settings, however they require high quality data, extensive compute, longer implementation times, and often have reduced or lacking interpretability [20]. The trade-off is that statistical methods are faster and easy to interpret, ML methods are better than statistical methods and make some trade-offs on interpretability and resource use, and DL methods give the best predictive quality but worse resource use and interpretability. Traditional signal analyses still have a place for use with monitoring vibrations and acoustics because of their ease of use and speed of implementation, but they cannot provide the predictive capabilities of ML or DL[29].

3. Explainable AI (XAI) Methods

Popular Explainable AI (XAI) methods differ in scope, computational cost, and reliability, each with trade-offs. LIME [21] offers local approximations of black-box models with low computational cost but shows medium reliability due to instability across runs. SHAP [55] extends this approach for local and global interpretability using cooperative game theory principles; it is more computationally expensive but highly reliable and widely used in sensitive domains like finance and healthcare. Attention-based mechanisms provide global interpretability within deep learning models, balancing efficiency and usability, though reliability depends on model complexity [56]. Rule-based methods provide readily interpretable explanations across the globe, especially through decision trees or expert systems; however, they typically apply to simpler and narrower domains [57]. Gradient-based methods produce local explanations for neural networks in an efficient manner, though they may lack reliability [58]. Ultimately, the choice of an XAI method will require a trade-off between interpretability, generalizability, and computational expense of the method.

4. Multi-modal and Fusion Methods

Early fusion methods involve filtering and fusing together several raw or low-level features into a single representation prior to model training. These methods maintain maximum amounts of available information with low computational cost; however, they can also remain sensitive to noise or modality failures. Early fusion methods provide simplicity and efficiency of the applied processing when compared to separate models. Applicability of early fusion approaches range, among others, to the domains of audio-visual emotion recognition [59] or speech processing [59]. Late fusion is a method combining the output from separate models trained on pre-extracted features. This method ensures reliability through cross-validation, and offers moderate computational requirements, at the expense of lower amounts of raw information. In education the late fusion approach has been applied to multimedia retrieval [60] and healthcare diagnostics [60]. Hybrid fusion employs both a filtering step and a late fusion filtered decision from the individually trained models. Hybrid approaches afford the ability to maintain higher amounts of available information but at a greater computational cost, and require greater synchronization in analysis methods. While more expensive to implement, hybrid analysis methods can have utility in complicated domains like autonomous driving or medical imaging [61].

DISCUSSION

Decision-making and predictive analytics can involve many different methods, each with their own strengths and trade-offs depending on the goals of the analysis. The use of traditional statistical models such as ARMA lends itself to quick implementation, includes an easier use, and offers the stated results above 50% accuracy. They can be good for low data solutions [11], [29]. Compared to traditional statistical methods, machine learning (ML) leverages the use of multi-modal data to achieve a more significant accuracy potential. Within the family of ML models, different considerations such as feature engineering and large datasets may hinder performance outcomes. Deep learning (DL) attain a high level of performance, with increased accuracy, efficiency, and cost savings at scale; DL models may sacrifice speed, resource intensity, and/or interpretability along the way. Unlike deep learning, CNNs leverage a significant portion of their power from the identification of spatial features (with diagnostic information encoded).And clustering methods of unsupervised learning, such as DBSCAN, remain relevant and applicable while offering strong interpretability [62], [53], [54], [2], [20]. In a complementary way, edge computer gives rise to new avenues of real-time operational analytics through model compression and federated (edge) learning frameworks [63], [64], [65]. Quantum computing also offers promise of application within the predictive maintenance framework through their capabilities related to optimization or advanced ML methods [66], [67]. To manage related issues of transparency, many XAI (explainable AI) methods perform analysis through additives from or in concurrence to black box methods. These models vary in levels of interpretability, cost, and reliability; with examples such as LIME, SHAP, attention mechanisms, and rule-based systems [21], [55], [56], [57].

In the end, multimodal fusion methods combine different data sources: early fusion offers high information and low robustness, late fusion improves reliability of decision but with moderate complexity, and hybrid fusion gains the best of both methods at a high computation cost especially within autonomous vehicles or medical imaging [59], [60], [61].

CONCLUSION AND FUTURE RESEARCH DIRECTIONS

Predictive maintenance (PdM) has evolved from simple time-series models to advanced AI systems capable of learning complex patterns from multimodal industrial data. Although traditional statistical methods retain their appeal owing to their simplicity and interpretability compared to ML and DL, which, although typically yield superior prediction accuracy in exchange for a reduced interpretability and greater computational cost. Effective explainable AI and causal discovery contribute significantly to enhancing the trustworthiness associated with making consequently significant PdM interventions. The rapid developments in multimodal fusion, edge and federated learning, and quantum optimization are anticipated to contribute further to enhancements in scalability, latency, and adaptability thereby constituting intelligent, explicable, and resilient Industry 4.0 maintenance ecosystems.

The widespread adoption of IoT has increased the interest in edge analytics, offering the potential to minimize latency, afford savings in the use of bandwidth and will drive improvements in reliability. Edge AI enables inference to be undertaken in real-time through the use of lightweight models accompanied by techniques such as pruning and knowledge distillation. Federated learning enables a collaborative approach to model training between different distributed sites, while protecting owner privacy. Quantum computing, including quantum machine learning and optimization, may further enhance pattern recognition, scheduling, and resource allocation in complex industrial systems, complementing edge and distributed AI approaches for next-generation predictive maintenance.

References

1. T. Zonta, C. A. da Costa, R. da Rosa Righi, M. J. de Lima, E. S. da Trindade, and G. P. Li, "Predictive maintenance in the Industry 4.0: A systematic literature review," Computers & Industrial Engineering, vol. 150, p. 106889, 2020.
2. T. P. Carvalho, F. A. Soares, R. Vita, R. D. P. Francisco, J. P. Basto, and S. Alcalá, "A systematic literature review of machine learning methods applied to predictive maintenance," Computers & Industrial Engineering, vol. 137, p. 106024, 2019.
3. R. K. Mobley, An Introduction to Predictive Maintenance, 2nd ed. Oxford, UK: Butterworth-Heinemann, 2002.
4. Y. Ran, X. Zhou, P. Lin, Y. Wen, and R. Deng, "A survey of predictive maintenance: Systems, purposes and approaches," arXiv preprint arXiv:1912.07383, 2019.
5. H. M. Hashemian and W. C. Bean, "State-of-the-art predictive maintenance techniques," IEEE Trans. Instrumentation and Measurement, vol. 60, no. 10, pp. 3480-3492, 2011.
6. Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, "Machinery health prognostics: A systematic review from data acquisition to RUL prediction," Mechanical Systems and Signal Processing, vol. 104, pp. 799-834, 2018.
7. K. Jardine, D. Lin, and D. Banjevic, "A review on machinery diagnostics and prognostics implementing condition-based maintenance," Mechanical Systems and Signal Processing, vol. 20, no. 7, pp. 1483-1510, 2006.
8. G. A. Susto, A. Schirru, S. Pampuri, S. McLoone, and A. Beghi, "Machine learning for predictive maintenance: A multiple classifier approach," IEEE Trans. Industrial Informatics, vol. 11, no. 3, pp. 812-820, 2015.
9. W. Zhang, C. Li, G. Peng, Y. Chen, and Z. Zhang, "A deep convolutional neural network with new training methods for bearing fault diagnosis under noisy environment and

different working load," Mechanical Systems and Signal Processing, vol. 100, pp. 439-453, 2019.

10. V. H. Tran, B. S. Yang, M. S. Oh, and A. C. Tan, "Machine condition prognosis based on regression trees and one-step-ahead prediction," Mechanical Systems and Signal Processing, vol. 25, no. 7, pp. 2849-2861, 2021.
11. M. Baptista, S. Sankararaman, I. P. de Medeiros, C. Nascimento Jr, H. Prendinger, and E. M. Henriques, "Forecasting fault events for predictive maintenance using data-driven techniques and ARMA modeling," Computers & Industrial Engineering, vol. 115, pp. 41-53, 2018.
12. J. Wang, Y. Ma, L. Zhang, R. X. Gao, and D. Wu, "Deep learning for smart manufacturing: Methods and applications," Journal of Manufacturing Systems, vol. 48, pp. 144-156, 2020.
13. L. D. Xu, E. L. Xu, and L. Li, "Industry 4.0: State of the art and future trends," Int. Journal of Production Research, vol. 56, no. 8, pp. 2941-2962, 2018.
14. R. Liu, B. Yang, E. Zio, and X. Chen, "Artificial intelligence for fault diagnosis of rotating machinery: A review," Mechanical Systems and Signal Processing, vol. 108, pp. 33-47, 2018.
15. H. Oliff and Y. Liu, "Towards industry 4.0 utilizing data-mining techniques: A case study on quality improvement," Procedia CIRP, vol. 63, pp. 167-172, 2017.
16. Z. Li, Y. Wang, and K. S. Wang, "Intelligent predictive maintenance for fault diagnosis and prognosis in machine centers: Industry 4.0 scenario," Advances in Manufacturing, vol. 5, no. 4, pp. 377-387, 2017.
17. Z. H. Zhou, J. Wu, and W. Tang, "Ensembling neural networks: Many could be better than all," Artificial Intelligence, vol. 137, no. 1-2, pp. 239-263, 2019.
18. S. Khan and T. Yairi, "A review on the application of deep learning in system health management," Mechanical Systems and Signal Processing, vol. 107, pp. 241-265, 2018.
19. R. Zhao, R. Yan, Z. Chen, K. Mao, P. Wang, and R. X. Gao, "Deep learning and its applications to machine health monitoring," Mechanical Systems and Signal Processing, vol. 115, pp. 213-237, 2019.
20. Z. Chen, W. Li, and D. Dong, "Deep learning-based rolling bearing fault diagnosis under variable conditions using LiConvNet, generalization and interpretability," Knowledge-Based Systems, vol. 188, p. 104948, 2020.

21. M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?: Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135-1144.
22. D. Gunning and D. W. Aha, "DARPA's explainable artificial intelligence program," AI Magazine, vol. 40, no. 2, pp. 44-58, 2019.
23. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE
24. M. Sharp, R. Ak, and T. Hedberg Jr, "A survey of the advancing use and development of machine learning in smart manufacturing," Journal of Manufacturing Systems, vol. 48, pp. 170-179, 2018.
25. R. E. Randall, Vibration-Based Condition Monitoring: Industrial, Aerospace and Automotive Applications. Chichester, UK: John Wiley & Sons, 2011.
26. Y. Lei, B. Yang, X. Jiang, F. Jia, N. Li, and A. K. Nandi, "Applications of machine learning to machine fault diagnosis: A review and roadmap," Mechanical Systems and Signal Processing, vol. 138, p. 106587, 2020.
27. X. S. Si, W. Wang, C. H. Hu, and D. H. Zhou, "Remaining useful life estimation – A review on the statistical data driven approaches," European Journal of Operational Research, vol. 213, no. 1, pp. 1-14, 2011.
28. G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, Time Series Analysis: Forecasting and Control, 5th ed. Hoboken, NJ: John Wiley & Sons, 2015.
29. Brandt, Noise and Vibration Analysis: Signal Analysis and Experimental Procedures. Chichester, UK: John Wiley & Sons, 2011.
30. F. J. Harris and A. G. Piersol, Harris' Shock and Vibration Handbook, 5th ed. New York: McGraw-Hill, 2002.
31. P. D. Welch, "The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," IEEE Trans. Audio and Electroacoustics, vol. 15, no. 2, pp. 70-73, 1967.
32. P. D. McFadden and J. D. Smith, "Model for the vibration produced by a single point defect in a rolling element bearing," Journal of Sound and Vibration, vol. 96, no. 1, pp. 69-82, 1984.
33. L. Cohen, Time-Frequency Analysis. Upper Saddle River, NJ: Prentice Hall, 1995.
34. D. Griffin and J. Lim, "Signal estimation from modified short-time Fourier transform," IEEE Trans. Acoustics, Speech, and Signal Processing, vol. 32, no. 2, pp. 236-243, 1984.

35. Daubechies, Ten Lectures on Wavelets. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1992.
36. S. Mallat, A Wavelet Tour of Signal Processing, 2nd ed. San Diego, CA: Academic Press, 1999.
37. N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, and H. H. Liu, "The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis," Proc. Royal Society of London A, vol. 454, no. 1971, pp. 903-995, 1998.
38. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, and F. Herrera, "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82-115, 2020.
39. W. Samek and K. R. Müller, "Towards explainable artificial intelligence," in Explainable AI: Interpreting, Explaining and Visualizing Deep Learning, W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K. R. Müller, Eds. Cham, Switzerland: Springer, 2019, pp. 5-22.
40. F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," arXiv preprint arXiv:1702.08608, 2017.
41. C. Molnar, Interpretable Machine Learning: A Guide for Making Black Box Models Explainable. Self-published, 2020.
42. Pearl, Causality: Models, Reasoning, and Inference, 2nd ed. Cambridge, UK: Cambridge University Press, 2009.
43. [J. Peters, D. Janzing, and B. Schölkopf, Elements of Causal Inference: Foundations and Learning Algorithms. Cambridge, MA: MIT Press, 2017.
44. P. Spirtes, C. Glymour, and R. Scheines, Causation, Prediction, and Search, 2nd ed. Cambridge, MA: MIT Press, 2000.
45. C. Glymour, K. Zhang, and P. Spirtes, "Review of causal discovery methods based on graphical models," Frontiers in Genetics, vol. 10, p. 524, 2019.
46. B. Schölkopf, "Causality for machine learning," arXiv preprint arXiv:1911.10500, 2019.
47. Kaddour, A. Lynch, Q. Liu, M. J. Kusner, and R. Silva, "Causal machine learning: A survey and open problems," arXiv preprint arXiv:2206.15475, 2022.
48. N. Hassanpour and R. Greiner, "Counterfactual reasoning for process optimization using structural causal models," in Lecture Notes in Computer Science, vol. 11956, 2019, pp. 91-106.

49. J. Runge, P. Nowack, M. Kretschmer, S. Flaxman, and D. Sejdinovic, "Detecting and quantifying causal associations in large nonlinear time series datasets," Science Advances, vol. 5, no. 11, p. eaau4996, 2019.
50. B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio, "Toward causal representation learning," Proc. IEEE, vol. 109, no. 5, pp. 612-634, 2021.
51. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5-32, 2001.
52. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785-794.
53. Z. Chen, K. Gryllias, and W. Li, "Mechanical fault diagnosis using convolutional neural networks and extreme learning machine," Mechanical Systems and Signal Processing, vol. 133, p. 106272, 2019.
54. Ester, H. P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in Proc. 2nd Int. Conf. Knowledge Discovery and Data Mining, 1996, pp. 226-231.
55. S. M. Lundberg and S. I. Lee, "A unified approach to interpreting model predictions," in Proc. 31st Int. Conf. Neural Information Processing Systems, 2017, pp. 4765-4774.
56. D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," arXiv preprint arXiv:1409.0473, 2015.
57. R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, "A survey of methods for explaining black box models," ACM Computing Surveys, vol. 51, no. 5, pp. 1-42, 2018.
58. K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," arXiv preprint arXiv:1312.6034, 2014.
59. P. K. Atrey, M. A. Hossain, A. El Saddik, and M. S. Kankanhalli, "Multimodal fusion for multimedia analysis: A survey," Multimedia Systems, vol. 16, no. 6, pp. 345-379, 2010.
60. T. Baltrušaitis, C. Ahuja, and L. P. Morency, "Multimodal machine learning: A survey and taxonomy," IEEE Trans. Pattern Analysis and Machine Intelligence, vol. 41, no. 2, pp. 423-443, 2019.
61. D. Xu, Y. Zhu, C. B. Choy, and L. Fei-Fei, "Scene graph generation by iterative message passing," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, 2017, pp. 5410-5419.