IOSR Journal of Business and Management (IOSR-JBM)
e-ISSN:2278-487X, p-ISSN: 2319-7668. Volume 27, Issue 12. Ser. 6 (December. 2025), PP 48-58
www.iosrjournals.org

Strategic Asset Management In Industry 4.0: Enhancing Fault Diagnosis Efficiency Through Statistical Feature Engineering And Machine Learning

Brehme D'napoli Reis De Mesquita[1], Roberto Célio Limão De Oliveira[2]
(Federal Institute Of Education, Science And Technology Of Maranhão, Campus Açailândia, Maranhão, Brazil)
(Graduate Program In Electrical Engineering, Federal University Of Pará, Campus Belém, Pará, Brazil)

Abstract:
Predictive Maintenance has become a foundational element of Industry 4.0, with the objective of enhancing asset availability and reducing operational expenditures. However, the widespread adoption of automated diagnostic systems is hindered by significant barriers, including the substantial computational demands of Deep Learning models and the opaque nature of algorithmic decisions ("black box") that can impede managerial confidence. This study proposes an efficient methodology based on classical Machine Learning and statistical signal processing to address these limitations. Vibration data from an experimental setup were utilized, simulating diverse operational conditions and unbalance fault types. The methodological approach centered on the extraction of time-domain statistical features (e.g., Root Mean Square, Kurtosis, Standard Deviation, and Crest Factor) to construct a robust feature vector. The performance of various algorithms, including Random Forest, Support Vector Machine, k-Nearest Neighbors, and Multilayer Perceptron, was evaluated through a comparative analysis. This evaluation employed cross-validation and F1-Score metrics to address the issue of class imbalance. The experimental results demonstrated that the Random Forest model, when utilizing a set of selected statistical features, exhibited superior performance in comparison to the alternative architectures, attaining a global F1-Score of approximately 92%. Visual analysis via t-Distributed Stochastic Neighbor Embedding confirmed high separability of fault classes within the latent feature space. The model demonstrated superior robustness in differentiating between normal and critical states, thereby minimizing both false positives and false negatives, without requiring computationally intensive graphics hardware. The conclusion drawn from this analysis indicates that feature engineering based on physically grounded statistics is more decisive for diagnostic accuracy than the architectural complexity of the classifier. The proposed approach validates a cost-effective pathway for the industry, enabling the implementation of high-precision diagnostics on edge devices (Edge Computing) or Programmable Logic Controllers (PLCs) with low computational power. This has the potential to democratize access to predictive maintenance, offering managers a reliable and interpretable tool for strategic decision-making.

Key Word: *Predictive Maintenance; Feature Engineering; Decision Support Systems; Operational Efficiency; Industry 4.0.*

Date of Submission: 13-12-2025 Date of Acceptance: 23-12-2025

I. Introduction

The concept of Industry 4.0 has evolved beyond the confines of mere marketing terminology. In the contemporary global context, it has become the fundamental foundation of international competition[1,2]. The integration of digital tools is not merely a matter of deployment; rather, it signifies the intricate convergence of physical systems with virtual reality. This paradigm shift necessitates a constant state of vigilance and optimization of operations for companies[3]. In our contemporary, highly interconnected environment, the capacity to extract valuable insights from vast quantities of raw data has evolved from a mere competitive advantage to an essential prerequisite for economic survival[4,5].

Historically, strategic asset management was regarded as a mere support cost. Recently, however, it has emerged as a pivotal element in ensuring business continuity[6]. Yet, within the context of the factory floor, managers find themselves entrapped in a costly predicament perpetuated by conventional maintenance strategies. On one hand, the "run-to-failure" approach (i.e., corrective maintenance) can be financially detrimental. The financial repercussions resulting from these events extend beyond the immediate costs of emergency repairs. They also encompass the production downtime, the loss of revenue, and the presence of genuine safety hazards[7].

The prevailing alternative, calendar-based preventive maintenance, offers only marginal financial advantages. The practice of replacing components solely based on the schedule, irrespective of their actual condition, results in the premature disposal of components with remaining useful life (RUL) and the misallocation of skilled labor[7]. This operational disconnect creates a significant management gap. The continued adherence to reactive or calendar-based models has become untenable in the contemporary landscape. The industry has necessitated a radical transition towards predictive strategies. The identification of faults before they escalate is crucial for the prevention of operational waste, ensuring that resources are allocated precisely to their intended destinations[8,9].

Demand for efficiency has brought Condition-Based Maintenance (CBM) to the forefront. Industrial sensors are currently producing an immense quantity of data, which is unparalleled in its scope. This influx is redefining the established paradigm, with agile, data-driven insights eclipsing the cumbersome, intricate traditional physical models[7,8,9].

The analysis of signals for the purpose of fault diagnosis can be categorized into three distinct domains: time, frequency, and time-frequency domains[10]. For rotating machinery, the Fast Fourier Transform (FFT) remains the standard approach for identifying fault signatures[11]. However, FFTs encounter challenges when confronted with the chaotic, non-stationary signals that are characteristic of a dynamic plant. While sophisticated methods, such as Wavelet Transforms (WT) or Empirical Mode Decomposition (EMD), have been developed to address the resolution problem, these methods require a substantial computational burden[12,13]. This barrier frequently precludes the implementation of continuous, large-scale monitoring[10,14]. Consequently, there has been a resurgence of "old-school" statistical time-domain indicators, such as root mean square (RMS), kurtosis, and skewness. These metrics have been shown to effectively remove noise and reduce data size while preserving the essential health signal[15,16,17].

The Machine Learning (ML) toolkit for maintenance has also expanded. Researchers have found evidence in the form of experimental results to support the notion that "shallow" algorithms like Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Random Forests (RF) achieve a pragmatic balance between accuracy and speed[18,19,20]. In recent developments, Deep Learning (DL) (e.g., convolutional neural networks (CNNs), recurrent neural networks (RNNs)) has emerged as a prominent approach, offering the promise of direct data processing, thereby relieving engineers of the task of manual feature creation[10,21].

However, this level of sophistication is often accompanied by a substantial financial investment. A significant limitation of DL models is their requirement for substantial labeled datasets, which are scarce in the manufacturing sector. These models also necessitate expensive hardware, such as graphics processing units (GPUs), exacerbating the issue. In the context of projects emphasizing a rapid return on investment (ROI) and minimal operational expenditures (OPEX), these specifications frequently serve as a pivotal determinant, often dictating the overall success or failure of the project[22,23,24].

The transition from the abstract realm of theory to the practical domain of real-world application often presents a series of challenges. The process of handling high-frequency raw data creates a bottleneck, resulting in bandwidth constraints and a significant demand on available processing resources. For factories operating on limited IT budgets or facing strict latency limits, scaling these heavy ML solutions is not a viable option[25,26,27].

Additionally, the interpretability challenge, often referred to as the "black-box" problem, is a significant concern[28]. While a model may exhibit high degrees of accuracy, the concealment of its underlying logic can engender a state of suspicion. It is rare for managers and engineers to authorize a machine shutdown based on an opaque diagnosis that cannot be physically verified[29,30]. The industry has consequently identified a need for a middle ground: methods that effectively balance high precision with transparency and computational lightness[23,30].

The economic viability of a solution is jeopardized if it requires high-performance computing clusters or lacks interpretability. Any methodology that fails to satisfy these operational constraints is likely to face rejection on economic grounds. The deployment of opaque models can compromise the financial viability of the investment and the anticipated return thereof. Consequently, resolving this dilemma necessitates a strategic compromise between precision and applicability. The tools in question must possess adequate precision to avert system failures while also maintaining the simplicity required to facilitate integration with existing legacy infrastructure. Ensuring widespread accessibility of predictive technology demands finding this equilibrium[23,25,31].

The present paper proposes and validates a strategic fault diagnosis model aimed at maximizing operational efficiency. The proposed methodology employs a combination of Statistical Feature Engineering and Machine Learning algorithms. Rather than pursuing raw complexity, our emphasis is placed on the informational quality of the extracted features. This research conducts an experimental case study to demonstrate the efficacy of a smart statistical treatment in delivering reliable diagnostics. By comparing classical classifiers (specifically RF, SVM, k-NN, and Multilayer Perceptron (MLP)) fed by temporal indicators, the results indicate that this approach delivers robust diagnostics without the computational overhead associated with other methods, offering managers a viable tool for Industry 4.0.

II. Material And Methods

This section delineates the experimental methodology employed for the development and validation of the proposed diagnostic strategy. The approach was engineered to translate physical vibration signals into automated managerial decisions, guaranteeing both low computational cost and high reliability. To ensure the study's reproducibility and the logical comprehension of the system, the workflow was organized into five distinct phases: data acquisition, signal pre-processing, feature engineering, experimental design, and computational modeling.

Figure 1 provides a schematic representation of the developed methodological framework. The diagram delineates the complete data pipeline, illustrating how information flows from physical capture at the cooler fan (left) to the performance evaluation of the machine learning algorithms (right).

Figure 1: Proposed experimental framework and data processing pipeline. The workflow highlights the transition from raw signal acquisition to statistical feature extraction and machine learning classification.

Dataset Description and Experimental Test Rig

The dataset employed in this investigation was obtained from the UCI Machine Learning Repository[32], which was originally developed for the purpose of predicting engine faults through vibration analysis. The experimental mechanism was designed to simulate the vibrational behavior of industrial rotating machinery under varying operating conditions and load imbalances.

The experimental test rig comprised a computer cooling fan, model Akasa AK-FN059 12 cm Viper, serving as the primary rotary actuator. For the purpose of acquiring vibration signals, a digital accelerometer (model MMA8452Q) was affixed directly to the fan structure. The sensor in question features a 12-bit resolution and utilizes the I2C (Inter-Integrated Circuit) protocol to communicate with an Arduino UNO microcontroller. The latter is responsible for regulating rotational speed and the data extraction. The acquisition system was orchestrated by custom software developed in the Processing language, which captured and stored the data transmitted via the serial port.

In order to emulate realistic degradation scenarios and failure due to unbalance, physical modifications were introduced to the fan blades. To induce weight differentials, small magnets were attached to the blades, thereby generating controlled vibrations during rotation. The experimental protocol defined three discrete blade weight distribution configurations, meticulously designed to encompass varying levels and patterns of vibrational severity. For each of these configurations, the motor was operated at 17 distinct rotational speeds, ranging from 20% to 100% of the cooler's maximum capacity, in 5% increments.

The data acquisition protocol was configured to operate with a sampling interval of 20 milliseconds (equivalent to a sampling frequency of 50 Hertz) over a duration of one minute for each configured speed. This process yielded a total of 3,000 temporal records for each speed. The raw dataset, which is the sum of all the vibration instances, contains 153,000 measurements across the orthogonal axes (x, y, z), rotational speed, and operating condition. This dataset providing a solid foundation for the training and validation of the proposed machine learning models.

Pre-processing and Feature Engineering

Given the inherently stochastic and noisy nature of raw vibration signals, the direct application of time series data in classification models risks compromising generalization due to the "curse of dimensionality"[33,34]. To circumvent this issue and extract latent patterns associated with machine condition, a time-domain Statistical Feature Engineering approach was adopted[14,35,36]. This strategy has been demonstrated to reduce data complexity and enhance model generalization capabilities, while also fostering the interpretability of results. Such transparency is a critical aspect for the reliability of decision-making in industrial environments.

The initial processing stage involved signal windowing. The data stream was segmented into non-overlapping windows comprising 50 samples. Given the 50-hertz sampling frequency, each window corresponds precisely to one second of asset operation. This granularity was strategically established to enable low-latency diagnostics, facilitating near real-time decision-making[14,36,37]. Prior to feature extraction, a linear detrending technique was applied to each window to eliminate DC offsets or signal drifts unrelated to actual mechanical faults. In the field of time-series analysis, detrending is widely recognized as a critical pre-processing step[38]. Its incorporation into vibration signal analysis has been shown to enhance system robustness and reduce false positive detections in anomaly detection systems[39]. This initial treatment ensures that signals are standardized and representative, thereby establishing a solid foundation for statistical feature extraction.

For each processed window across the three spatial axes (x, y, z), a statistical feature vector was computed. This vector comprised six key indicators. The selection of these indicators was based on their demonstrated capacity to characterize the physical characteristics associated with mass imbalance and the intensity of vibrations[14,36,40]:

- Root Mean Square (RMS): This index is indicative of the total energetic content of the signal. In the context of unbalance, the RMS is directly proportional to the magnitude of the centrifugal forces generated by mass asymmetry in the rotor.
- Standard Deviation: The metric quantifies the dispersion of data relative to the mean. It serves to indicate the amplitude of dynamic oscillation around the equilibrium point.
- Kurtosis: The fourth statistical moment is fundamental for analyzing the shape of the signal distribution. While pure unbalance tends to generate sinusoidal signals, Kurtosis allows for the identification of waveform distortions (flattening) caused by system non-linearities or sensor saturation at critical vibration levels.
- Peak Value: The maximum absolute amplitude of the signal is a critical metric for identifying vibration excursion limits and potential structural impacts.
- Mean Absolute Value (MAV): The indicator provides a robust measure of the average amplitude of the rectified signal is robust. There is a correlation between this measure and the overall severity of the fault.
- Crest Factor: Defined as the ratio between the peak value and the RMS value. This dimensionless metric is effective for monitoring changes in vibration patterns, indicating how "impulsive" or saturated the signal is relative to its average energy.

This transformation process yielded a structured matrix in which each instance represents the machine's state over a one-second interval. These instances are characterized by low-dimensional feature vectors that are primed for ingestion by the machine learning algorithms.

Model Configuration and Validation Strategy

For the classification and diagnostic phase, the study employed the Python programming language along with the Scikit-Learn[41] and SciPy[42] libraries. Four classical supervised learning algorithms were implemented and selected for their computational efficiency relative to deep neural networks. The following algorithms are considered: Random Forest[43], Support Vector Machines[44], k-Nearest Neighbors[45], and Multilayer Perceptron[46].

The performance evaluation process was meticulously executed in accordance with a two-phase approach, characterized by its rigor and systematic nature. In the initial phase, a benchmark was established using default hyperparameters. In the subsequent phase, the GridSearchCV technique was employed for the purpose of hyperparameter optimization, with the objective of maximizing the F1-Macro metric. This metric offers a balanced trade-off between precision and recall, making it ideal for multi-class problems[47].

To ensure the statistical reliability of results and mitigate overfitting, stratified k-fold cross-validation with $k = 5$ was employed. This approach ensures that the proportion of failure classes remains constant across all training and testing subsets. As a result, it provides a robust estimate of the models' generalization capability on unseen data[48]. This methodological configuration ensures that models are evaluated fairly and reliably, thereby enhancing their applicability in real industrial environments.

Experimental Design and Scenario Definition

In order to isolate the influence of different signal components and assess the impact of measurement artifacts on diagnostic accuracy, the experiment was structured into four main analytical scenarios (C1 to C4), including variations for handling sensor saturation (clipping).

The rationale for defining these scenarios is rooted in the electromechanical characteristics of the test rig[32]. It was observed that the accelerometer's mounting orientation subjects the x-axis to constant gravitational acceleration, thereby introducing a DC offset (zero-frequency component) that has the potential to mask high-frequency oscillatory components characteristic of mechanical faults. Furthermore, under conditions of high rotation or severe unbalance, vibration amplitude may exceed the sensor's dynamic range (saturation), thereby inducing the clipping phenomenon, where signal peaks are truncated[49,50,51].

To investigate these variables, the scenarios were designed as follows:

- Scenario C1 (Baseline Triaxial): The complete feature vector (18 features) is extracted from the three axes (x, y, z). In this scenario, detrending is employed to eliminate the DC level, thereby testing the hypothesis that combined information from three spatially orthogonal axes maximizes class separation, despite potential noise in the x-axis[14,38].
- Scenario C2 (Axis Selection – y/z): The analysis is restricted to the y and z axes (12 features), with the x-axis being deliberately discarded. This is a deliberate choice, as the x-axis is often neglected in correlated

studies due to its low dynamic informational content. The objective is to achieve more generalizable and less noisy models[32,36].

- Scenarios C3 and C4 (DC Component Analysis): In contrast to previous models, these scenarios do not utilize second-order or higher-order statistics. C3 employs the x-axis mean, whereas C4 utilizes the means of x, y and z. The objective was to investigate whether static displacement (or bias), which is often removed by detrending, carries residual predictive information regarding the machine state[49,51].

Furthermore, to address the non-linearity issue introduced by sensor saturation, sub-scenarios were created based on a clipping threshold that was experimentally defined at $7.8g$ (close to the sensor's physical limit of $\pm 8g$):

- Variation “a” (Saturation Filtering): The following time windows were excluded from the study: those in which any axis exceeded the $7.8g$ threshold (e.g., C1a, C2a). This approach sought to train models exclusively with "high-fidelity" data, thereby eliminating samples that were corrupted by signal truncation[50,51].
- Variation “b” (Saturation Flag Inclusion): In this variation, all windows were retained, but with the incorporation of an additional binary feature. This feature is designed to assume a value of 1 in the event of saturation, and 0 in the absence thereof (e.g., C1b, C2b). This approach tests the hypothesis that the saturation event itself is a severe failure indicator and should be incorporated into the model as information, rather than treated as noise to be discarded[23,50,52,53].

This experimental matrix, composed of combinations of axes and signal treatment strategies, allows for not only finding the most accurate model but also understanding the physics of the monitoring problem.

Computational Modeling and Performance Metrics

The classification stage was implemented using the Scikit-Learn library within a Python environment[41]. To ensure the complete reproducibility of the experiments and to eliminate stochastic variations in weight initialization (as in MLP) or data sampling (as in RF and cross-validation), a global random state (seed) was fixed at 42 throughout all process stages.

Concurrently, the objective was to cultivate a robust comparative analysis that considers disparate learning paradigms and decision boundaries. To this end, four algorithms with distinct inductive biases were selected:

- k-Nearest Neighbors (k-NN): The algorithm under consideration is of an instance-based nature, designed to classify samples based on proximity in the feature space. This serves as a simple, non-parametric baseline[45].
- Support Vector Machines (SVM): The selection of this method is primarily attributed to its demonstrated efficacy in high-dimensional spaces, as well as its capacity to optimize the separation margin between classes. This property is of paramount importance for the discernment of subtle vibration states[44].
- Random Forest (RF): This ensemble method based on decision trees was selected for its resistance to overfitting and its ability to handle non-linear relationships between statistical variables[43].
- Multilayer Perceptron (MLP): A classical feedforward artificial neural network is employed to evaluate whether the universal approximation capability of non-linear functions offers a performance gain over traditional methods[46].

In order to ensure that the results reflected the maximum potential of each algorithm rather than merely default configurations, a GridSearchCV strategy coupled with stratified cross-validation ($k = 5$) was adopted[41,48].

With regard to the evaluation criteria, although global accuracy (the ratio of total correct predictions) is reported for the sake of comparability with the extant literature, the primary decision metric adopted in this study was the F1-Score (Macro). This approach is substantiated by the critical nature of fault diagnosis.

In industrial environments, accuracy can be deceptive, particularly in the presence of class imbalance. The F1-Score, defined as the harmonic mean of precision and recall, imposes a stringent penalty on "false diagnoses"[47]. In the context of maintenance, this is vital to avoid two risk scenarios that have been widely discussed in the machinery prognosis literature[36]:

1. False Negative (Type II Error): The model's inability to detect a genuine defect may result in catastrophic failures and safety hazards.
2. False Positive (Type I Error): The generation of false alarms by the model causes unnecessary downtime and increased operational costs, thereby reducing system efficiency.

Therefore, the objective is to maximize the F1-Score in order to obtain a model that is conservative enough to avoid spurious alarms, yet sensitive enough not to overlook incipient faults.

III. Results And Analysis

This section presents and discusses the quantitative results derived from the classification experiments. The evaluation framework was designed to assess not merely the final accuracy of the models, but also the robustness of the models against variations in data architecture (Scenarios C1 and C4) and algorithm parameterization. The analysis is grounded in the F1-Score (Macro) metric, selected for its ability to penalize both false positives and false negatives in imbalanced classes, as well as Global Accuracy.

The ensuing tables offer a synopsis of the performance. Table 1 presents the initial benchmark utilizing default hyperparameters, while Table 2 displays the results following exhaustive optimization via GridSearchCV.

Table 1: Baseline classifier performance evaluation. This table summarizes the diagnostic metrics (F1-Macro and Accuracy) obtained using default hyperparameters, serving as a reference benchmark to assess the intrinsic stability and "out-of-the-box" capability of each algorithm across the four experimental scenarios.

Scenario	RF	SVM	k-NN	MLP
C1	0.917992 / 0.918301	0.716645 / 0.725163	0.754999 / 0.755229	0.794167 / 0.795425
C1a	0.917847 / 0.918675	0.716287 / 0.724573	0.757479 / 0.758026	0.796977 / 0.798187
C1b	0.917314 / 0.917647	0.714010 / 0.722876	0.753341 / 0.753595	0.790079 / 0.791176
C2	0.874943 / 0.875490	0.699724 / 0.709150	0.724175 / 0.724510	0.738393 / 0.742157
C2a	0.875264 / 0.876512	0.696981 / 0.706832	0.718709 / 0.719208	0.748472 / 0.752014
C2b	0.875261 / 0.875817	0.697901 / 0.707516	0.723871 / 0.724183	0.740533 / 0.744444
C3	0.398979 / 0.400000	0.379325 / 0.453922	0.396704 / 0.395752	0.418143 / 0.466993
C3a	0.403736 / 0.406962	0.363111 / 0.454485	0.402316 / 0.402282	0.388090 / 0.466199
C3b	0.408540 / 0.409477	0.380611 / 0.456209	0.401649 / 0.401634	0.407573 / 0.467647
C4	0.647013 / 0.648039	0.539851 / 0.565033	0.642660 / 0.643464	0.636215 / 0.640523
C4a	0.646179 / 0.647595	0.548291 / 0.571293	0.630675 / 0.631198	0.637325 / 0.641898
C4b	0.648103 / 0.649020	0.518646 / 0.549673	0.645566 / 0.646405	0.643837 / 0.647059

Table 2: Optimized classifier performance following hyperparameter tuning. This table presents the maximum diagnostic metrics achieved after exhaustive GridSearchCV, illustrating the potential gain in predictive power and the robustness of the models when subjected to fine-tuning strategies.

Scenario	RF	SVM	k-NN	MLP
C1	0.921292 / 0.921569	0.813539 / 0.815359	0.754999 / 0.755229	0.794167 / 0.795425
C1a	0.920176 / 0.921021	0.820917 / 0.822955	0.765821 / 0.767066	0.796977 / 0.798187
C1b	0.918366 / 0.918627	0.809770 / 0.811765	0.753341 / 0.753595	0.790079 / 0.791176
C2	0.876237 / 0.876797	0.752067 / 0.756863	0.733547 / 0.734314	0.738393 / 0.742157
C2a	0.876266 / 0.877516	0.745112 / 0.751012	0.724585 / 0.725901	0.748472 / 0.752014
C2b	0.875261 / 0.875817	0.743807 / 0.749673	0.734978 / 0.735621	0.740533 / 0.744444
C3	0.415326 / 0.436928	0.393308 / 0.458824	0.411639 / 0.412745	0.418143 / 0.466993
C3a	0.419895 / 0.448461	0.376986 / 0.453822	0.426823 / 0.430059	0.388202 / 0.466868
C3b	0.426989 / 0.452288	0.393876 / 0.465033	0.406758 / 0.406536	0.407573 / 0.467647
C4	0.658402 / 0.660784	0.618788 / 0.624837	0.653820 / 0.654902	0.639090 / 0.643791
C4a	0.654430 / 0.656964	0.618858 / 0.625494	0.634774 / 0.635879	0.645559 / 0.649930
C4b	0.660522 / 0.662418	0.614977 / 0.621569	0.655265 / 0.656536	0.643837 / 0.647059

Learning Dynamics: RF Robustness vs SVM Sensitivity

To facilitate the visualization of performance distribution and model evolution, Figures 2 and 3 present the performance heatmaps.

A comparative analysis reveals distinct machine learning paradigms. The RF exhibited remarkable stability. The performance variation between the default version (Table 1) and the optimized version (Table 2) in scenario C1 was negligible ($\Delta < 0{,}4\%$). From a theoretical standpoint, this occurrence can be attributed to the utilization of the Bagging (Bootstrap Aggregating) principle by RF, a technique that aims to minimize model variance by integrating multiple uncorrelated decision trees. This architecture enables the model to inherently manage the non-linearity of vibration data without necessitating intricate fine-tuning, thereby validating its suitability for expeditious industrial implementation ("plug-and-play")[43,54,55].

Conversely, the SVM proved to be highly sensitive to the definition of the separating hyperplane. In the default configuration, the SVM failed to capture complex decision boundaries ($F1 \approx 0.71$). However, the optimization of parameters C (error penalty) and γ (RBF kernel coefficient) allowed the model to adjust the separation margin, elevating the F1-Score to 0.8135. This indicates that while the SVM possesses the mathematical capability for diagnosis, its practical application is contingent upon the prior computational cost of grid searching[44,56,57].

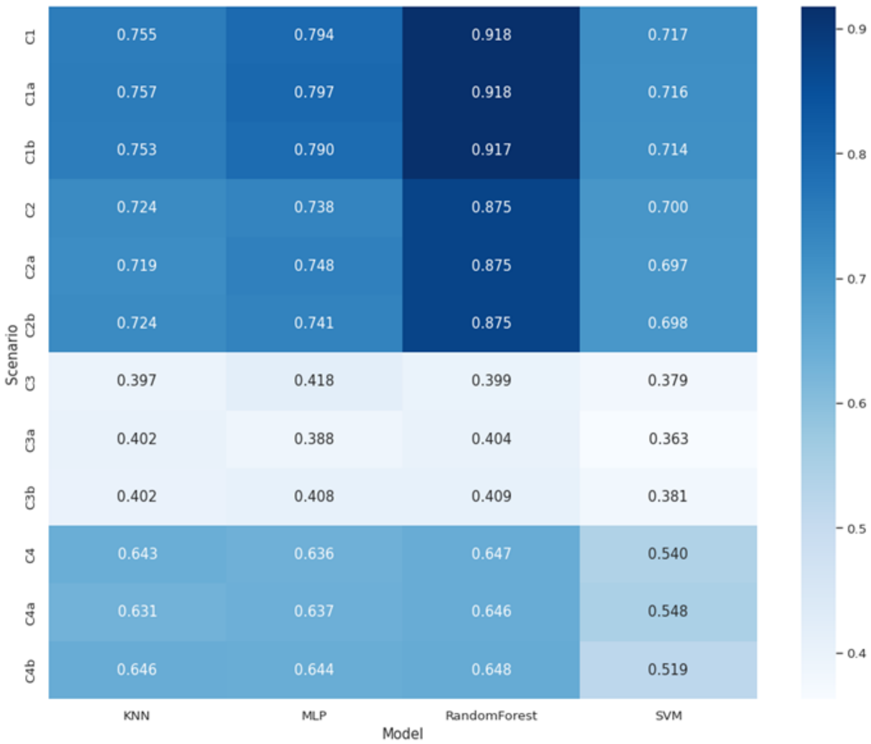

Figure 2: Heatmap visualization of baseline classifier performance using default hyperparameters. The color gradient represents the F1-Macro score, where darker shades indicate superior diagnostic capability, highlighting the dominance of RF across time-domain scenarios (C1-C2).

Figure 3: Heatmap visualization of classifier performance using optimized hyperparameters via GridSearchCV. The color gradient represents the F1-Macro score, where darker shades indicate superior diagnostic capability. The figure highlights the significant improvement in SVM performance compared to the baseline (Figure 2) and the continued dominance of RF across key scenarios.

Physics of Vibration and the Relevance of the Gravitational Axis (x)

A seminal contribution of this study is the refutation of the hypothesis that the vertical axis (x) of MEMS accelerometers should be discarded due to the gravitational offset ($1g$). This hypothesis was based on prior works that utilized the same dataset[32], in which the x-axis was not considered in the analyses.

A comparison of the optimized results is presented below:

- Scenario C1 (Triaxial x, y, z): The RF F1 score reached approximately 0.921.
- Scenario C2 (Biaxial y, z): The RF F1 score reached approximately 0.876.

The physical phenomenon of rotor mass asymmetry (unbalance) induces a rotating centrifugal force vector that excites the structure in all spatial directions, not merely in the horizontal plane. The performance degradation observed upon removing the x-axis (Scenario C2) demonstrates the efficacy of the pre-processing stage, particularly the detrending process, in eliminating the static gravity component (DC) while isolating the dynamic vibration component (AC)[14,38].

Consequently, the x-axis did not introduce noise; rather, it provided discriminatory information regarding the spatial projection of the unbalance force, thereby validating that low-cost MEMS sensors can be utilized in a full triaxial configuration[36].

Feature Engineering and Data Topology (t-SNE Analysis)

The comparison between scenarios based on higher-order statistics (C1) and those based solely on signal mean (C3/C4) underscores the nature of the fault signal. Scenarios C3 and C4 demonstrated suboptimal performance ($F1 < 0.66$), thereby corroborating the hypothesis that the fault "signature" resides within the energy and dispersion (RMS and Standard Deviation) and the distribution shape (Kurtosis), rather than in the mean displacement, which tends toward zero in oscillatory systems[14,36,40].

To graphically explore this separability, Figure 4 presents the t-SNE (t-Distributed Stochastic Neighbor Embedding) visualization[58], which projects the multidimensional complexity of the data onto a 2D plane.

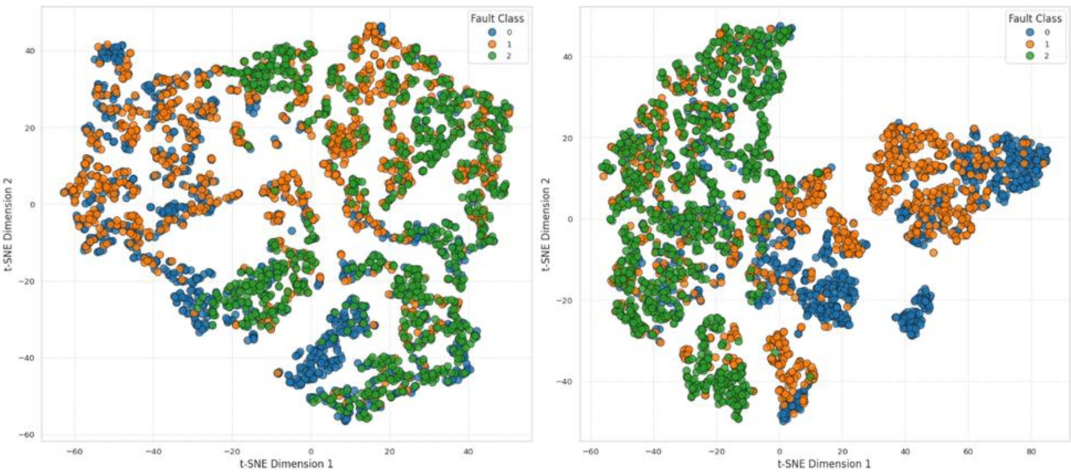

Figure 4: t-SNE visualization of the feature space topoly. (a) In Scenario C4 (Mean-based), the classes overlap significantly, creating a chaotic cluster that explains the classifiers' inability to exceed 65% accuracy. (b) In Scenario C1 (Statistical Features), clear boundaries and distinct clusters emerge for each fault condition, visually validating why the Random Forest model achieved $> 92\%$ F1-Score.

IV. Managerial Implications

The experimental results obtained in this study extend beyond the mere technical validation of algorithms, offering direct strategic repercussions for asset management within the context of Industry 4.0.

The initial major implication pertains to risk mitigation and operational reliability. The F1-Score of approximately 92% achieved by the RF model translates, in managerial terms, into a high degree of confidence for decision-making. In a production environment, a predictive system must balance two critical risks: the False Negative risk (failure to detect a defect), which precipitates catastrophic failure and unplanned downtime, and the False Positive risk (false alarm), which generates unnecessary maintenance stoppages and wastes specialized labor. The elevated F1-Macro score attained signifies that the proposed system concurrently minimizes errors, as previously discussed in machinery prognosis reviews[36]. For the plant manager, this signifies the capability to transition from a reactive stance to a planned approach. In lieu of suspending operations in the face of an emergency scenario, such as a sudden equipment malfunction, the manager has the option to schedule the intervention within a designated time frame, such as the interval between shifts. This approach is predicated on the principle of ensuring operational continuity and optimizing Overall Equipment Effectiveness (OEE)[59,60].

Secondly, this study validates a cost-effective pathway for the democratization of predictive maintenance. A persistent obstacle to the integration of artificial intelligence within industrial contexts pertains to the prevailing perception that its implementation demands advanced hardware, such as graphics processing units (GPUs), and intricate cloud infrastructure, which is commonly associated with DL models. The present study refutes the need for exorbitant computational resources for this class of problem by demonstrating that a RF algorithm fed by statistical attributes achieves competitive performance on tabular data[61,62]. This suggests that the proposed solution can be integrated into low-cost industrial microcomputers (Edge Computing) or existing Programmable Logic Controllers (PLCs), enabling the retrofitting of legacy machinery without substantial Capital Expenditure (CAPEX). This lightweight architecture enables Small and Medium-sized Enterprises (SMEs) to implement monitoring systems that operate locally, reducing bandwidth costs and latency associated with cloud transmission, ultimately ensuring that diagnostics are available on the shop floor in real-time[25,63,64].

The transition from raw data to statistical diagnosis alters the dynamics of information visualization and shop floor empowerment. From a User Experience (UX) perspective, the proposed system functions as a translator, converting complex vibratory signals—which are unintelligible to non-specialists—into a clear diagnostic dashboard. Rather than presenting abstract error codes, the system provides a "health status" grounded in robust physical statistics (e.g., "High Probability of Unbalance due to Kurtosis deviation"). This enhanced transparency serves to reduce the cognitive load experienced by operators and maintenance technicians, thereby enabling them to validate the diagnostic findings of artificial intelligence by comparing them to their own experiential knowledge. Consequently, the tool ceases to be a "black box" and becomes a trusted decision support partner, fostering a data-driven culture where maintenance actions are justified by objective evidence rather than solely by intuition or rigid schedules[28,65].

V. Conclusion

This work presented a robust approach for fault diagnosis in rotating machinery by integrating time-domain statistical signal processing with ML algorithms. The comparative analysis demonstrated that the RF model, fed by a vector of meticulously selected statistical features, outperformed alternative architectures, achieving an F1-Score of approximately 92%. The principal scientific contribution of this study is the empirical demonstration that the quality of Feature Engineering frequently plays a more decisive role in diagnostic accuracy than the architectural complexity of the classifier itself. In summary, the study demonstrated that statistically well-founded classical models can exhibit performance that is either superior or equivalent to "black-box" approaches, with the added advantage of interpretability.

In terms of industrial impact, the proposed methodology offers a viable and efficient pathway for the implementation of Industry 4.0. By validating a solution that demands low computational cost, thereby obviating the need for computationally intensive hardware such as GPUs, this study significantly lowers the entry barriers for predictive maintenance. This facilitates the implementation of real-time asset monitoring for enterprises of varying scales, whether utilizing Edge Computing infrastructure or existing PLCs. The capacity to translate raw vibration data into actionable diagnostics enhances managerial decision-making, driving the transition from unplanned corrective downtime to proactive and economically sustainable interventions.

Despite the promising results, it is necessary to acknowledge the limitations inherent in the use of controlled experimental data, where exogenous variables—such as extreme thermal variations and industrial electromagnetic noise—are mitigated. As a distinct avenue for future research, the deployment of this method in a real-world shop floor environment (in-situ) is suggested, thereby exposing the algorithm to noisy and unstructured operational conditions. Furthermore, future investigations may explore sensor fusion (e.g., combining vibration and electrical current data) to enhance diagnostic robustness against the uncertainties of the manufacturing environment.

References

[1]. Piccarozzi M, Aquilani B, Gatti C. Industry 4.0 In Management Studies: A Systematic Literature Review. Sustainability. 2018; 10 (10): 3821. DOI: Https://Doi.Org/10.3390/Su10103821

[2]. Santos BP, Alberto A, Lima TDFM, Charrua-Santos FMB. Indústria 4.0: Desafios E Oportunidades. Revista Produção E Desenvolvimento. 2018;4(1):111–24. DOI: Https://Doi.Org/10.32358/Rpd.2018.V4.316

[3]. Tortorella GL, Cauchick-Miguel PA, Li W, Staines J, Mcfarlane D. What Does Operational Excellence Mean In The Fourth Industrial Revolution Era? International Journal Of Production Research. 2022;60(9):2901–17. DOI: Https://Doi.Org/10.1080/00207543.2021.1905903

[4]. Romero CAT, Ortiz JH, Khalaf OI, Prado AR. Business Intelligence: Business Evolution After Industry 4.0. Sustainability. 2021;13(18):10026. DOI: Https://Doi.Org/10.3390/Su131810026

[5]. Janík S, Szabó P, MÍkva M, Mareček-Kolibiský M. Effective Data Utilization In The Context Of Industry 4.0 Technology Integration. Applied Sciences. 2022;12(20):10517. DOI: Https://Doi.Org/10.3390/App122010517

[6]. Oyekan M, Jinadu SO, Enyejo JO. The Role Of Strategic Asset Management In Accelerating The Energy Transition. International Journal Of Innovative Science And Research Technology. 2025;10(9):1405–21. DOI: Https://Doi.Org/10.38124/Ijisrt/25sep792

[7]. Gholipour Y, Zare M, Sereshk MV, Gholipour Y. A Comprehensive Review Of Maintenance Strategies: From Reactive To Proactive Approaches. Central Asia And The Caucasus. 2025;26(1):70–83. DOI: Https://Dx.Doi.Org/10.2139/Ssrn.5349871

[8]. Thomas DS, Weiss B. Maintenance Costs And Advanced Maintenance Techniques: Survey And Analysis. International Journal Of Prognostics And Health Management. 2021;12(1). DOI: Https://Doi.Org/10.36001/Ijphm.2021.V12i1.2883

[9]. Hector I, Panjanathan R. Predictive Maintenance In Industry 4.0: A Survey Of Planning Models And Machine Learning Techniques. Peerj Computer Science. 2024;10:E2016. DOI: Https://Doi.Org/10.7717/Peerj-Cs.2016

[10]. Lei Y, Yang B, Jiang X, Jia F, Li N, Nandi AK. Applications Of Machine Learning To Machine Fault Diagnosis: A Review And Roadmap. Mechanical Systems And Signal Processing. 2020;138:106587. DOI: Https://Doi.Org/10.1016/J.Ymssp.2019.106587

[11]. Xu M, Yu Q, Chen S, Lin J. Rolling Bearing Fault Diagnosis Based On CNN-LSTM With FFT And SVD. Information. 2024;15(7):399. DOI: Https://Doi.Org/10.3390/Info15070399

[12]. Li Y, Gan Y, Wu Y, Chen Z, Lv Q, Xiong Q. A Wavelet-Enhanced Multiscale Dense Capsule Network For Rotating Machinery Fault Diagnosis. Journal Of Mechanical Science And Technology. 2025;39:5763–76. DOI: Https://Doi.Org/10.1007/S12206-025-0914-X

[13]. Kumar S, Singh S, Yelve NP. Integrating Machine Learning With Empirical Mode Decomposition For Multiple Fault Diagnosis In Rotating Machinery. In: Andhare A, Ganesh Kumar P, Editors. Recent Advances In Condition Monitoring. NCCM 2024. Lecture Notes In Mechanical Engineering. Singapore: Springer; 2025. DOI: Https://Doi.Org/10.1007/978-981-96-5791-9_10

[14]. Randall RB, Antoni J. Rolling Element Bearing Diagnostics—A Tutorial. Mechanical Systems And Signal Processing. 2011;25(2):485–520. DOI: Https://Doi.Org/10.1016/J.Ymssp.2010.07.017

[15]. Yadav OP, Pahuja GL. Bearing Health Assessment Using Time Domain Analysis Of Vibration Signal. International Journal Of Image, Graphics And Signal Processing. 2020;12(3):27–40. DOI: Https://Doi.Org/10.5815/Ijigsp.2020.03.04

[16]. Altaf M, Akram T, Khan MA, Iqbal M, Iqbal Ch MM, Hsu CH. A New Statistical Features Based Approach For Bearing Fault Diagnosis Using Vibration Signals. Sensors. 2022;22(5):2012. DOI: Https://Doi.Org/10.3390/S22052012

[17]. Jebur NA, Soud WA. Fault Diagnosis Of Journal Bearings Using Vibration Analysis And Machine Learning Techniques. Journal Of Failure Analysis And Prevention. 2025;25:1167–80. DOI: Https://Doi.Org/10.1007/S11668-025-02157-Z

[18]. Nainwal A, Hussein L, Dixit KK, Michaelson J, Sehgal A, Gupta P. Predictive Maintenance For Industrial Equipment In Big Data: An Implementation With Random Forest. In: 2024 International Conference On Communication, Computer Sciences And Engineering (IC3SE). IEEE; 2024. P. 1417–22. DOI: Https://Doi.Org/10.1109/IC3SE62002.2024.10593295

[19]. Tzionis G, Ntemi M, Gialampoukidis I, Vrochidis S, Kompatsiaris I, Vlachopoulou M. Comparative Analysis Of Feature Importance In Machine Learning Models For Predictive Maintenance. In: Vrochidis S, Ducq Y, Gialampoukidis I, Hribernik K, Editors. Enterprise Interoperability XI I-ESA 2024 Proceedings Of The I-ESA Conferences. Springer, Cham; 2025. DOI: Https://Doi.Org/10.1007/978-3-031-93995-2_3

[20]. Yang Y, Iqbal MZ. Cost-Optimised Machine Learning Model Comparison For Predictive Maintenance. Electronics. 2025;14(12):2497. DOI: Https://Doi.Org/10.3390/Electronics14122497

[21]. Wang H, Zhang W, Yang D, Xiang Y. Deep-Learning-Enabled Predictive Maintenance In Industrial Internet Of Things: Methods, Applications, And Challenges. IEEE Systems Journal. 2023;17(2):2602–15. DOI: Https://Doi.Org/10.1109/JSYST.2022.3193200

[22]. Bhaskar RK, Kumaraswamy B. Harnessing CNN And LSTM Models For Enhanced Predictive Maintenance In Industrial Iot Systems. In: 2025 International Conference On Automation And Computation (AUTOCOM). IEEE; 2025. P. 247–51. DOI: Https://Doi.Org/10.1109/AUTOCOM64127.2025.10956547

[23]. Garcia J, Rios-Colque L, Peña A, Rojas L. Condition Monitoring And Predictive Maintenance In Industrial Equipment: An NLP-Assisted Review Of Signal Processing, Hybrid Models, And Implementation Challenges. Applied Sciences. 2025;15(10):5465. DOI: Https://Doi.Org/10.3390/App15105465

[24]. Li W, Li T. Comparison Of Deep Learning Models For Predictive Maintenance In Industrial Manufacturing Systems Using Sensor Data. Scientific Reports. 2025;15. DOI: Https://Doi.Org/10.1038/S41598-025-08515-Z

[25]. Kapoor D, Gupta D, Uppal M. Analyzing The Impact Of Edge, Fog And Cloud Computing On Predictive Maintenance In The Industrial Internet Of Things. Discover Computing. 2025;28:207. DOI: Https://Doi.Org/10.1007/S10791-025-09653-8

[26]. Keshta I, Majeed ME, Al-Shaikhli TR, Adel AH, Jabbar KA, Soni M. AI-Based Predictive Maintenance For Machine Vibration And Condition Monitoring In Industrial Applications. In: 2025 4th OPJU International Technology Conference (OTCON) On Smart Computing For Innovation And Advancement In Industry 50. IEEE; 2025. P. 1–7. DOI: Https://Doi.Org/10.1109/OTCON65728.2025.11070934

[27]. Wilk-Jakubowski JL, Pawlik L, Frej D, Wilk-Jakubowski G. The Evolution Of Machine Learning In Vibration And Acoustics: A Decade Of Innovation (2015–2024). Applied Sciences. 2025;15(12):6549. DOI: Https://Doi.Org/10.3390/App15126549

[28]. Molnar C. Interpretable Machine Learning: A Guide For Making Black Box Models Interpretable. Morisville, North Carolina: Lulu; 2020.

[29]. Samek W, Montavon G, Vedaldi A, Hansen LK, Müller KR, Editors. Explainable AI: Interpreting, Explaining And Visualizing Deep Learning. Cham: Springer; 2019.

[30]. Alshkeili HMHA, Almheiri SJ, Khan MA. Privacy-Preserving Interpretability: An Explainable Federated Learning Model For Predictive Maintenance In Sustainable Manufacturing And Industry 4.0. AI. 2025;6(6):117. DOI: Https://Doi.Org/10.3390/Ai6060117

[31]. Tsallis C, Papageorgas P, Piromalis D, Munteanu RA. Application-Wise Review Of Machine Learning-Based Predictive Maintenance: Trends, Challenges, And Future Directions. Applied Sciences. 2025;15(9):4898. DOI: Https://Doi.Org/10.3390/App15094898

[32]. Sampaio GS, Vallim Filho AR De A, Da Silva LS, Da Silva LA. Prediction Of Motor Failure Time Using An Artificial Neural Network. Sensors. 2019;19(19):4342. DOI: Https://Doi.Org/10.3390/S19194342

[33]. Bellman RE. Adaptive Control Processes. Princeton, New Jersey: Princeton University Press; 1961.

[34]. Bishop CM. Pattern Recognition And Machine Learning. New York, NY: Springer; 2006.

[35]. Jia F, Lei Y, Lin J, Zhou X, Lu N. Deep Neural Networks: A Promising Tool For Fault Characteristic Mining And Intelligent Diagnosis Of Rotating Machinery With Massive Data. Mechanical Systems And Signal Processing. 2016;72-73:303–15. DOI: Https://Doi.Org/10.1016/J.Ymssp.2015.10.025

[36]. Lei Y, Li N, Guo L, Li N, Yan T, Lin J. Machinery Health Prognostics: A Systematic Review From Data Acquisition To RUL Prediction. Mechanical Systems And Signal Processing. 2018;104:799–834. DOI: Https://Doi.Org/10.1016/J.Ymssp.2017.11.016

[37]. Oppenheim AV, Schafer RW. Discrete-Time Signal Processing. 3rd Ed. New York, NY: Pearson; 2013.
[38]. Hyndman RJ, Athanasopoulos G. Forecasting: Principles And Practice. 2nd Ed. Heathmont, Vic.: Otexts; 2018.
[39]. Zhang Y, Hutchinson P, Lieven NAJ, Nunez-Yanez J. Adaptive Event-Triggered Anomaly Detection In Compressed Vibration Data. Mechanical Systems And Signal Processing. 2019;122:480–501. DOI: Https://Doi.Org/10.1016/J.Ymssp.2018.12.039
[40]. Mcfadden PD, Toozhy MM. Application Of Synchronous Averaging To Vibration Monitoring Of Rolling Element Bearings. Mechanical Systems And Signal Processing. 2000;14(6):891–906. DOI: Https://Doi.Org/10.1006/Mssp.2000.1290
[41]. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Et Al. Scikit-Learn: Machine Learning In Python. The Journal Of Machine Learning Research. 2011;12.
[42]. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Et Al. Scipy 1.0: Fundamental Algorithms For Scientific Computing In Python. Nature Methods. 2020;17:261–72. DOI: Https://Doi.Org/10.1038/S41592-019-0686-2
[43]. Breiman L. Random Forests. Machine Learning. 2001;45:5–32. DOI: Https://Doi.Org/10.1023/A:1010933404324
[44]. Cortes C, Vapnik V. Support-Vector Networks. Machine Learning. 1995;20:273–97. DOI: Https://Doi.Org/10.1007/BF00994018
[45]. Cover T, Hart P. Nearest Neighbor Pattern Classification. IEEE Transactions On Information Theory. 1967;13(1):21–7. DOI: Https://Doi.Org/10.1109/TIT.1967.1053964
[46]. Rumelhart DE, Hinton GE, Williams RJ. Learning Representations By Back-Propagating Errors. Nature. 1986;323:533–6. DOI: Https://Doi.Org/10.1038/323533a0
[47]. Sokolova M, Lapalme G. A Systematic Analysis Of Performance Measures For Classification Tasks. Information Processing And Management. 2009;45:427–37. DOI: Https://Doi.Org/10.1016/J.Ipm.2009.03.002
[48]. Kohavi R. A Study Of Cross-Validation And Bootstrap For Accuracy Estimation And Model Selection. In: Proceedings Of The 14th International Joint Conference On Artificial Intelligence. Morgan Kaufmann Publishers Inc.; 1995. P. 1137–43.
[49]. Varanis M, Silva A, Mereles A, Pederiva R. MEMS Accelerometers For Mechanical Vibrations Analysis: A Comprehensive Review With Applications. Journal Of The Brazilian Society Of Mechanical Sciences And Engineering. 2018;40. DOI: Https://Doi.Org/10.1007/S40430-018-1445-5
[50]. Miyoshi S. Statistical-Mechanical Analysis Of Adaptive Filter With Clipping Saturation-Type Nonlinearity. IEEE Transactions On Signal Processing. 2022;70:4867–82. DOI: Https://Doi.Org/10.1109/TSP.2022.3212527
[51]. Ghemari Z, Belkhiri S, Reguieg SK. Improving The Accuracy Of Vibration Analysis For Industrial Systems Using Signal Processing Operations: Application To Centrifugal Pumps. Transactions On Electrical And Electronic Materials. 2025;26:822–44. DOI: Https://Doi.Org/10.1007/S42341-025-00628-6
[52]. Wang J, Li C, Han S, Sarkar S, Zhou X. Predictive Maintenance Based On Event-Log Analysis: A Case Study. IBM Journal Of Research And Development. 2017;61(1):11:121–32. DOI: Https://Doi.Org/10.1147/JRD.2017.2648298
[53]. Aitsam M, Goyal G, Bartolozzi C, Di Nuovo A. Vibration Vision: Real-Time Machinery Fault Diagnosis With Event Cameras. In: Del Bue A, Canton C, Pont-Tuset J, Tommasi T, Editors. Computer Vision – ECCV 2024 Workshops ECCV 2024 Lecture Notes In Computer Science. Springer; 2025. DOI: Https://Doi.Org/10.1007/978-3-031-92460-6_18
[54]. Breiman L. Bagging Predictors. Machine Learning. 1996;24:123–40. DOI: Https://Doi.Org/10.1023/A:1018054314350
[55]. Hastie T, Tibshirani R, Friedman J. The Elements Of Statistical Learning: Data Mining, Inference, And Prediction, Second Edition. New York, NY: Springer New York; 2009.
[56]. Schölkopf B, Smola AJ. Learning With Kernels: Support Vector Machines, Regularization, Optimization, And Beyond. Cambridge, Massachusetts: The MIT Press; 2001.
[57]. Hsu CW, Chang CC, Lin CJ. A Practical Guide To Support Vector Classification. Taipei, Taiwan: Department Of Computer Science, National Taiwan University; 2003.
[58]. Maaten L Van Der, Hinton G. Visualizing Data Using T-SNE. Journal Of Machine Learning Research. 2008;9(86):2579−2605.
[59]. Nakajima S. Introduction To TPM: Total Productive Maintenance. Cambridge, Massachuttes: Productivity Press; 1988.
[60]. Muchiri P, Pintelon L. Performance Measurement Using Overall Equipment Effectiveness (OEE): Literature Review And Practical Application Discussion. International Journal Of Production Research. 2008 Jul;46(13):3517–35.
[61]. Fernández-Delgado M, Cernadas E, Barro S, Amorim D. Do We Need Hundreds Of Classifiers To Solve Real World Classification Problems? Journal Of Machine Learning Research. 2014;15(1).
[62]. Shwartz-Ziv R, Armon A. Tabular Data: Deep Learning Is Not All You Need. Information Fusion. 2022;81:84–90. DOI: Https://Doi.Org/10.1016/J.Inffus.2021.11.011
[63]. Shi W, Cao J, Zhang Q, Li Y, Xu L. Edge Computing: Vision And Challenges. IEEE Internet Of Things Journal. 2016;3(5):637–46. DOI: Https://Doi.Org/10.1109/JIOT.2016.2579198
[64]. Achouch M, Dimitrova M, Ziane K, Karganroudi SS, Dhouib R, Ibrahim H, Et Al. On Predictive Maintenance In Industry 4.0: Overview, Models, And Challenges. Applied Sciences. 2022;12(16):8081. DOI: Https://Doi.Org/10.3390/App12168081
[65]. Ribeiro MT, Singh S, Guestrin C. "Why Should I Trust You?": Explaining The Predictions Of Any Classifier. In: KDD '16: Proceedings Of The 22nd ACM SIGKDD International Conference On Knowledge Discovery And Data Mining. New York, NY, USA: Association For Computing Machinery; 2016. P. 1135–44. DOI: Https://Doi.Org/10.1145/2939672.2939778