International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(20s) 2026 pp. 1116-1139
https://doi.org/10.70917/ijcisim-2026-5255
Article

Adaptive unsupervised anomaly detection with dynamic recalibration for scalable industrial asset reliability monitoring

Dhirendra Singh[1]*, Durga Prasad Panday[1], Manish Kumar[1,2]

[1]Sustainability Cluster, School of Advanced Engineering, UPES, Dehradun-248007, Uttarakhand, India
[2]Escuela de Ingeniería y Ciencias, Tecnológico de Monterrey, Campus Monterrey, Monterrey 64849, Nuevo León, Mexico
*Corresponding Author Email: dhireens@gmail.com; ORCID: 0009-0005-6680-230X

Abstract: Anomaly-detection methods typically fail to detect changes in operating conditions or are unresponsive in detecting changes in the aging of industrial assets. This paper outlines an unsupervised lightweight machine learning-based approach to adaptively detect anomalies in the reliability of industrial assets. It incorporates in its framework operating-range normalisation, principal component analysis, Euclidean-distance anomaly scoring, Gaussian-based dynamic thresholds, sensor-contribution ranking and autonomous baseline recalibration. The validation was performed by taking 10 sensor variables initially screened in a critical P-101 A/B pump in an HDPE slurry polymerisation process, sampled approximately every 15 min and nine variables were selected for sensor-quality screening and then used for the validation. The framework identified all 44 reported functional failures resulting in a 100% recall, 84.62% precision, an average warning lead time of 72 h and an F1-score of 91.67%. The findings show that a self-recalibrating unsupervised model can successfully detect faults in a scalable manner, with an interpretable model, low computational complexity and can adapt to ageing assets and varying operating baselines without labelled fault information.
Keywords: Predictive maintenance; Anomaly detection; Unsupervised learning; Dynamic recalibration; Asset reliability; Fault diagnosis; Petrochemical systems

1. Introduction

In advanced manufacturing, oil and gas, utilities and petrochemical systems, process safety, production continuity and the reliability of industrial assets are critical to environmental protection. An abnormal process condition, shutdown, hazardous release and cascading operational disturbances are not necessarily restricted to disruption of equipment like pumps and compressors or heaters and pipelines; they can be caused by failure of SIF's or other components within the tightly coupled technological system. Enlightened with historian data, industrial Internet of Things (IIoT) infrastructure and machine-learning-based analytics, that shift in the culture of factories, from reactive to condition-based and predictive maintenance, has become even quicker. These strategies will enable continuous operational data to be utilised for detecting deviations from the equipment's expected behaviour, enabling maintenance to be achieved in advance [1,2,22,23,30].

Reliability-centred maintenance, failure mode and effects analysis, and time-based preventive maintenance are all traditional methods that continue to be crucial in the realm of industrial asset management. However, the operating conditions of industrial assets is not necessarily stationary one. Normal operating behaviour can vary over time as a result of equipment ageing, maintenance interventions, product changes, plant restarts, feedstock variation, and environmental and process disturbances. However, the fixed alarms and manually set thresholds can cause nuisance alarms if the operating conditions change or if there isn't adequate sensitivity to sharp changes in the operating condition that occurred before the alarms were set. Such challenges have given rise to the need for adaptive approaches, interpretable, and robust ones that are able to characterise asset behaviour from data streams from industry under changing operating conditions [3,4,25,27,29].

In recent years, various machine learning techniques for predictive maintenance and fault diagnosis have been developed for industrial systems such as pipelines, temperature prediction [5,6], rotating-equipment prognosis, and maintenance diagnosis in power systems [7]. Supervised learning methods offer a high improvement of prediction potential, but usually coded examples and failure histories are needed for modelling. Failure histories may be challenging to compile for critical industrial assets, as failure events are relatively rare, and maintenance procedures can differ in practice from asset to asset and among different operating periods. The use of hybrid and deep learning architectures can yield complex process behavior modelling, but can also create increased computation and implementation needs [8-10]. This is especially important if many assets are involved and used over many years.

In recent studies in the domain of industrial anomaly detection and predictive maintenance, the focus also shifted towards adaptive modelling, real-time streams of industrial data, integrating a digital-twin paradigm, concept-drift, and explainable anomaly detection. Several anomaly detection review works in the IIoT domain have pointed out that most of the anomaly detection models are actually static and do not cope well with variation of operational conditions, and further several explainability approaches in the IIoT domain have been reported that instead of direct physical diagnosis, offer statistical feature attribution [18]. A number of adaptive predictive-maintenance (APM) approaches have shown promise when adapting models as conditions of the equipment and environment evolve [19]. In the context of industrial monitoring and hybrid learning or digital-twin, preferably a mixture that is neither too heavy in the detection task nor excessively complex, the trade-off between detection sensitivity, interpretability and computational efficiency has been accentuated in other studies [20,21]. These developments suggest that actual methods of anomaly detection need to be designed that will adapt to varying operating baselines without relying on extensive failure details in labeled data or on complex architectures to process massive amounts of data at runtime.

However, in industrial deployment, the approach of just using anomaly detection isn't enough. The utility of a monitoring system also relies on the maintenance of an appropriate definition of ‘normal operation’ as equipment and process conditions change over time. An effective monitoring framework should then describe typical operating conditions through operational data, detect key deviations, constantly redefine the baseline as soon as a new stable operating condition is realised and give useful and understandable information to aid decisions by operators and maintenance personnel. This adaptation would ideally eliminate the need for frequent manual tuning of thresholds on specific assets to the exclusion of requiring labeled examples of each failure mode.

The use of unsupervised anomaly detection is not claimed as a primary methodological novelty of this study, as PCA-based monitoring and adaptive thresholding are already well established. Instead, the contribution lies in the integration of three complementary elements within a lightweight monitoring framework: (i) a low-dimensional statistical representation of normal multivariate operating behaviour, (ii) an independent baseline recalibration mechanism that updates the reference operating state and associated thresholds after a new stable regime is established, and (iii) sensor-level contribution analysis that provides interpretable information on the variables driving each anomaly. Unlike conventional PCA-based monitoring approaches that typically rely on a fixed reference model, or adaptive-threshold methods that modify decision limits without explicitly redefining the underlying normal operating baseline, the proposed framework updates both the baseline representation and its corresponding anomaly thresholds while retaining a simple and computationally inexpensive scoring structure. This enables the monitoring system to accommodate persistent changes arising from equipment condition, process operation, maintenance interventions, or environmental conditions without requiring labelled failure data or repeated manual threshold adjustment. The novelty therefore resides in the coordinated use of dynamic baseline recalibration, lightweight multivariate anomaly scoring, and interpretable sensor-level diagnostics as a practical self-adapting reliability-monitoring architecture.

This study has three implications. It first investigates whether a lightweight unsupervised approach will be able to identify documented functional failures with no labelled data being provided as part of the development of the model. It is secondly designed to feature dynamic re-calibration to minimise reliance on operating baselines and to enable tracking of operating changes. Third, it includes sensor contribution analysis to give interpretable information concerning the measured process variables of monitored anomalies which helps to support diagnostic investigation and maintenance decision making. Thus, the study emphasizes the design of the practical application of unsupervised anomaly detection, adaptive baseline management and sensor-level interpretation in the context of an industrial asset-monitoring system.

This work evaluates the framework using three years of historian data from a critical P-101 A/B pump in a high-density polyethylene (HDPE) slurry polymerisation process. The framework includes data cleaning, operating-range normalisation, principal component analysis (PCA), Euclidean-distance anomaly scoring, Gaussian based warning and alert thresholds, sensor-contribution ranking and autonomous baseline recalibration [9,17,24,26,28,31]. In model performance evaluation, the modeled performance is compared with documented failures in the real system, but no failure information is used in model learning. The evaluation, therefore, gets to see whether the proposed framework can authenticate errors corresponding to known functional failures, and is able to adjust the baseline to changing operating conditions.

2. Industrial System and Data

2.1 Process Description and Equipment under Study

The industrial application is a slurry polymerisation process of HDPE; the highlight of this project is focused on the critical P-101 A/B pump system. The pumps are part of the heat-removal loop that co-relates to the Reactor R-101 and therefore are directly responsible for maintaining the safe operation of the reactor and stable polymerization temperature.

In this study, polymer lump formation in the slurry is regarded as an abnormal condition. An agglomeration or blockage of polymer may partially or totally restrict the flow of circulation and cause pressure surges and, in extreme cases, plant or pump shutdown. Because proper slurry circulation is a critical aspect of the reactor for cooling the charge, abnormal pump behaviour should be identified as soon as possible to ensure uninterrupted production and process safety.

This system serves as a meaningful industrial illustration of reliability and process-safety analysis, since degradation may impact not only the reliability of the pump, but also affect the stability of the larger process and reactor heat removal. Instead of forming a diagnostic model based on a pre-determined historical failure signature, this study takes a generic, unsupervised anomaly-detection approach. The aim is to decide whether a model that has been built with a focus on normal operating behaviour can reliably recognize the significant multivariate deviations that are associated with a developing process or equipment abnormality while at the same time being adequate enough to normal variations in the operating regime.

This formulation enables the study to appraise anomaly detection in terms of an early-warning approach, as opposed to a classifier limited to a previously labeled failure mode.

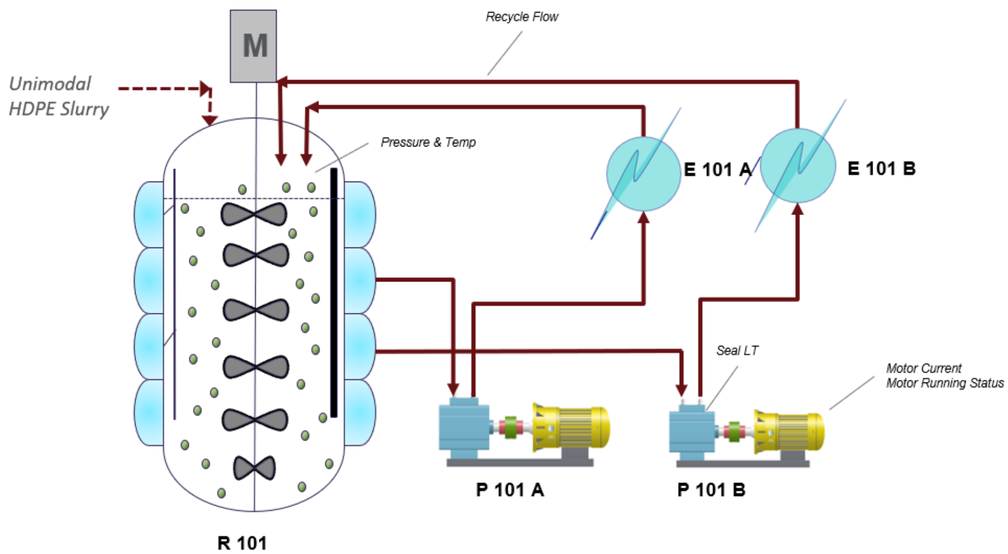

Fig. 1. Schematic representation of the HDPE slurry polymerisation process and the position of the P-101 A/B pump in the heat-exchange loop.

2.2 Historian Data and Process Variables

Process variables are continuously recorded by industrial historians integrated with the Distributed Control and Supervisory Control Systems, which serve as a basis for reconstructing events, validating models and performing long-term reliability analysis [11]. Data on the selected sensor variables were sampled at an approximately 15 min interval, which resulted in about 35,040 observations per variable per year and over 105,120 observations per variable over 3 years of study. This multi-year set of data included stable data, cold start-up, transitioning in the context of maintenance or any event prior to a failure. The data base provided thus, an adequate foundation for assessment of the anomaly-detection framework under normal operating conditions and under abnormal process behaviour.

Two complementary approaches were used to select the input variables. The review of the process paperwork and instrumentation documentation was to document measurements in the pump system boundary or directly involved in the operation of P-101 A/B. Secondly exploratory statistical screening was conducted to evaluate each of the potential candidates for three major factors: variability, operating range, data quality and interdependency with other process variables.

Initially, 10 candidate variables were identified. After this assessment, 9 variables were selected for use in the PCA model. This comprised the range of available vibration measurements, where it was judged that this range was of sufficient value to characterize pump behavior. Motor current was not included in the PCA input set as its distribution was observed to fall into the range of concern for sensor reliability and data quality. It is important to accentuate that it was not excluded because of no physical relevance to pump condition but because of measurement quality considerations.

This selection process enabled the selection of the pump's main hydraulic and mechanical behavior and minimized the effect of unreliable sensor readings on the final PCA input set.

Table 1. Process variables used for screening and model development.

Variable	**Process measurement**	**Unit**	**Model use**
X1	Motor current	A	Screened; excluded from PCA due to sensor reliability issue

X2	Pump DE vibration	m/s² pk	Retained
X3	Pump NDE vibration	m/s² pk	Retained
X4	Pump hexane flush inlet	m³/h	Retained
X5	Seal pot level	%	Retained
X6	Pump flow	m³/h	Retained
X7	Suction temperature	°C	Retained
X8	Suction pressure	psi	Retained; also used for functional-failure definition
X9	Discharge pressure	psi	Retained
X10	Discharge temperature	°C	Retained

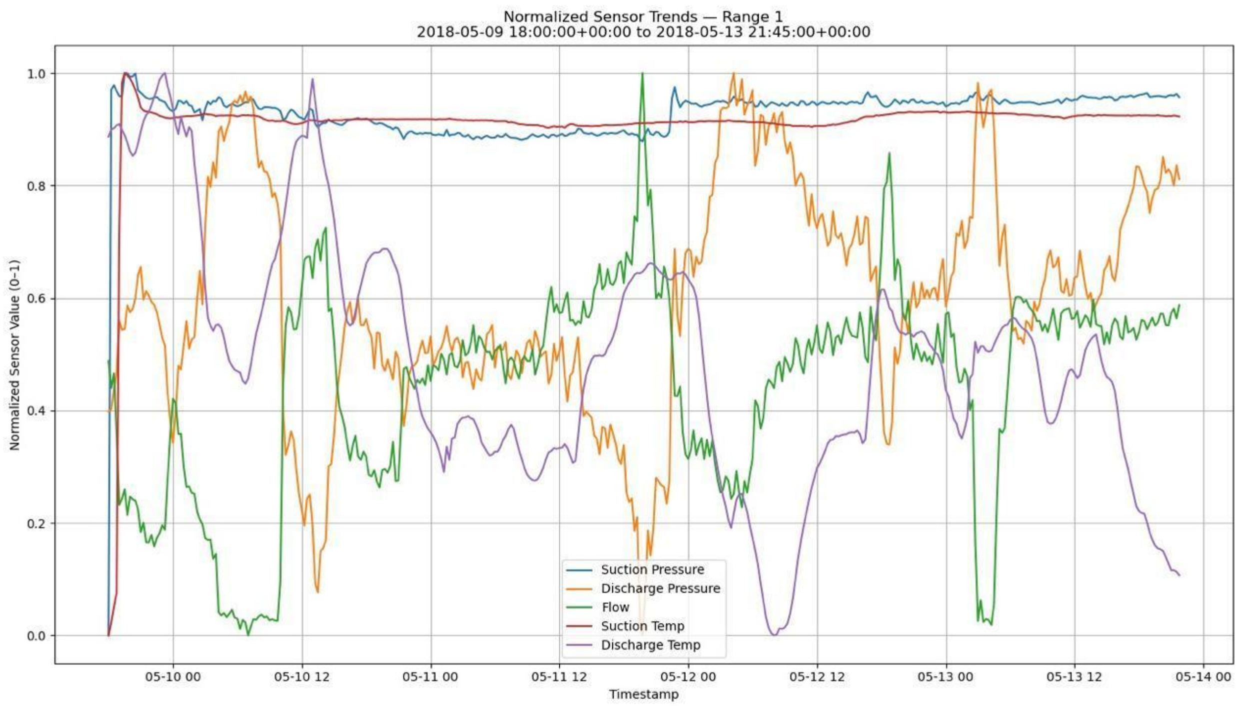

Figure 2. Raw historian sensor trends during industrial pump operation

This figure presents the time-series behaviour of key process variables collected from the industrial historian system. The observed variations demonstrate the dynamic operating conditions used for developing the unsupervised anomaly-detection framework.

2.3 Data Cleaning and Failure Definition

The historian data were pre-processed before model development to improve data quality and prevent sensor artefacts from influencing the learned representation of normal operation. A small number of isolated missing observations that were not associated with known failure events were imputed using local-window means, providing a pragmatic approach for maintaining continuity in long industrial time-series records [12]. Extended flat-line periods, physically implausible values, and clearly corrupted sensor segments were excluded from the training dataset to prevent defective sensor behaviour from being incorporated into the model as normal operating behaviour. The same pre-processing rules and normalisation parameters established during model development were subsequently applied consistently during inference.

Functional failure was defined according to the reliability principle that an item is considered to have failed when it can no longer perform its required function [13]. Based on operational, safety, and process-engineering considerations, the functional-failure criterion for the P-101 A/B pump system was defined using a suction-pressure threshold of 1200 psi. A binary running-status indicator (RS) was therefore assigned as:

$$RS = \{1, \quad P\_suction >= 1200\,psi\; 0, \quad P\_suction < 1200\,psi$$

Here, $RS = 1$ represents operation above the defined functional threshold, whereas $RS = 0$ indicates that the suction pressure has fallen below the functional-failure criterion. Importantly, this status indicator was not used to train the anomaly-detection model. It was retained exclusively as an independent operational reference for post hoc validation of the unsupervised framework. This separation preserves the unsupervised nature of the proposed methodology while allowing detected anomalies to be evaluated against an engineering-defined failure condition.

3. Proposed Adaptive Anomaly-Detection Framework

3.1 Framework Overview

This proposed framework was created for the scalable condition and reliability monitoring of industrial assets for which there is limited labelled failure data or normal operation behavior that is subject to change over time. Instead of predefining fault signatures, a framework-based multivariate representation of normal process behavior is generated, and subsequent observations are tested by comparing them with an adaptive normal-operation baseline.

The methodology comprises five major steps: (1) engineering based sensor screening and data cleaning, (2) operating-range normalization of data by statistics, (3) lowering the dimension of the correlated process variables by PCA and, (4) computing the anomaly score for each observation based on the Euclidean distance from the current normal-baseline PCA centroid, and (5) warning and alert generation via a gaussian thresholding plus operating range-based dynamic re-calibration. Moreover, a sensor-contribution layer is used to order the above-mentioned process variables according to their correlation with the detected anomalies, aiding the engineering interpretation and then, root cause analysis.

The integrated data, model, insight, and feedback flows of the proposed methodology are illustrated in Fig. 3.

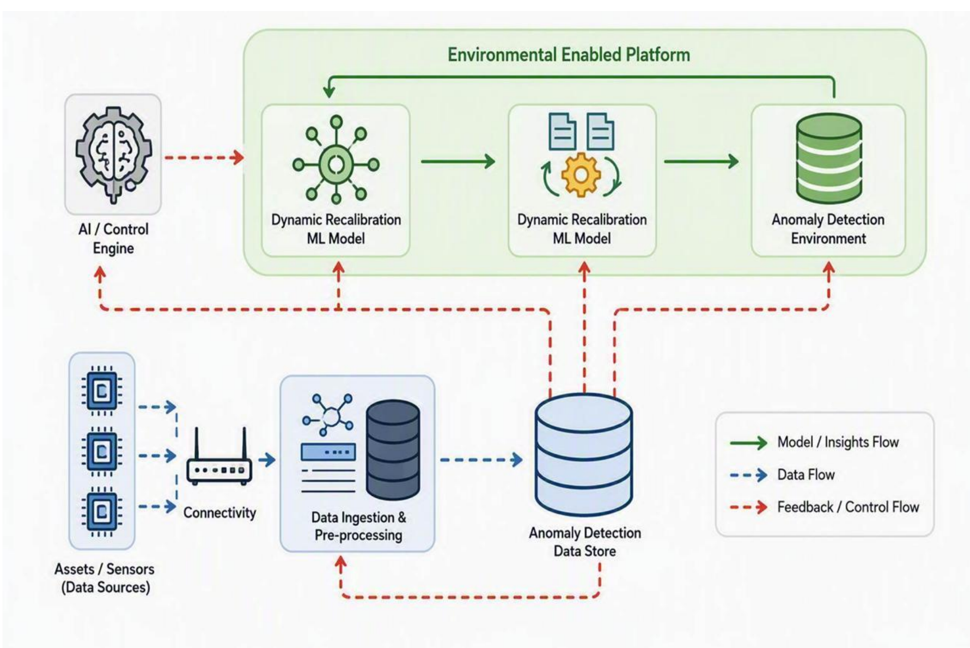

Fig. 3. Integrated methodology architecture showing data ingestion, dynamic recalibration, anomaly detection, and feedback/control flows.

An important aspect of this approach is that it is not pump specific and does not rely on a previous labelled failure signature. Rather, the inputs to be considered are historian data, a set of engineering relevant process variables, an initial operating window for stable normal operation, and a criterion for future functional failure that is set independently of the process. The framework is therefore designed to detect large multivariate deviations from the learned normal behavior, not just classifications of faults with which it has been seen before.

The logic used in the modelling, therefore, is to be transferable to any type of industry equipment when needed, provided that appropriate sensor selection, data-quality evaluation, baseline definition, and engineering validation of that equipment is carried out before deployment. Note that the framework has a generic level of method and selected variables, operating ranges, and recalibration criteria and validation thresholds are always dependent on the characteristics and primary operating context of the monitored asset.

3.2 Operational Range Estimation

The first step was to establish a representative operating range for each retained sensor variable. Where appropriate engineering or design limits were unavailable, or where such limits did not adequately represent the observed operating regime, statistical operating limits were estimated from the filtered historical data. For each sensor , the lower and upper statistical limits were defined as:

$$Low\_i = \mu\_i - 2\sigma\_i \qquad (1)$$

$$High\_i = \mu\_i + 2\sigma\_i \qquad (2)$$

where $\mu_i$and $\sigma_i$denote the mean and standard deviation, respectively, of sensor i over the screened reference operating period. The $\mu_i \pm 2\sigma_i$ interval was used as a data-derived reference operating range rather than as a hard physical or safety limit. This distinction is important because the statistical limits describe the variability observed in the reference data and should not be interpreted as equipment design constraints.

3.3 Normalization

The retained sensor variables were measured in different physical units and had substantially different numerical scales. Normalization was therefore performed before PCA to prevent variables with larger numerical magnitudes from disproportionately influencing the principal components [14]. For each sensor at time, operating-range normalization was performed as:

$$x'\_\{i,t\} = (x\_\{i,t\} - Low\_i) \,/\, (High\_i - Low\_i) \qquad (3)$$

where $x\_\{i,t\}$ is the data from sensor i at time t. For online scoring, baseline minimum and maximum values were also recorded for the subsequent observations to be transformed consistently with baseline:

Under this transformation, values at the lower and upper statistical limits correspond to 0 and 1, respectively. Values outside the reference interval were not necessarily clipped, because excursions beyond the learned operating range may contain information relevant to anomaly detection.

For subsequent online observations, the scaling parameters were fixed from the corresponding baseline or stable normal operating window and then applied to new observations. When baseline minimum and maximum scaling was used, the transformation was:

$$x'\_\{i,t\} = (x\_\{i,t\} - Min\_i) \,/\, (Max\_i - Min\_i) \qquad (4)$$

where Min_i and Max_i are obtained from the reference baseline and retained during subsequent scoring. Importantly, these parameters should not be recomputed independently for each incoming observation period, because doing so could rescale abnormal behaviour toward the normal range and reduce anomaly sensitivity.

In the implemented framework, Eq. (3) is used to determine the normalization parameters from the reference baseline, while Eq. (4) applies these fixed baseline-derived parameters to subsequent online observations. Therefore, the two equations do not represent alternative normalization methods; rather, they correspond to the baseline normalization and online scoring stages, respectively. The normalization parameters remain fixed during online monitoring and are updated only when a new stable baseline is accepted during recalibration.

3.4 PCA Transformation and Anomaly Scoring

Suppose that x_t is the normalised multivariate sensor vector at time t. The current stable normal operating window (SLOW) was used to fit a PCA, which was then applied to reduce x_t to a lower dimensional k dimensional score vector y_t. The first two components accounted for a large proportion of the baseline variance for each window of recalibration and therefore were used for anomaly scoring and visual interpretation for this study, $k = 2$.

The selection of two principal components was not based solely on visualization. It was chosen to retain the dominant structure of normal operating variability while maintaining a low-dimensional and computationally lightweight representation for anomaly scoring. A sensitivity analysis was also performed using different numbers of retained principal components to assess whether the anomaly-detection results were materially affected by this choice. The results showed that the overall detection performance remained comparable across the tested configurations, supporting the use of two principal components in the final framework.

Given the loading matrix W_k of the first k principal component loadings the PCA score vector is given by:

$$y\_t = W\_k\text{^}T\, x\_t \qquad (5)$$

The current normal operating window centroid in PCA space is found to be:

$$c = (1/n)\, \Sigma\_\{t = 1\}\text{^}\{n\}\, y\_t \qquad (6)$$

The anomaly score, at the time t, is then equal to the Euclidian distance between the PCA score at t and the baseline centroid:

$$D\_t = ||\, y\_t - c\, ||\_2 \qquad (7)$$

A small distance is said to be similarity to the current learned normal state, while a large distance is said to be multivariate deviation from normal operation.

3.5 Justification for PCA-Based Scoring and Gaussian Thresholding

Our goal was to create a generic, transparent and computationally efficient anomaly-detection framework, and not an asset-specific black-box classifier, so that's why we chose PCA. Process variables such as flow, pressure, temperature, vibration and seal variables are correlated and control the behaviour of an industrial pump. This is an appropriate application of PCA as it is used to map correlated multivariate sensor data into a reduced dimensional space, while retaining the dominant variance structure of normal operation [15]. This provides the framework to identify unusual interactions of sensor behaviour rather than the univariate violations of thresholds.

Due to its relatively simple PCA-based structure and Euclidean-distance anomaly scoring, the proposed framework is expected to require lower computational resources than more complex deep-learning approaches such as deep autoencoders, recurrent neural networks, or transformer-based models. Its simpler model structure also facilitates baseline recalibration and provides more direct interpretability for maintenance engineers. This helps to meet the envisioned objective of the framework as a scalable reliability-monitoring model which is replicated to other asset types while minimizing site-specific model engineering. Architectures that can deliver high accuracy is more complex, and require large numbers of labelled or simulated anomaly samples to train, or have high computation cost, limiting the architectures to deployment in industrial plants where asset populations are not necessarily homogeneous [18-21].

A simple, interpretable and dynamically updateable alarm logic was provided with the use of the Gaussian thresholding. The distribution of the Euclidean distances from the baseline centroid was modelled as a normal distribution in the anomaly-scoring step. This assumption is substantiated with practical and theoretical elements. The distance scores are affected by a number of small variations across correlated sensor variables, so that the overall spread of distance scores for the baseline distance values is fairly evenly distributed around a mean distance operating radius. The histogram of distances in this study had little skewness in the baseline data and a normal approximation was used. Second, Gaussian thresholding is in harmony with accepted industrial process-monitoring practice, which involves reporting statistical unusual behaviour when the process falls outside of three standard deviations on either side of the mean.

The warning threshold was thus taken as μ_D + 3σ_D and the alert threshold was taken as μ_D + 5σ_D, where μ_D and σ_D are the mean and standard deviation of the baseline distance scores, respectively:

$$T\_warning = \mu\_D + 3\sigma\_D \qquad (8)$$

$$T\_alert = \mu\_D + 5\sigma\_D \qquad (9)$$

The μ_D + 3σ_D warning threshold provides an early indication of abnormal behaviour while remaining robust to normal variation. The higher μ_D + 5σ_D alert threshold indicates a more substantial deviation from the baseline. Although Euclidean distances are non-negative and may theoretically follow a chi or Rayleigh distribution under a perfect multivariate-normal assumption, the Gaussian approximation remains appropriate for the high-percentile warning threshold used here. Robustness was assessed by comparing the Gaussian-based warning threshold with the empirical 99.7th percentile of the baseline distance scores; the two values were similar. The Gaussian fit therefore provides a reliable, lightweight, and readily recalibrated thresholding mechanism for industrial anomaly detection.

This thresholding selection is also consistent with the overall goal of the study to show that a simple, yet computationally very efficient model can attain good reliability performance without the massive computing power of deep architectures or the labeled data needs of supervised models.

3.6 Dynamic Recalibration

Dynamic recalibration is the key component in the framework. Upon a restart, maintenance action, operating transition or post-anomaly recovery, the algorithm looks for a new stable operating window of ~100h. When the candidate window meets the data quality and stability, the PCA fitting, centroid estimation and gaussian threshold fitting are performed again. A new baseline is then computed and the next baseline used for the future scoring.

This logic will minimise false alarms arising from legitimate operating changes whilst maintaining the sensitivity for degradation when compared to the latest healthy operating condition. It is not possible to go back to design during unstable/close to failure times as this normalise incipient degradation. The stable-window requirement and/or data-quality screening thus serve as a protection against undesirable updating of the baseline.

3.7 Sensor Contribution Anomaly Interpretation

After each alert, the PCA loadings and standardised deviations from each of the sensors were analysed to improve the interpretability of the diagnosis. The PCA loadings show which of the variables are most important to the main modes of normal operating variation, and the z-scores at the sensor level show which process variables exhibited the largest deviations from normal for every alert. This interpretation is two layered which enables the model to differentiate between statistical anomaly detection and engineering diagnosis.

For instance, alerts with high scores for negative flow deviation and discharge-pressure deviation would indicate a slurry restriction/impediment or obstruction of the impeller. Transient process disturbances are more likely to be the cause of alerts that are primarily based on discharge-temperature excursions. This interpretation layer further enhances the applicability of this framework as it enables maintenance personnel to progress from “anomaly detected” to “probable process area requiring inspection.”

The sensor contribution analysis identifies variables that are most strongly associated with a detected multivariate deviation and is intended to support engineering investigation. It does not establish causality or provide a confirmed root-cause diagnosis. The identified sensor contributions should therefore be interpreted together with process knowledge, maintenance records, and other engineering evidence.

3.8 Implementation and Validation Metrics

This framework has been carried out in python with the help of the libraries pandas, NumPy, scikit-learn, and SciPy for the time-series analysis, numerical computation, PCA and Gaussian fitting, respectively [16]. Validation was undertaken to demonstrate its ability to detect functional failures (where documented) as well as minimize false alarms during normal operating transitions and provide adequate lead time for maintenance planning.

True positive was defined as model alerts made before a known functional failure during the evaluation period. A false positive was when a model alert did not result in a functional failure during the evaluation period. There were documented functional failures that were not preceded by a model alert, termed as False negatives. Precision, recall and F1-score were defined as:

$$Precision = TP \,/\, (TP + FP) \qquad (10)$$

$$Recall = TP \,/\, (TP + FN) \qquad (11)$$

$$F1 = 2 \times Precision \times Recall \,/\, (Precision + Recall) \qquad (12)$$

The first warning/alert was defined as the lead time, which was defined as the period of time between the first warning/alert and the subsequent failure. This is directly relevant to the reliability-centred maintenance because it is the time window in which the pump is inspected, loaded or the pump is changed or planned intervention is required.

4. Results

4.1 Sensor Screening and Baseline Visualization

All the process variables that remained in the process were found to be within plausible ranges by exploratory screening. Motor current was omitted due to unreliable sensor behaviour in comparison to its operating range. Variables of vibration, hexane flush, seal-pot level, flow, pressure and temperature variables were kept. This allowed the baseline learned by the model to not be affected by the known defect in the sensors' quality.

The first stable normal baseline was identified from 2017-01-09T12:45:00Z to 2017-01-13T16:30:00Z. During this 100 h period, the operation of the pumps remained steady and the conservatively operated variables remained within the operation limits. A two-dimensional baseline for 78.36% of variance was obtained by PCA of this window. The initial warning threshold, alert threshold and maximum standard deviation of distances were 0.3077, 0.4637 and 0.0862, respectively.

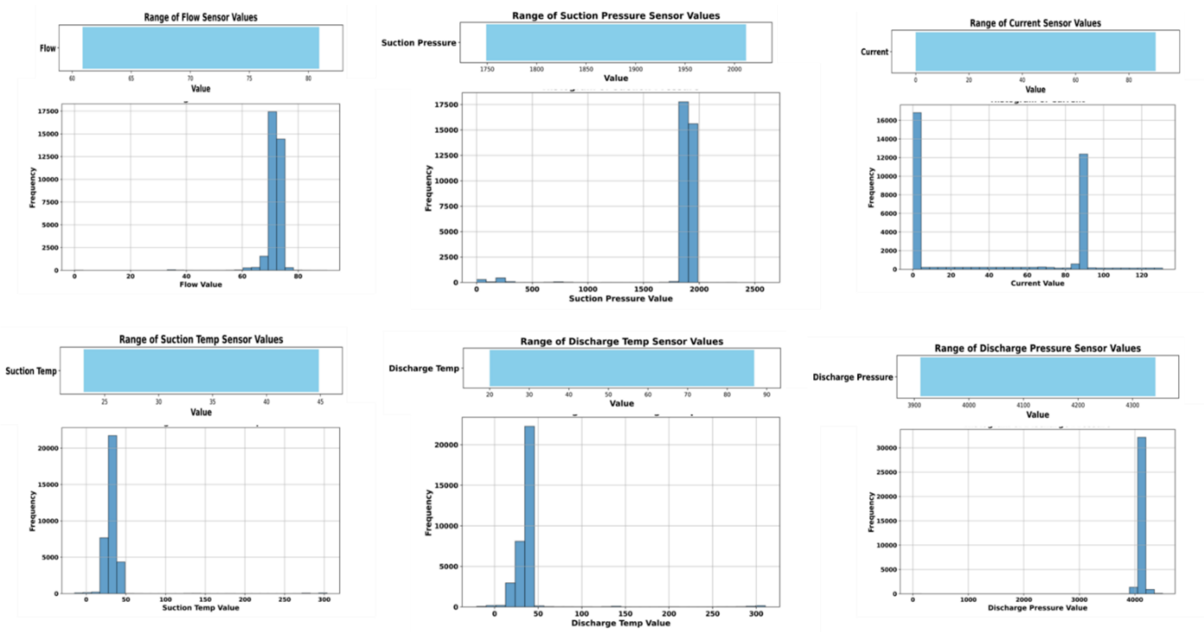

Fig. 4. Histograms of sensor data collected from the historian system during operational periods. The current signal was excluded from model building because of sensor reliability issues.

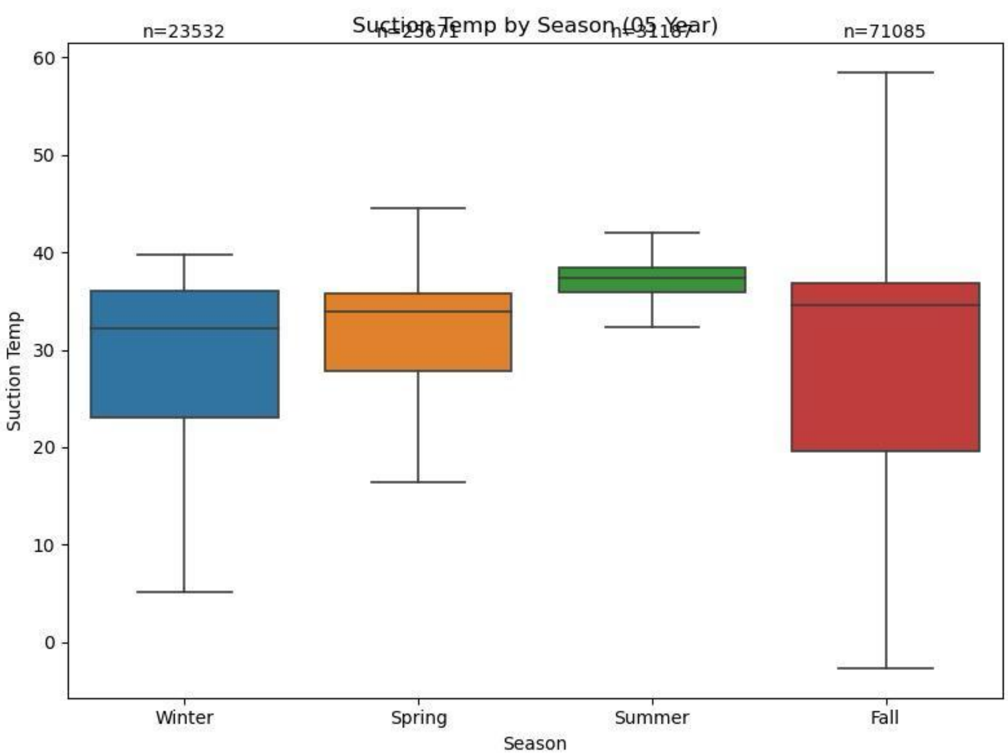

Figure 5. Seasonal distribution of suction temperature measurements before anomaly modelling.

This figure illustrates the variability and seasonal distribution of suction temperature measurements across different operating periods. The observed spread and median differences highlight changes in process behaviour that are considered during baseline characterization.

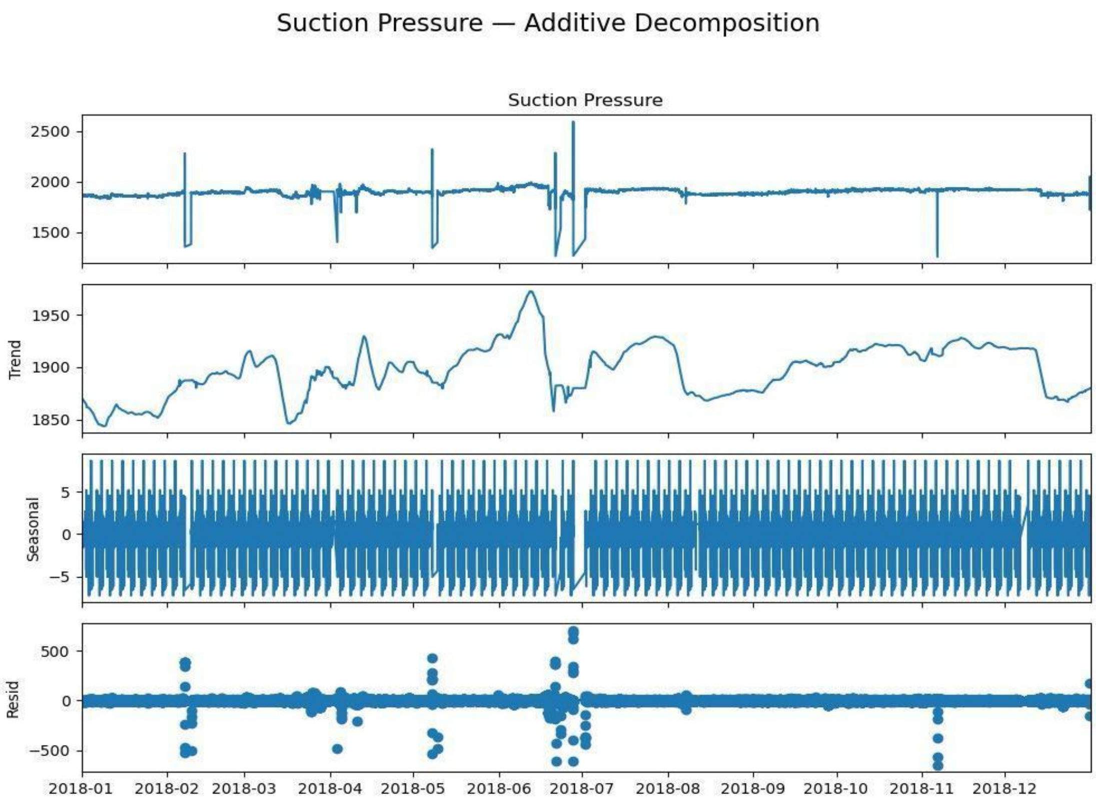

Figure 6. Temporal decomposition of suction pressure behavior during baseline operating conditions.

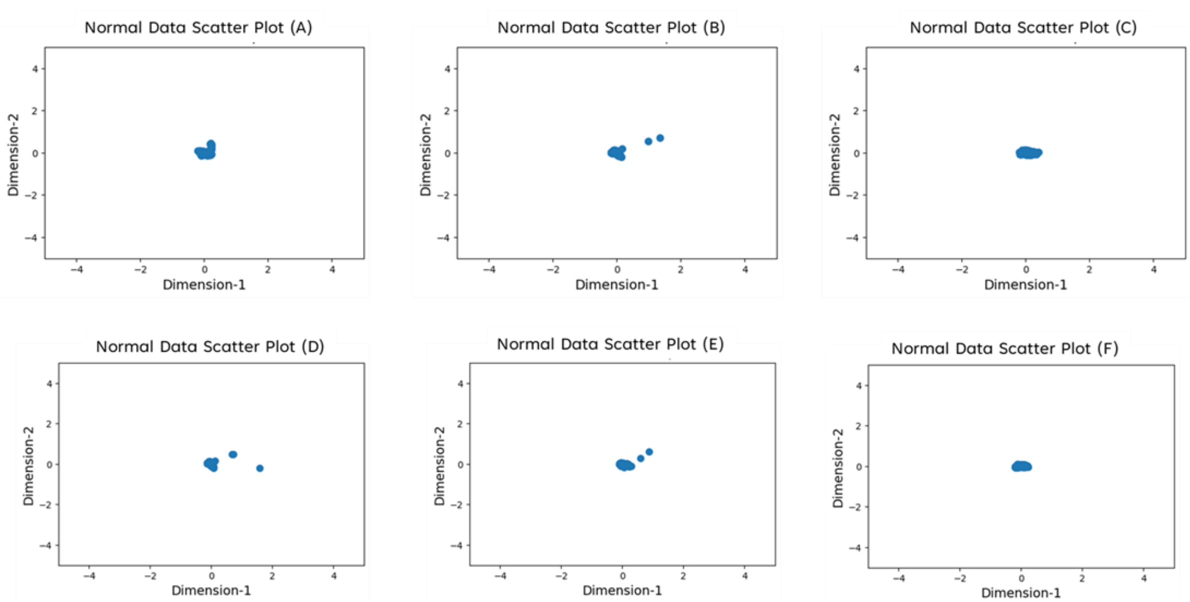

Fig. 7. PCA-transformed scatter plots of normal operating baselines identified during dynamic recalibration. Each cluster represents a stable operating window in two-dimensional PCA space.

4.2 Dynamic Recalibration Events

A variety of normal operating baselines were found over the three-year period of validation. The events for recalibration, explained variance, warning thresholds, alert thresholds and maximum SDs of the baseline distance scores are summarized in Table 2. The cumulative variance explained by the first two PCA components was between 76.13% and 92.03%, thus showing that a significant amount of the structure in different operating states was captured in this representation. The differences in the thresholds between baselines also demonstrate the need for not using a fixed threshold for long-term deployment.

Table 2. Key parameters for identified normal operating baselines and recalibration events

Recalibration event	Normal operating range (start-end)	Explained variance	Warning threshold	Alert threshold	Year
Initial baseline	2017-01-09T12:45:00Z - 2017-01-13T16:30:00Z	78.36%	0.3077	0.4637	2017
Recalibration 1	2017-07-22T02:00:00Z - 2017-07-26T06:00:00Z	76.13%	0.3185	0.4634	2017
Recalibration 2	2017-11-23T21:45:00Z - 2017-11-28T01:30:00Z	91.20%	0.1676	0.2548	2017
Recalibration 3	2018-01-02T01:00:00Z - 2018-01-06T04:45:00Z	77.58%	0.2554	0.3654	2018
Recalibration 4	2018-05-09T18:00:00Z - 2018-05-13T21:45:00Z	89.23%	0.3990	0.6170	2018
Recalibration 5	2018-11-06T14:15:00Z - 2018-11-10T18:00:00Z	83.75%	0.2515	0.3786	2018
Recalibration 6	2019-01-25T09:30:00Z - 2019-01-29T13:30:00Z	89.36%	0.3352	0.5015	2019
Recalibration 7	2019-09-07T09:45:00Z - 2019-09-11T13:30:00Z	92.03%	0.3422	0.5350	2019

The value in the proposed approach is illustrated by the changes in the threshold after recalibration. The model did not impose changes to the 2017 baseline to fit future data. Instead it recalculated the centroid and the thresholds after stable operating windows were found, thus retaining sensitivity to degradation while not triggering the alarm system when the system was in normal operating mode.

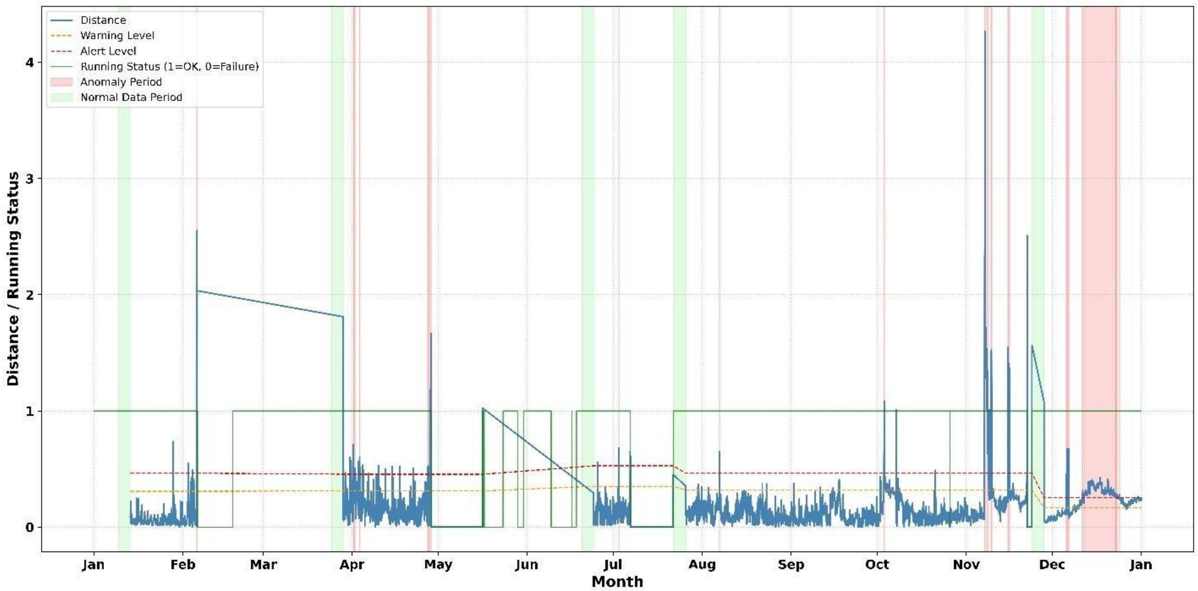

Fig. 8. Deviation plot for 2017 showing anomaly scores with dynamically adjusted warning and alert thresholds.

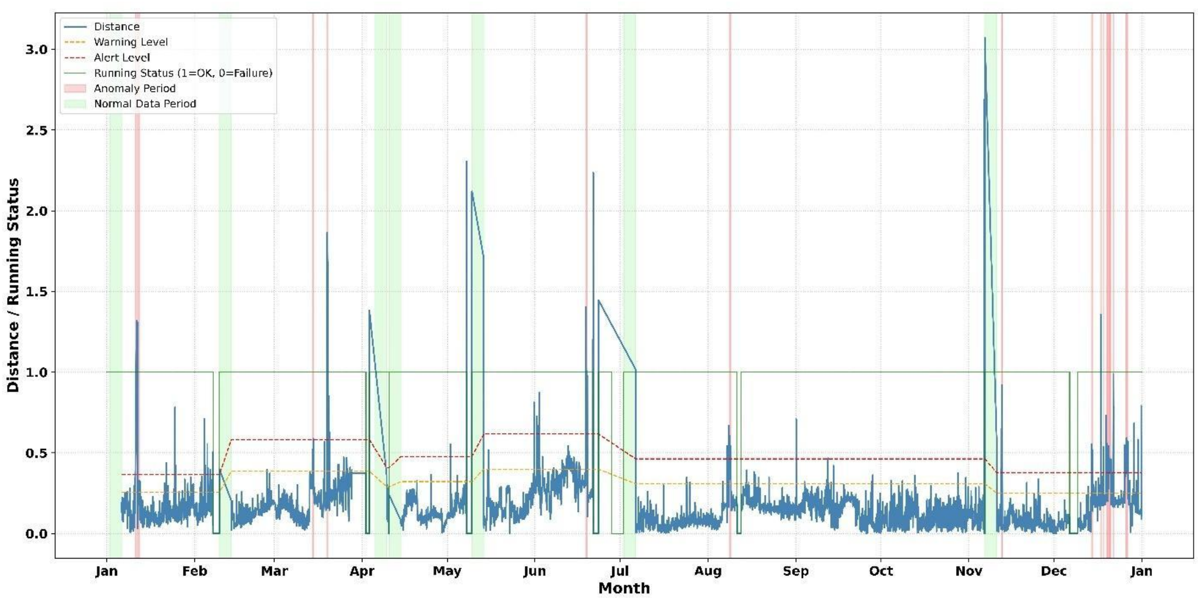

Fig. 9. Deviation plot for 2018 showing adaptive anomaly detection and recalibration across operational changes.

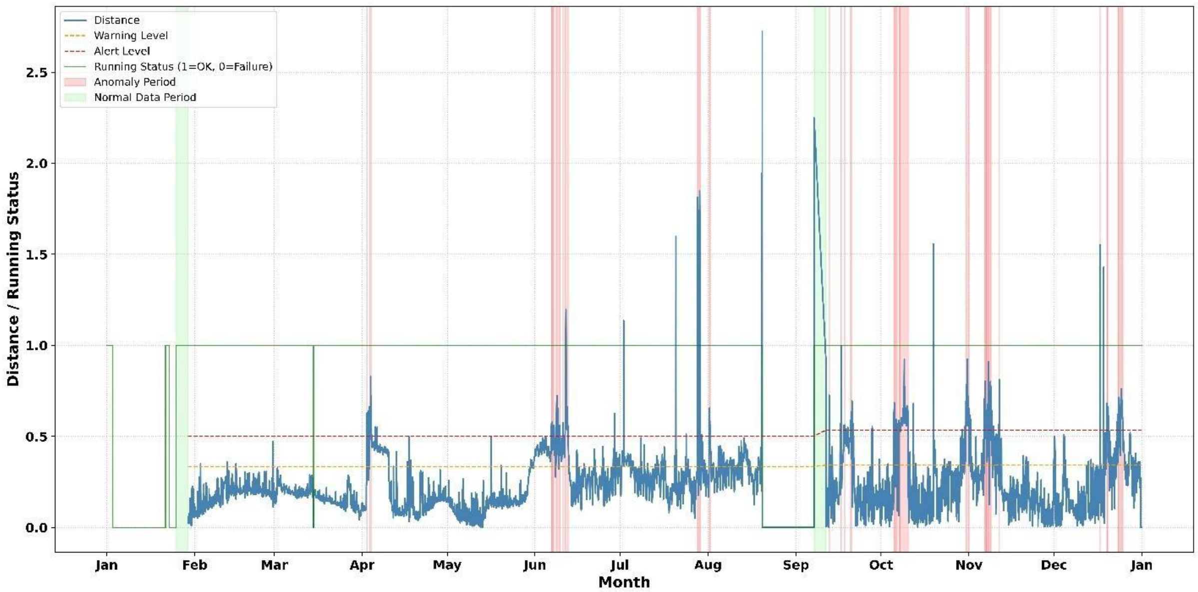

Fig. 10. Deviation plot for 2019 showing the persistence of adaptive thresholding across the final validation year.

4.3 Alert Classification and Diagnostic Interpretation

The framework identified all the 44 functional failures in the retrospective validation set. The model achieved 44 true positives, 8 false positives and zero false negatives. This resulted in 100% recall, 84.62% precision and an F1-score of 91.67%. Compared to the average warning lead time of about 72 h, there was a more realistic time frame for action (inspection, load adjustment, etc. pump changeover and planned maintenance).

A 100% recall score is within the range of this dataset, asset and failure definition. It demonstrates failure sensitivity over three years of varying operation in the data available for validation. It is not to be read as a guarantee that the model will identify all the failure modes for all asset classes and or all operating situations. Such a widespread statement needs to be confirmed.

The true positive alerts represented showed that many of these were dominated by a negative flow deviation and discharge-pressure deviation, which is physically consistent with slurry restriction/impeller obstruction. A few false positives were related to either a discharge-temperature excursion or short-lived operational disturbances that did not develop into a functional failure based on the suction-pressure. The following cases show the value of using the combination of anomaly scoring and interpretation of the sensors.

Table 3. Representative alerts, sensor contributions and classifications

Alert	**Date/time range**	**Top sensor**	**Z-score**	**Second sensor**	**Z-score**	**Classification**
1	2018-01-11T00:15Z to 2018-01-11T20:00Z	Discharge Temp	89.33	Suction Pressure	-1.60	False positive
2	2018-01-12T04:30Z to 2018-01-12T12:15Z	Discharge Temp	-20.01	Flow	-2.10	True positive
5	2018-03-19T14:15Z to 2018-03-19T21:45Z	Flow	-18.77	Discharge Pressure	6.41	True positive

6	2018-06-18T21:45Z to 2018-06-19T09:30Z	Flow	-22.93	Discharge Pressure	4.42	True positive
10	2018-11-12T16:00Z to 2018-11-12T19:00Z	Flow	-5.04	Discharge Pressure	3.06	False positive
21	2018-12-26T11:45Z to 2018-12-27T00:30Z	Flow	-6.17	Discharge Pressure	1.87	False positive

The PCA loading structure was found to be different for the pressure flow variables compared to the temperature-related variables, and this was used to make a root cause interpretation of the PCA anomaly score rather than using the anomaly score as a black box output. For instance, Alert 1 was characterized by a false positive caused by temperature deviation in the discharge and Alert 6 was characterized by a true positive, which was mainly due to flow reduction and discharge-pressure change. This behavior is consistent with that expected in engineering of slurry-loop restriction, and it demonstrates the ability of the framework to aid in operational diagnosis.

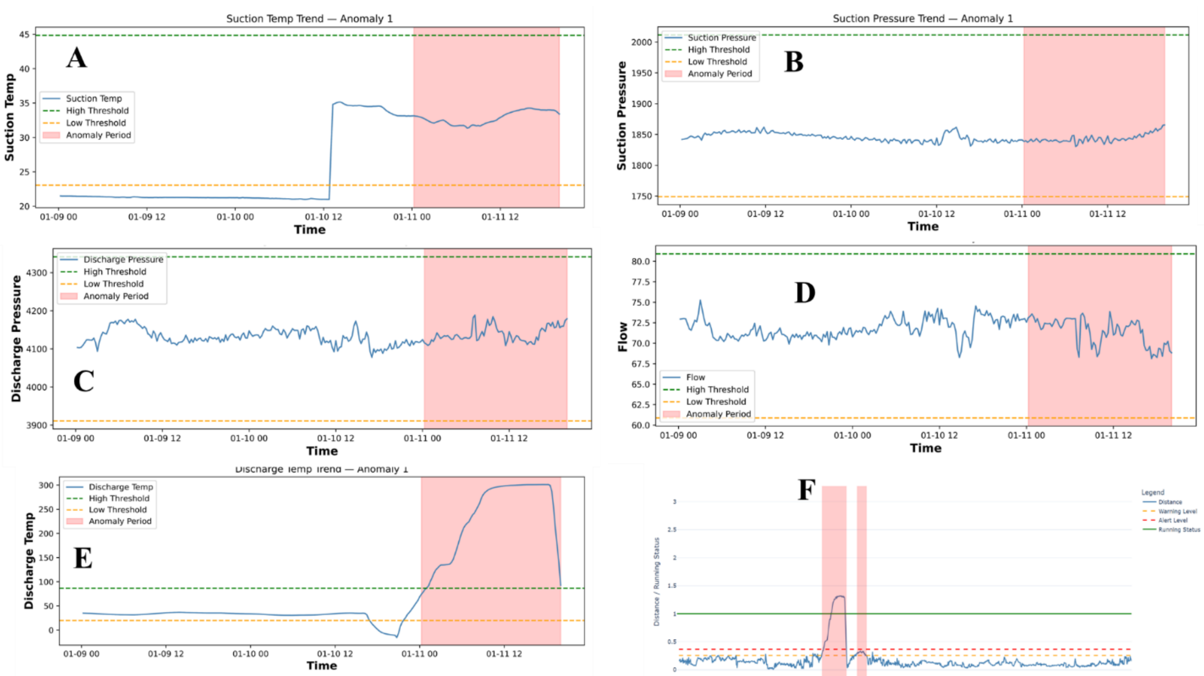

Fig. 11. Detailed sensor behaviour and deviation plot for Alert 1, a false-positive event dominated by discharge-temperature deviation.

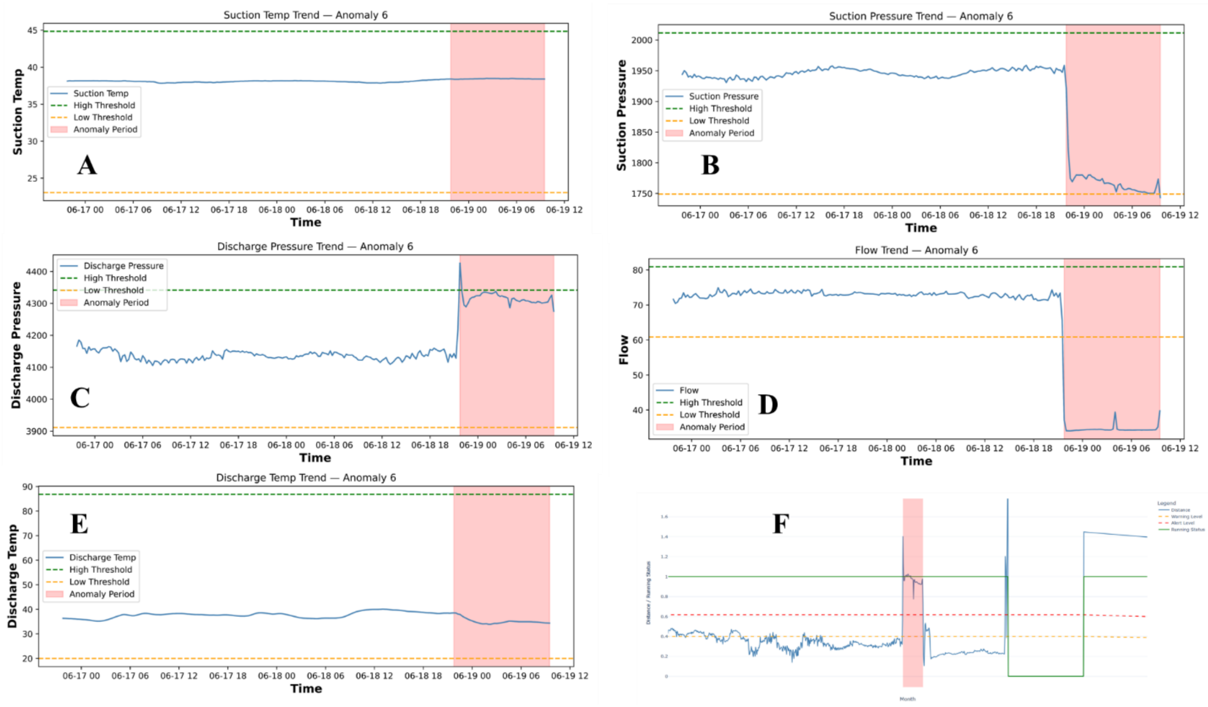

Fig. 12. Detailed sensor behaviour and deviation plot for Alert 6, a true-positive event dominated by flow and discharge-pressure deviations.

5. Discussion

The results indicate that dynamic recalibration is an important component of long-term unsupervised anomaly detection in industrial systems with changing operating regimes. A model mounted on the initial baseline in 2017 would need to be thresholded manually whenever it shifted from one healthy operating state to another; otherwise it would seem anomalous if it changed between normal operating regions. By contrast, the proposed framework uses a framework to determine a new stable operating window after which the reference centroid and alert thresholds are updated, and applying anomaly scores with the present baseline of operation. The results of this work align with previous research which affirms the significance of adaptation within systems that are influenced by varying operating conditions, environmental variations and concept drifts [18,19,32]. The results presented in this paper, however, make it more clear why the recalibration was useful for the present case study, but do not imply that dynamic recalibration is always a requirement for reliability.

Additionally, the framework includes a discussion about the limited availability of fail marked data. Supervised techniques are effective when the training sets are large and representative and labelled data is available, but as we know, large sets of labelled data are not easy to acquire for critical industrial assets as failures are undesirable and failure modes are rare. For this reason, the proposed approach does not need a labelled failure example in the model training process to learn the multivariate structure of normal operation. These functional-failure classifications are made post hoc. The method could be applied to industrial situations where the labels of the faults are either incompleteness, are delayed or are sparsely documented. The PCA- and distance-based architecture is also relatively low in computational load compared to some deep-learning and hybrid based on digital twins approaches that could call for larger training sets, more model development and computational resources [20,21,33]. However, this difference between the two should not be considered as a measure of computational performance, but rather as an architectural difference, because computational performance was not directly evaluated in the present study.

All 44 failure instances were found when the dataset was retrospectively validated, giving a recall (sensitivity) of 100% on the evaluated dataset for failures defined by function. This result is especially pertinent to a pump found in association with a reactor heat removal system, where unreported abnormal conditions could mean significant impacts on process continuity and on safety. But the accuracy result should not just be defined as “100% accuracy”, since there is a specific meaning behind “accuracy” in classification, and that also includes true-negative observations. The reported result in contrast shows that there were no failures in the retrograde validation set that were not functionally defined. Furthermore, it is based on one pump system, one HDPE slurry-polymerisation process, one functional-failure criterion, and one set of historical validations. It thus serves as an indicator of high sensitivity in the case investigated and should not be considered as a blanket assurance that failure detection is possible in other assets, processes or failure mechanisms.

The interpretations of the eight false-positive alerts are also critical. These alerts were asserted when the statistically significant deviations occurred in an operating state which was not associated with the determined functional failure point for suction pressure being below 1200 psi. Other were linked to transient temperature or pressure deviations, which indicated that the structure was reacting to real process deviations observed even though the deviations were not defined as functional failures using the definition provided. It would therefore be premature to claim any of the eight alerts as meaning-less false alarms - or to attribute them to being genuine warnings of a failure - without further operational data. On a reliability standpoint, only a few of these alerts may be tolerable provided the primary purpose is to ensure that critical failures are detected highly sensitively. However, too many alerts that are not actionable may result in operator workload and loss of confidence in the monitoring system. Transient process disturbances must then be separated from continued equipment degradation by adding process-mode tags, maintenance work orders, operator annotations and the like to future implementations.

The sensor-contribution layer makes it easier to understand and use the anomaly-detection framework. The Euclidean anomaly score measures how far the observed multivariate data are from the current multivariate normal baseline, but it cannot identify which of the process variables is responsible for the anomaly. An engineering-based ranking of sensor contributions offers a metric for the relationship between the determinant variables of an alert, and can contribute to the diagnostic workup. Capture logic anomalies that were true-positives in the analysed alerts are consistent with the likely effects of the slurry restriction or potential impeller obstruction. Non-failure alerts, on the other hand, were linked with the rheology deviation tag of discharge temperature, suggesting that process perturbations can lead to statistically unusual process conditions without breaking the functional-failure threshold [38,39]. These contribution rankings are therefore not necessarily indicative of root cause. If a particular failure mechanism is confirmed, supporting evidence from maintenance records, inspections, operator observations and/or other condition-monitoring measures would be required.

The second contribution of the study is that the developed workflow for reliability monitoring is rather simple in its general structure. The framework does not demand a signature (or labelled training examples) to be defined manually per each failure mode. On the contrary, the implementation must be done with a set of variables from the sensors and since the sensor variables differ they do require some kind of data quality screening, must have a stable normal operating region for these data (representative) and must have a functional failure criterion defined independently for validation [34,35]. These features suggest that the same modelling architecture may be transferable to other industrial assets, subject to asset-specific engineering assessment, recalibration, and validation. Selection criteria for the sensor, operating regimes, baseline definition, recalibration criteria and sensor thresholds and validations are still asset/process dependent. Therefore, there would be extra engineering assessment and validation required for deployment on a different pump or equipment class [37].

The findings should be viewed with a certain precaution. First, only one P-101 A/B pump system in one single HDPE slurry-polymerisation process was validated. Although the methodology was developed as a generic framework, the present study does not establish its generalizability or scalability beyond the investigated pump system because validation was limited to a single industrial case. Second, operationally, failure was defined as suction pressure < 1200psi. This criterion is repeatable for later assessment, but may not be used in determining whether different physical failure mechanisms are present or whether each successful deciphering signifies the presence of polymer lumping or impairment in the impellers. More research should then be performed to correlate anomaly alerts with maintenance work orders, inspection items, replaced components, operator reporting, and confirmed failure modes, in the future.

Thirdly, motor current was rejected due to sensor reliability concerns. However, the removal of this piece of information about pump loading and electromechanical behaviour may be unfortunate, so future deployments will need the evaluation of better measures of current. Fourthly, the study was retrospective. Retrospective results have shown that the framework can highlight the past deviations, however they cannot confirm that prospective application will lead to better maintenance decisions, lower unplanned downtime or better process safety. Such outcomes must be judged during an actual industrial deployment.

Another methodological limitation is the dependence on the assumptions when determining the threshold values and dimensionality reduction. When the distribution of anomalies-scores is reasonably close to the assumed distribution, Gaussian-based thresholds can be used in a convenient fashion, although this is not true in all operating regimes. Similarly, if only the first two PCs are used for anomaly scoring, the resulting score can be interpreted and captures a good amount of the baseline variance, but it may be possible to have anomalies (remove background noise) in lower-variance directions [36]. For future work, Gaussian thresholding should be compared with statistically appropriate thresholding procedures and empirical thresholding by the quantiles and the number of the retained principal components should be explored as the number of principal components decreases.

Future studies should test hypotheses that can be directly tested with industrial data, while larger conclusions are not definitive and remain beyond the scope for the data used in the present study. The framework should first be validated on several pumps, and then on other types of equipment, including compressors, heat exchangers and rotating components, among others. To establish if an alert is a real degradation or a transient process disturbance, the maintenance histories, operator feedback and process-mode information together with relevant operating-condition tags should be combined. Alternative thresholding strategies should also be compared under non-Gaussian baseline distributions and quantify the impact of PCA component-retaining criteria.

Finally, prospective studies should evaluate the framework in near real-time operational settings and not just the statistical detection performance, but also operational measures like warnings issued, false-calling ratio, maintenance actions taken, downtime avoided, computational demands, and operator response. Afterwards, integration into edge, cloud and/or digital-twin architectures can even be considered as a deployment paradigm, rather than an expectation of better detection performance. For broader conclusions to be drawn, the proposed adaptive anomaly-detection framework should be validated across multiple assets, operating environments, and prospective datasets. The present findings therefore demonstrate potential transferability rather than experimentally validated scalability or generalizability.

6. Conclusion

This study presents and tests a light-weight, generic and dynamically calibrated unsupervised anomaly-detection framework for monitoring the reliability of industrial assets. The framework combines operating-range normalisation, PCA based dimensionality reduction, euclidean-distance anomaly evaluation, gaussian anomaly based warning and alert thresholds, sensor-contribution anomaly rank, and finally dynamic recalibration of baseline in case of changes in normal operating condition.

The critical P-101 A/B pump in an HDPE slurry polymerisation process was used for evaluating the framework with three years of historical data. The framework achieved a perfect recall (100%) and a high precision (84.62%) and F1 (91.67%) when tested on the 44 functionally defined failure cases in the retrospective validation dataset, as well as an average warning lead time of 72 h. The results in this case study clearly support the claim that the framework is capable of capturing all functional failures recorded under the condition and failure definition studied here, without requiring to create labelled failure examples for model training.

One of the key findings of the study is the incorporation of the dynamic baseline recalibration algorithm with an easily interpretable and efficient unsupervised monitoring architecture. The framework follows a historical baseline that doesn't rely on constant reference and adapts the baseline level each time a stable operating window is identified during the monitoring period. The sensor-contribution layer can be further used to identify the most associated process variables with each alert and help in interpretation. These contributions serve as a practical reference for the support of the engineering investigation and maintenance prioritisation. Observe however, that the contribution of the sensor should not be treated as a definitive conclusion of a root cause.

The findings should be interpreted within the limitations of the study. The reported performance is based on the performance of one particular pump system, one particular industrial process, one particular history of performance data, and a functional-failure definition recognizing suction pressure levels below 1200 psi. The 100% recall here is not meant to imply that the same level of recall is guaranteed for any other asset, operating conditions, or mode of failure. Likewise, the structure of the framework is simpler to compute than some deep-learning methods, but computational performance has not been specifically compared to other methods; so, gathering that it is faster to compute than other deep-learning methods at this time cannot be asserted.

For future studies, prospective and mult-asset validation should be performed. The framework should also be tested for other equipment categories and/or with other pumps, incorporating information from work orders, inspection reports, operator feedback, and known failures in the system to validate. Future studies should also compare Gaussian based thresholds with other empirically appropriate and/or statistically proper thresholding methods, evaluate sensitivity on the number of retained PCs, and quantify operating results (false alert burden, warning lead time, avoided downtime, maintenance actions, computational burden, etc.). More generally, there is a need for such prospective validation prior to making wider claims about the generalisability, scalability and safety of the proposed framework in real-time industrial applications.

An additional direction for future research is to integrate environmental variables, such as ambient temperature, humidity, seasonal conditions, feedstock characteristics, and process-mode indicators, with the sensor-based anomaly-detection model. As illustrated in Fig. 13, this workflow would assess whether a detected deviation is more likely to reflect equipment degradation or environmental and operational variation. The resulting environmental-influence validation could support recalibration when the influence is negligible, or reassessment of environmental coupling when it is significant. This extension should be evaluated prospectively to determine whether it improves model robustness and reduces non-actionable alerts.

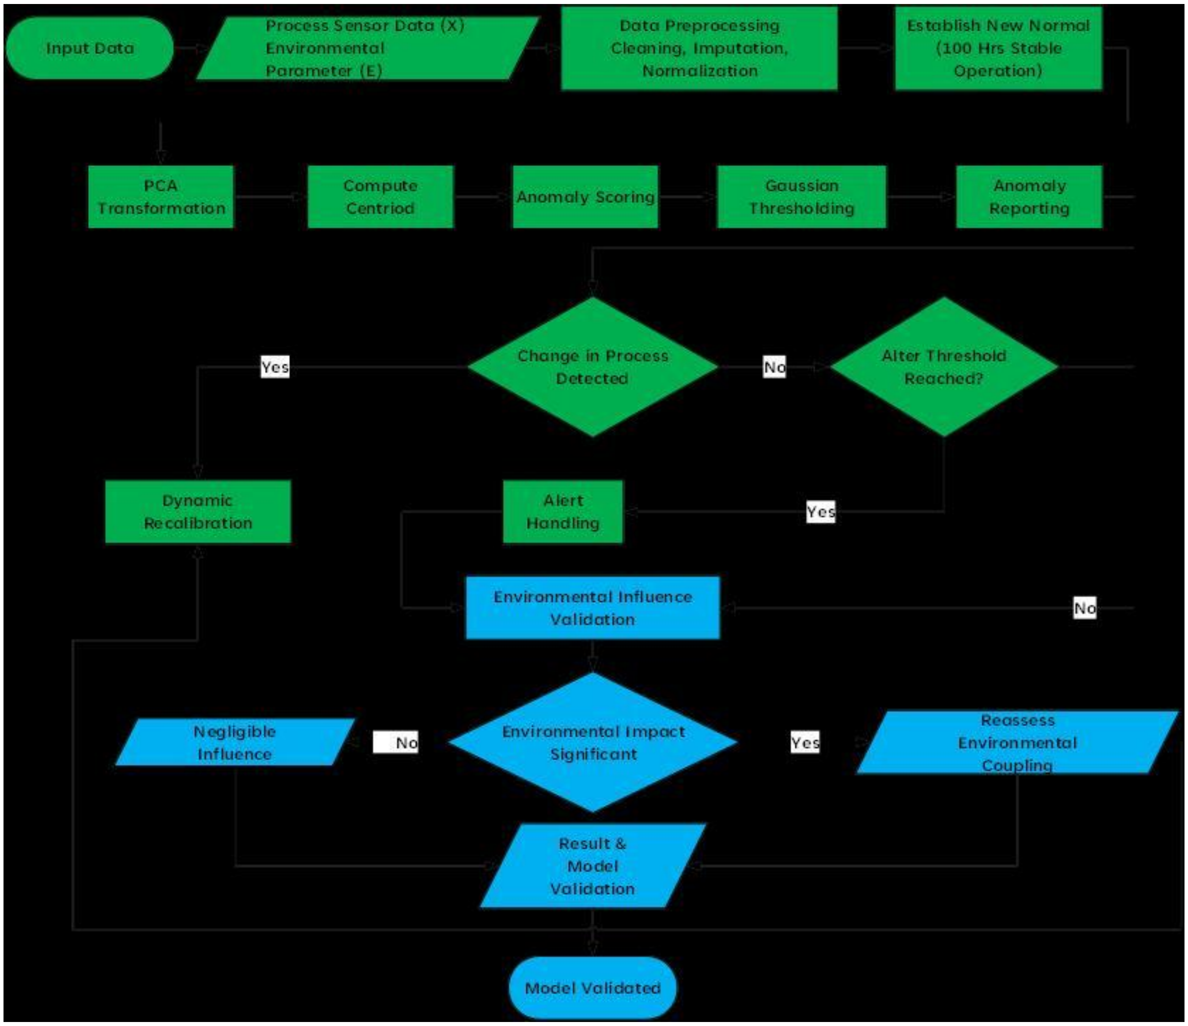

Fig. 13. Proposed future workflow for integrating environmental variables into dynamic recalibration, anomaly detection, and model-robustness validation.

Declarations

Funding: This work has not been supported by any specific grant funding from public, commercial or not-for-profit funding body.

Declaration of Competing Interest: The authors have no known competing financial interests or personal relationships that could have a potential or apparent influence on the work reported in this paper.

Data availability: The information in this study is from confidential industrial historian records. The aggregated data and processed data that support the conclusions of the work presented in this paper, obtained by the corresponding author upon reasonable request and with institutional and industrial approval.

Ethics Approval: Not applicable. No human, animal or personal data was involved in this study. It was derived from information from the sensors on industrial equipment.

Author Contributions: D.S.: Conceptualization, Methodology, Data curation, Software, Formal analysis, Investigation, Visualization, Writing - original draft. Writing - review and editing, D.P.P.: Supervision, Methodology, Validation, Project administration. M.K.: Conceptualization, Methodology, Supervision, Writing - review and editing. All the authors have read and approved the final manuscript.

Acknowledgements: The authors appreciate the assistance of the plant personnel in collecting and interpreting the data used for developing the plant historians.

Highlights

- A lightweight unsupervised approach that can identify anomalies in the pump without failure data labels.
- Dynamic recalibration automatically adjusts thresholds with changes in operating baseline.
- Long-horizon applicability has been validated with three years of petrochemical pump historian data.
- All 44 of the functional failures found were in the retrospective validation set.
- Average warning lead time of 72 h allows for decisions to be made based on reliability of maintenance.

Appendices

Code	Meaning	Unit
X1	Motor current	A
X2	Pump DE vibration	m/s^2 pk
X3	Pump NDE vibration	m/s^2 pk
X4	Pump hexane flush inlet	m^3/h
X5	Seal pot level	%
X6	Pump flow	m^3/h
X7	Suction temperature	°C
X8	Suction pressure	psi
X9	Discharge pressure	psi
X10	Discharge temperature	°C

So, for example, **X6** represents pump flow, **X8** represents suction pressure, and **X9** represents discharge pressure. The manuscript also states that **X1 motor current was monitored but excluded from the PCA model due to sensor reliability issues**.

REFERENCES

1. Ahmed Murtaza A, Saher A, Zafar MH, Moosavi SKR, Aftab MF, Sanfilippo F. Paradigm shift for predictive maintenance and condition monitoring from Industry 4.0 to Industry 5.0: A systematic review, challenges and case study. Results Eng 2024;24:102935. https://doi.org/10.1016/j.rineng.2024.102935.
2. Ahmed S, Terefe T, Hailemariam D. Machine learning for base transceiver stations power failure prediction: A multivariate approach. e-Prime Adv Electr Eng Electron Energy 2024;100814. https://doi.org/10.1016/j.prime.2024.100814.
3. Alsulieman A, Ge X, Zeng Z, Butenko S, Khan F, El-Halwagi M. Dynamic risk analysis of evolving scenarios in oil and gas separator. Reliab Eng Syst Saf 2024;243:109834. https://doi.org/10.1016/j.ress.2023.109834.
4. Wahono T, Purniawan A, Mukhlash I, Putri E. Risk-based asset integrity management in the oil and gas industry from traditional to machine learning approaches: A systematic review. Results Eng 2025. https://doi.org/10.1016/j.rineng.2025.107287.
5. Al-Sabaeei AM, Alhussian H, Abdulkadir SJ, Jagadeesh A. Prediction of oil and gas pipeline failures through machine learning approaches: A systematic review. Energy Rep 2023;10:1313-1338. https://doi.org/10.1016/j.egyr.2023.08.009.
6. Cabarcos A, Paz C, Suarez E, Vence J. Application of supervised learning algorithms for temperature prediction in nucleate flow boiling. Appl Therm Eng 2024;240:122155. https://doi.org/10.1016/j.applthermaleng.2023.122155.
7. Kazim M, Pirim H, Le C, Le T, Yadav O. Multilayer GNN for predictive maintenance and clustering in power grids. iScience 2025. https://doi.org/10.1016/j.isci.2025.113490.

8. Duan D, Han S, Wang Z, Pang C, Yao L, Liu W, et al. Multivariate state estimation-based condition monitoring of slurry circulating pumps for wet flue gas desulfurization of power plants. Eng Fail Anal 2024;159:108099. https://doi.org/10.1016/j.engfailanal.2024.108099.
9. Wang K, Yuan X, Cao Z, Chen G, Yuan X, Yang C, et al. Capturing principal features in slow industrial processes for anomaly detection application. J Process Control 2025;153:103487. https://doi.org/10.1016/j.jprocont.2025.103487.
10. Kim J, Cho S-B. Unsupervised video anomaly detection by memory network with autoencoders in Euclidean and non-Euclidean spaces. Pattern Recognit 2025;167:111759. https://doi.org/10.1016/j.patcog.2025.111759.
11. Morsi I, Mohy El-Din L. SCADA system for oil refinery control. Measurement 2014;47:5-13. https://doi.org/10.1016/j.measurement.2013.08.032.
12. Xiao H, Cao R, Chen Z, Hong C, Wang J, Yao M, et al. Handling missing data in large-scale TBM datasets: Methods, strategies, and applications. Intell Geoeng 2025;2:109-125. https://doi.org/10.1016/j.ige.2025.07.001.
13. IEC 60050. International Electrotechnical Vocabulary. Geneva: International Electrotechnical Commission; 1990.
14. Nicosia GP. Machine learning, optimization, and data science. Cham: Springer Nature; 2020.
15. Wang X, Chen Y. Unsupervised detection of multivariate geochemical anomalies using a high-performance deep autoencoder Gaussian mixture model. J Geochem Explor 2025;271:107671. https://doi.org/10.1016/j.gexplo.2025.107671.
16. López-Flores FJ, Ochoa-Barragán R, Raya-Tapia AY, Ramírez-Márquez C, Ponce-Ortega JM. Machine learning with Python. In: Machine Learning Tools for Chemical Engineering: Methodologies and Applications. Elsevier; 2025. p.31-55. https://doi.org/10.1016/B978-0-443-29058-9.00004-5.
17. Crespo Márquez A, Serra Parajes J, de la Fuente Carmona A, Sola Rosique A. Integrating complex asset health modelling techniques with continuous time simulation modelling: A practical tool for maintenance and capital investments analysis. Comput Ind 2021;133:103507. https://doi.org/10.1016/j.compind.2021.103507.
18. Ben Haha MA, Bohli A, Haddour N, Bouallegue R. A systematic review of industrial IoT anomaly detection and the forensic interpretability gap. Electronics 2026;15(11):2240. https://doi.org/10.3390/electronics15112240.
19. Subashree S, Rajakumaran M, Pushpa G, Manivannan R. Adaptive machine learning models for predictive maintenance in industrial internet of things (IIoT) systems. Sci Rep 2026;16:12451.
20. Mavrelis I, et al. Hybrid deep learning for predictive maintenance in industrial machinery using LSTM and MLP models. Machines 2026;14(2):191.
21. Khan T, et al. Data-driven digital twin framework for predictive maintenance of smart manufacturing systems. Machines 2025;13(6):481. https://doi.org/10.3390/machines13060481.
22. El Hammoumi Z, et al. Predictive maintenance approaches: A systematic literature review. Eng Proc 2025;112(1):70.
23. Elkateb S, Métwalli A, Shendy A, Abu-Elanien AEB. Machine learning and IoT-based predictive maintenance approach for industrial applications. Alex Eng J 2024;88:298-309. https://doi.org/10.1016/j.aej.2023.12.065.
24. Gao J, Liu Y, Zhao Q, Ai X, Shan L. Research on sensor anomaly detection and fault location in nuclear power plants based on GAN-IAAKR model. Ann Nucl Energy 2025;219:111462. https://doi.org/10.1016/j.anucene.2025.111462.
25. Guo L, Gao J, Yang J, Kang J. Criticality evaluation of petrochemical equipment based on fuzzy comprehensive evaluation and a BP neural network. J Loss Prev Process Ind 2009;22:469-476. https://doi.org/10.1016/j.jlp.2009.03.003.
26. Mohammed A. Data driven-based model for predicting pump failures in the oil and gas industry. Eng Fail Anal 2023;145:10719. https://doi.org/10.1016/j.engfailanal.2022.107019.
27. Ratnayake RM. Application of a fuzzy inference system for functional failure risk rank estimation: RBM of rotating equipment and instrumentation. J Loss Prev Process Ind 2014;29:216-224. https://doi.org/10.1016/j.jlp.2014.03.002.
28. Wu S, Yang P, Li D, Chen G, Wang Z. Towards predictive maintenance: A performance evaluation framework for cooling towers in HVAC systems. Build Environ 2025;284:113443. https://doi.org/10.1016/j.buildenv.2025.113443.
29. Xie L, Håbrekke S, Liu Y, Lundteigen MA. Operational data-driven prediction for failure rates of equipment in safety instrumented systems: A case study from the oil and gas industry. J Loss Prev Process Ind 2019;60:96-105. https://doi.org/10.1016/j.jlp.2019.04.005.
30. Yucesan YA, Dourado A, Viana FA. A survey of modeling for prognosis and health management of industrial equipment. Adv Eng Inform 2021;50:101404. https://doi.org/10.1016/j.aei.2021.101404.
31. Dereci U, Tuzkaya G. An explainable artificial intelligence model for predictive maintenance and spare parts optimization. Supply Chain Anal 2024;8:100078. https://doi.org/10.1016/j.sca.2024.100078.
32. Akil, M., Martucci, L., & Hoepman, J.-H. (2023). Non-interactive privacy-preserving Sybil-free authentication scheme in VANETs. In Proceedings of the Symposium on Vehicle Security and Privacy (VehicleSec) 2023 (pp. 27–35). Internet Society. https://doi.org/10.14722/vehiclesec.2023.23007
33. Akil, M., Martucci, L., & Hoepman, J. (2024). A secure and privacy-preserving authentication scheme with a zero-trust approach to vehicle renting in VANETs. In Proceedings of the 21st International Conference on Security and Cryptography (SECRYPT 2024) (Vol. 1, pp. 114–127). SciTePress. https://doi.org/10.5220/0012759600003767
34. Akil, M., Naskar, S., Martucci, L. A., & Hoepman, J. (2024). A Privacy-Preserving approach to vehicle renting and driver accountability in VANETs. *IFIP Advances in Information and Communication Technology*, 192–210. https://doi.org/10.1007/978-3-031-57978-3_13
35. Akil, M., Mancini, L., V., & Venturi, D. (2019). *Multi-covert channel attack in the cloud*. DIVA. https://kau.diva-portal.org/smash/record.jsf?faces-redirect=true&aq2=%5B%5B%5D%5D&af=%5B%5D&searchType=SIMPLE&sortOrder2=title_sort_asc&language=

sv&pid=diva2%3A1369702&aq=%5B%5B%5D%5D&sf=all&aqe=%5B%5D&sortOrder=author_sort_asc&onlyFullText=false&noOfRows=50&dswid=-623

36. Bello, E. A., Estrada, U. L., Hernández, H. H., Ortiz, S. Y. C., & Vázquez, J. T. (2024). *Aplicación Web para el Monitoreo de Conversión de Agua a Materia Seca para el Desarrollo en Diferentes Genotipos de Soya*. Dialnet. https://dialnet.unirioja.es/servlet/articulo?codigo=9430317
37. Ortiz, S. Y. C., Vázquez, J. T., Bello, E. A., Estrada, U. L., & Erdosay, J. S. (2024). *La Soya: Tesoro versátil, proteico y saludable importante en la dieta mexicana*. Dialnet. https://dialnet.unirioja.es/servlet/articulo?codigo=9787138
38. Inakollu, M. V. (2025). Quantum-Resilient architectures for enterprise and cloud information systems, implications of quantum computing for enterprise cybersecurity and data integrity. *www.periodicos.ulbra.org*. https://doi.org/10.22178/acta.26.3.44
39. Inakollu, Manikantha Varaprasad. (2022). Decision Intelligence As A Strategic Capability: A New Mis Framework For AI-Driven Enterprises, A Comprehensive Approach To Integrating Artificial Intelligence Into Organizational Decision-Making. Acta Scientiae. 23. 66-79.