Modern Anomaly Detection Methods in Industry: A Comparative Analysis of Machine Learning Algorithms and Their Application to Improve the Efficiency of Manufacturing Processes

Kanan Mikayilov and Latafat Gardashova

Department of Computer Engineering, Azerbaijan State Oil and Industry University, Baku, Azerbaijan

Correspondence to:
Kanan Mikayilov,
kananmikayilov3@gmail.com

Additional material is published online only. To view please visit the journal online.

Cite this as: Mikayilov K and Gardashova L. Modern Anomaly Detection Methods in Industry: A Comparative Analysis of Machine Learning Algorithms and Their Application to Improve the Efficiency of Manufacturing Processes. Premier Journal of Science 2025;14:100154

DOI: https://doi.org/10.70389/PJS.100154

Peer Review

Received: 4 September 2025

Last revised: 8 October 2025

Accepted: 8 October 2025

Version accepted: 3

Published: 18 October 2025

Ethical approval: N/a

Consent: N/a

Funding: No industry funding

Conflicts of interest: N/a

Author contribution:
Kanan Mikayilov – conceptualization, methodology, data curation, writing – original draft preparation
Latafat Gardashova – visualization, investigation, supervision, software, validation, writing – reviewing and editing.
All authors read and approved the final manuscript

ABSTRACT

BACKGROUND

The growing complexity of industrial systems and the large volume of operational data have increased the demand for automated anomaly detection to ensure production efficiency and stability. Machine learning methods provide promising solutions for identifying process deviations, but their comparative effectiveness in manufacturing environments remains insufficiently evaluated.

MATERIALS AND METHODS

A comparative analysis of multiple machine learning algorithms was conducted for anomaly detection in industrial time-series systems. The methods included unsupervised models (Autoencoder, Isolation Forest), supervised classifiers (Random Forest, Support Vector Machine), and density-based approaches (Local Outlier Factor), with a focus on evaluating their performance in high-dimensional, noisy environments typical of industrial datasets.

RESULTS

Unsupervised algorithms, particularly Autoencoder (87–89% accuracy) and Isolation Forest (84–86%), showed strong performance in environments without labeled data, making them suitable for real-world applications. Supervised classifiers achieved the highest accuracy (Random Forest: 89–91%; SVM: 88–90%) but were constrained by the availability of labeled datasets. Density- and clustering-based methods performed moderately (70–80%) in low-dimensional data but lost effectiveness as feature dimensionality increased. Implementation of machine learning-based monitoring systems demonstrated potential reductions in downtime (up to 29%), repair costs (20%), and significant improvement in productivity.

CONCLUSION

Machine learning methods are effective tools for anomaly detection in manufacturing systems. Algorithm selection should be tailored to the availability of labelled data, system complexity, and processing constraints. Adaptive anomaly detection frameworks provide practical benefits for industrial process control, supporting efficiency, cost reduction, and operational stability. The research reduces operational costs while simultaneously enhancing real-time monitoring, thereby improving decision-making and overall system reliability in industrial environments.

Keywords: Cross-validation, Dimensionality reduction, Hyperparameter optimization, Intelligent monitoring, Noise robustness, Unsupervised models

Highlights

- Machine learning methods were comparatively evaluated for anomaly detection in industrial manufacturing systems using real and synthetic datasets.
- Unsupervised models (Autoencoder, Isolation Forest) achieved high accuracy (84–89%) in the absence of labeled data, suitable for real-world applications.
- Supervised classifiers (Random Forest, SVM) reached the highest accuracy (88–91%) but were limited by data availability.
- Intelligent monitoring systems based on these methods reduced downtime by up to 29%, cut repair costs by 20%, and improved productivity.

Introduction

The intensive development of industrial technologies and the widespread introduction of automated control systems have led to a rapid increase in the volume of data generated during production processes. Along with this, the complexity of technical systems has increased, and the sensitivity of production chains to any deviations from the norm has also increased. In the context of high competition and the need to ensure the continuity of technological operations, there is an urgent need for reliable mechanisms for the early detection of non-standard behaviour of equipment and systems.

Anomaly detection in industry has evolved from methods based on fixed thresholds and expert systems to more sophisticated approaches using machine learning.[1] Unlike traditional methods, which are limited by rigid criteria, machine learning offers dynamic and adaptive solutions, allowing for more efficient processing of large amounts of data and real-time detection of deviations. It is important to emphasise that in the context of Industry 4.0, machine learning allows data from various sensors and systems to be integrated to create flexible solutions for anomaly detection. Anomaly detection has become essential in industrial systems owing to the growing complexity and interconnectivity of contemporary industrial processes. As industries embrace advanced technologies such as automation, IoT, and machine learning, the influx of operational data has increased, complicating the manual identification of errors or anomalies. Prompt detection of these anomalies is crucial for decreasing downtime, enhancing safety, and lowering operational expenses, especially as systems become increasingly dynamic and interdependent. Effective anomaly detection facilitates early defect identification and anticipates

Guarantor: Kanan Mikayilov
Provenance and peer-review: Unsolicited and externally peer-reviewed
Data availability statement: N/a

probable failures, hence improving the reliability and resilience of industrial operations.

Elía and Pagola[2] mapped current approaches to manufacturing challenges, addressing the need for context-aware and scalable solutions. The study underscored the importance of integrating domain knowledge with data-driven methods to improve the interpretability and reliability of anomaly detection systems. This problem has become especially relevant in the context of the transition to the concepts of digital production and smart factories, where the integration of sensor data, telemetry and analytics algorithms is becoming the basis for the effective operation of industrial facilities.

Shaikh et al.[3] conducted a comprehensive review of machine learning techniques applied in smart manufacturing. The study identified key algorithms and categorised their applicability based on process types and operational needs. It was noted that supervised learning was predominant in applications related to quality prediction, whereas unsupervised and reinforcement learning methods were more frequently utilised in dynamic environments that necessitate real-time adaptability. Against the background of the growing importance of timely detection of deviations from the normal state of equipment and processes, traditional approaches based on fixed thresholds or expert rules have shown limited flexibility and scalability.

Zare et al.[4] provided a comprehensive overview of deep learning-based anomaly detection in cyber-physical systems (CPS). The review examined convolutional, recurrent, and autoencoder architectures in various CPS contexts. It was found that deep learning models excelled in capturing complex spatial-temporal correlations in multivariate data, though their black-box nature posed interpretability concerns. In this regard, special attention was paid to machine learning methods that can adapt to changing conditions, take into account complex dependencies between parameters and ensure high accuracy when working with a large number of variables.

Desani and Chittibala[5] proposed adaptive machine learning models for real-time anomaly detection in streaming data. The research focused on evolving data patterns in dynamic systems, introducing online learning techniques that could update models on the fly. The findings highlighted that adaptive models enhanced detection speed and accuracy, particularly in scenarios involving concept drift. Two key areas in this area were identified – supervised and unsupervised learning, each of which demonstrated certain advantages depending on the data structure and the goals of the analysis.

Kaur and Ranjan[6] performed a comparative analysis of supervised and unsupervised machine learning algorithms for fake news detection. Metrics such as accuracy, efficiency, and robustness were used to assess the performance of models including decision trees, and Support Vector Machine (SVM). It was determined that supervised methods provided higher accuracy but required labelled datasets, whereas unsupervised approaches were more adaptable but less precise.

Singh et al.[7] conducted a comparative study of anomaly detection and diagnosis techniques in manufacturing systems, analysing statistical, machine learning, and deep learning approaches. The research highlighted that traditional statistical methods were efficient in detecting simple anomalies, while machine learning offered better generalisation. Deep learning techniques, especially those using recurrent and convolutional neural networks, demonstrated superior performance in handling high-dimensional data and complex fault patterns.

Yan et al.[8] surveyed deep transfer learning methods for anomaly detection in industrial time series. The authors classified transfer learning approaches into feature-based, parameter-based, and instance-based techniques. The findings emphasised the importance of domain adaptation, as transferring knowledge from similar processes significantly enhanced detection accuracy in limited-data scenarios.

Velásquez et al.[9] proposed a hybrid ensemble machine learning framework for real-time anomaly detection in Industry 4.0 systems. The system combined decision trees, SVMs, and deep learning models to exploit their complementary strengths. The hybrid model demonstrated superior performance in identifying anomalies in industrial sensor data streams, reducing false alarms and improving detection latency.

Adapting machine learning algorithms to changing conditions (known as concept drift) is key to effective application in dynamic industrial environments.[10] Models that are capable of adjusting their predictions in real time based on changing data patterns ensure that systems are resilient to external and internal disturbances, such as changes in operating conditions or equipment failure. This theoretically contributes to improved anomaly detection accuracy and fewer false positives, which is important for preventing production losses.

This study differentiates itself from prior research on anomaly detection in industrial contexts by employing an experimental design that assesses a wider array of models under authentic industrial conditions, emphasising data scale (up to 468,000 data points) and utilising synthetic datasets to simulate various anomalies, including sensor failures and communication disruptions. This facilitates a more thorough comprehension of the performance of these models in demanding contexts. Moreover, the incorporation of delay analysis and real-time operational feasibility introduces an industrial significance that previous research has not fully addressed.

The study aimed to provide a systematic comparative evaluation of state-of-the-art machine learning algorithms for anomaly detection in production environments and to assess their relative effectiveness in enhancing industrial process efficiency. To achieve this goal, the following tasks were set: to analyse and classify the most common algorithms for detecting deviations, to implement an experimental comparison of their effectiveness based on open and synthetically generated data, and to determine the conditions under which the use of each method is most justified.

Materials and Methods

Data collection and preparation

The study used two types of source data. Three open datasets were used, reflecting different production scenarios and typical deviations in the operation of the equipment. Tennessee Eastman Process (TEP) was downloaded from the UCI Machine Learning Repository, which contained 4800 observations with 52 variables of process parameters.[11,12] TEP dataset simulates the operation of a chemical production facility, where process parameters such as temperature, pressure, and flow rate are critical for anomaly detection. Secure Water Treatment (SWaT) data from the iTrust platform included 468,000 records with 51 variables obtained during the operation of the water treatment system.[13] The SWaT dataset, on the other hand, models the water treatment process, where sensor failures and abnormal events can lead to serious system malfunction. These datasets were selected to represent different industrial sectors and evaluate the robustness of the anomaly detection models under various operational conditions. To test the robustness of the models to different types of anomalies, the Numenta Anomaly Benchmark (NAB) from the official GitHub repository was used, consisting of 58 time series with an average length of about 4000 points and 15 variables.[14,15] The data was collected taking into account real industrial conditions, which made it possible to model practical operating scenarios. A physics-statistical simulation of the production process was performed using the TEP model and the Monte Carlo method to generate time series of key parameters (temperature, pressure, flow rate).

Each variable was varied within ±15% of its nominal value following a normal distribution, and discrete disturbances lasting from 5 to 30 s with amplitudes up to 25% of the baseline were introduced to reproduce anomalies. To mimic sensor failures, Gaussian noise ($\sigma = 0.05$) and random packet loss with a 2% probability were added, reflecting realistic communication interruptions and measurement instability. The resulting synthetic samples provided sufficient scenario diversity for subsequent experimental analysis, enabling assessment of algorithm robustness to noise and parameter variation and influencing comparative anomaly-detection accuracy metrics. Synthetic data were enriched with artificially introduced noise factors, which made it possible to model sensor failures and communication system disruptions.

At the preliminary stage, data normalisation and standardisation were carried out, as well as the removal of obvious outliers, which made it possible to minimise the impact of abnormal values on further analysis. Principal Component Analysis was used on datasets containing over 20 characteristics, preserving components that accounted for 95% of the variation to assure consistency. Feature selection included the exclusion of variables exhibiting a variance inflation factor over 10 to mitigate multicollinearity, while Min-Max scaling was used for normalisation across all datasets (SWaT, NAB, TEP). Outliers were eliminated by removing values over three standard deviations from the mean. The Autoencoder used a reconstruction error threshold for thresholding, established by grid search to optimise the equilibrium between sensitivity and specificity. The Isolation Forest contamination parameter was optimised by walk-forward validation with values of 0.01, 0.05, and 0.1, while the threshold was established according to anomaly scores to maximise the F1 score. These approaches were uniformly implemented across all datasets (SWaT, NAB, TEP) to guarantee methodological consistency and repeatability. Filtering and elimination of correlated variables were used, which helped to reduce the dimensionality of the data and improve the quality of the input information. As a result, representative samples were formed that met the requirements for realism and diversity of industrial scenarios, which was a necessary condition for an objective assessment of anomaly detection algorithms.

Selection and Implementation of Anomaly Detection Algorithms

Based on the analysis of modern literature and preliminary experiments, the main algorithms used to detect anomalies in an industrial environment were selected. The choice of algorithms was grounded in a targeted review of seminal works, including the comprehensive survey by Elía and Pagola[1] and the detailed classification presented by Shaikh et al.,[2] which highlighted the strengths and limitations of density-based, unsupervised and supervised methods.

At the preliminary stage, small-scale experiments were executed on a representative subset of the Tennessee Eastman and SWaT datasets to calibrate key hyperparameters (e.g., tree counts, learning rates, layer sizes) and to validate correct implementation of each model. Density-based methods (Local Outlier Factor), unsupervised algorithms (Autoencoder, Isolation Forest) and supervised classifiers (Random Forest, SVM) were selected to compare performance under labelled and unlabelled data conditions, processing-speed requirements and noise resilience.

Among the methods based on density analysis, the Local Outlier Factor algorithm was implemented, allowing for the detection of rare and local deviations with minimal preliminary settings. Density analysis was applied to identify rare, localised deviations in multivariate data, with the Local Outlier Factor algorithm estimating each point's relative density to its neighbours. Unsupervised approaches (Autoencoder reconstruction and Isolation Forest isolation) processed unlabelled data, while supervised classifiers (Random Forest and SVM) were trained on labelled samples to achieve high classification accuracy. All models were implemented using standard libraries to ensure experiment reproducibility

Unsupervised learning algorithms were used to work with data that did not have preliminary labelling. In particular, Autoencoder models were implemented, capable of identifying nonlinear dependencies between parameters, and the Isolation Forest method, based on the construction of an ensemble of random trees. Along with this, in the presence of labelled data, supervised learning methods were used, such as Random Forest and SVM, which made it possible to achieve high accuracy in classifying anomalies. Each algorithm was implemented using standard machine learning libraries, which ensured the reproducibility of experiments and the possibility of further modification of the models.

Development of an Experimental Platform and Model Tuning

A specialised software platform was deployed to provide a full cycle of data processing: from preliminary preprocessing to algorithm execution and subsequent result analysis. The platform ran under Ubuntu 20.04 long-term support and was implemented in Python 3.8 using the scikit-learn and TensorFlow libraries. Computing resources comprised a dual-node cluster with two Intel Xeon Silver 4214 processors (12 cores each), 64 GB of random-access memory (RAM), and an NVIDIA Tesla V100 graphics processing unit (GPU) (16 GB video RAM). The GPU was utilised for deep learning models (e.g., Autoencoder), which contributed to reducing the latency, particularly for models requiring intensive computations. This configuration was selected to meet the computational demands of high-dimensional time-series data and to efficiently support the training and evaluation of deep learning models (e.g., Autoencoder) and unsupervised anomaly detection algorithms (e.g., Isolation Forest), particularly under the walk-forward validation approach used in this study.

The processing duration for each observation was assessed during the performance analysis, which is essential for real-time anomaly detection applications. The average delay for each model was assessed by 1000 repeated trials and documented using a high-resolution timer. The Isolation Forest and Autoencoder exhibited an average latency of 48 ms per observation, making them appropriate for use in high-velocity industrial settings where minimal latency is critical. Supervised models like Random Forest and SVM exhibited elevated latencies, with average processing durations above 110 ms per observation in the absence of GPU acceleration, hence limiting their use in real-time environments with stringent latency demands.

Training (80%, ~586,240 observations) and testing (20%, ~146,560 observations) subsets were created chronologically for time-series datasets (such as SWaT, NAB, and TEP) in order to prevent data leaking from later time steps. By ensuring that models are tested on future data and trained on historical data, this method preserves the dataset's temporal integrity. The platform included separate pipeline modules for normalisation, data cleaning, dimensionality reduction, detection algorithm implementation, and parallel computing tools, which accelerated the processing of large volumes of information – a total of over 732,800 records (including 10,000 synthetic observations) under real-time simulation conditions. These pipeline modules were distinct from the machine learning models themselves. Each comprised discrete steps for preprocessing, algorithm execution, and result aggregation, ensuring flexibility of configuration and extensibility of the system architecture.

Particular attention was paid to optimising the hyperparameters of each model. For the Isolation Forest algorithm, such parameters as the number of base trees and the anomaly detection threshold were selected, which increased accuracy and reduced computational complexity. When working with Autoencoder, the network architecture was configured, optimal activation functions were selected and the sizes of hidden layers were determined, which ensured a balance between computational costs and the quality of detecting abnormal samples. Autoencoders identified abnormalities by using the reconstruction of input data. Anomalies were detected by establishing a reconstruction error threshold, whereby data exhibiting elevated reconstruction errors (i.e., those markedly diverging from the learnt patterns) were designated as anomalies. This criterion was determined through a grid search to optimise the balance between model sensitivity and specificity.

Optimisation methods such as walk-forward validation and grid search were used for supervised algorithms (Random Forest, SVM) to minimise the risk of overfitting and increase the generalisation ability of the models. Hyperparameter optimisation was carried out separately for each model using fixed random seeds and walk-forward validation to ensure reproducibility and robust performance estimation. A grid search over the number of base trees (n_estimators ∈{100, 200, 500}) and contamination thresholds (contamination ∈{0.01, 0.05, 0.1}) was carried out for the Isolation Forest model using walk-forward validation, in which the model was tested on future data to respect the temporal dependencies and trained on historical data. The combination that produced the greatest mean F1-score was chosen after a total of 3(n_estimators) × 3(contamination) × number of training-test splits = total iterations were carried out. In Isolation Forest, anomalies were identified by computing anomaly scores that reflect the degree of isolation of data points inside the feature space. The classification threshold was established by testing various contamination levels (e.g., 0.01, 0.05, 0.1) using cross-validation to guarantee the quantity of identified anomalies corresponds with the anticipated contamination rate.

Two encoder-decoder depth variants ([64-32-16-32-64] and [128-64-32-64-128]) and two activation functions (ReLU and tanh) were used to compare architectural configurations for the Autoencoder model. With five repeats per configuration (totalling 2 architectures × 2 activations × 5 runs = 20 training runs), each configuration was trained for up to 100 epochs with early stopping (patience = 10) using walk-forward validation on 80% of the training set and validated on the remaining 20%. Weight initialisation and data shuffling were controlled by a set random seed.[42]

To make sure the model was evaluated exclusively on future data that it would meet in real-time applications, supervised models (Random Forest and SVM) were subjected to grid searches with walk-forward validation. For SVM, C ∈{0.1, 1, 10}, kernel in {linear, rbf} (3 × 2 × 5 × 2 = 60 iterations); for Random Forest, n_estimators ∈{100, 200}, max_depth ∈{None, 10, 20} (2 × 3 × 5 × 2 = 60 iterations). The sets of parameters with the highest average validation accuracy were kept.

Across the platform, models represented the actual trained machine-learning algorithms, whereas pipeline modules referred to distinct preparation and execution components (e.g., normalisation module, dimensionality-reduction module). This distinct division made sure that tuning processes only applied to model hyperparameters and did not change the pipeline's core components. An essential step in the creation of adaptable monitoring systems was made feasible by the setup's iterative character, which allowed for the recording of the effects of parameter modifications on the final performance indicators.

Methods of Performance Evaluation and Analysis of Experimental Data

The efficiency of the implemented algorithms was assessed using a set of criteria. The main evaluation metrics included anomaly detection accuracy, false positive rate, average processing time per observation, and robustness to noise. Accuracy was measured as the proportion of correctly identified abnormal samples relative to the total number of true anomalies in the test data. The false positive rate was determined by calculating the share of normal observations that were incorrectly classified as anomalies. The processing time per observation was defined as the average duration between receiving an input vector and completing the classification procedure, based on 1000 repeated trials and measured using a high-resolution timer under standardised hardware conditions.

To evaluate robustness to noise, artificial perturbations were introduced into the datasets. Specifically, Gaussian noise was applied to all numeric features at varying intensity levels (e.g., low, medium, high), and simulated packet loss was implemented by randomly removing a portion of data entries. These modifications imitated common sensor failures and communication issues in industrial systems. The sensitivity of each algorithm to noisy conditions was assessed by comparing the accuracy scores on clean datasets versus noise-augmented datasets. A smaller decline in performance was interpreted as higher resistance to noise effects. This approach provided a structured basis for assessing algorithm performance under realistic operating conditions and allowed for consistent comparisons across different detection models.

To quantitatively assess the performance of the algorithms, summary tables were created, which provided the indicators of each method. Particular attention was paid to comparing the performance of the methods on low- and high-dimensional data, as well as on samples with varying degrees of correlation between the parameters. Statistical analysis of the results was carried out, including using variance analysis methods, which made it possible to identify significant differences between the approaches. Statistical analysis of the experimental results was carried out to assess the significance of the observed differences in algorithm performance. Variance analysis (ANOVA) was applied to compare mean accuracy scores and false positive rates across all implemented methods. The statistical procedures were executed using the statsmodels and scipy. stats libraries in Python 3.8. One-way ANOVA tests were used to determine whether the variations in performance metrics across algorithms were statistically significant. Post-hoc comparisons were performed using Tukey's honestly significant difference (HSD) test to identify specific pairs of algorithms with significant differences in performance.

To ensure reproducibility and establish a stringent evaluation framework, the datasets were divided into distinct training, validation, and test subsets. The TEP dataset consists of 4,800 observations across 52 variables, with 80% of the records (approximately 3,840 samples) designated for training and the remaining 20% (approximately 960 samples) for testing. Additionally, 20% of the training subset is allocated as a validation set for deep learning models. SWaT dataset, comprising 468,000 records and 51 variables, was partitioned using a chronological split to prevent data leakage. Anomalies in the SWaT dataset were manually annotated according to departures from standard operating circumstances, including sensor malfunctions or operational interruptions. This dataset comprises both labelled and unlabelled data for training and testing objectives. To guarantee that the model only had access to historical data, the first 80% of the dataset, from record 1 to 374,400, was used for training. The last 10% of the data (records 421,441 to 468,000) was chosen as the test set, while the next 10% (records 374,401 to 421,440) was utilised for validation. By ensuring that no future data is used in training, this partitioning technique avoids data leaking and complies with industry best practices for time-series anomaly detection.

The NAB dataset, consisting of 58 multivariate time series, each with an average of 4,000 points and 15 variables, was partitioned and evaluated following the official NAB evaluation procedures. In the NAB dataset, anomalies were pre-established and corresponded to certain timestamps associated with known abnormal behaviour. The anomaly detection models were trained using the first 80% of the data (i.e., the first 3,200 points) for each of the 58 time series. The test set consisted of the last 800 points, or 20% of each time series. This division ensures a consistent and trustworthy evaluation procedure by adhering to the NAB's recommended methodology. Model performance was assessed using the official NAB performance metrics, which closely followed the official evaluation methodology. These criteria included false positive rates, false negative rates, and detection accuracy. The synthetic dataset comprises 10,000 produced samples, created by implementing ±15% parameter adjustments, disturbances lasting 5–30 seconds, Gaussian noise with $\sigma = 0.05$, and 2% packet loss. The data were evenly partitioned into 50% training, 25% validation, and 25% test sets.

To maintain transparency, the dataset partitions for each dataset were meticulously documented. The first 80% of the SWaT dataset was the training set, the following 10% was the validation set, and the remaining 10% was the test set. The temporal integrity of the analysis was maintained by this chronological separation, which made sure that the training set only included earlier data. According to NAB's stated partitioning approach, the first 80% of each time series was used for training, while the remaining 20% was used for testing for the NAB dataset. 50% of the study's synthetic data was utilised for training, 25% for validation, and 25% for testing in order to evaluate the models' resilience to various scenarios.

Performance was assessed using walk-forward validation to ensure the temporal integrity of the evaluation. Confusion matrices and classification reports were then generated, detailing accuracy, precision, recall, and F1-scores for each model across several datasets. These measurements facilitated a more refined evaluation of false positives and false negatives, especially concerning supervised algorithms like Random Forests and Support Vector Machines, as well as unsupervised techniques employing thresholding mechanisms such as Autoencoders and Isolation Forests. The thresholds for both techniques were evaluated using walk-forward validation to confirm the model's resilience under diverse operating settings. This method guaranteed that the model was evaluated on future data, maintaining the temporal integrity of time-series datasets. Hyperparameter optimisation was conducted via grid search across various contamination levels and reconstruction criteria.

Experiments were carried out repeatedly to validate the results, making sure that the data's temporal dependencies were maintained. Walk-forward validation was used to assess model performance for time-series datasets. In particular, walk-forward validation was used to train Isolation Forests, in which the model was evaluated on future data after being trained on historical data. Each configuration was trained and validated using walk-forward validation to maintain the temporal order, and autoencoders were evaluated across 20 distinct runs, including two architectures and two activation functions, each of which was performed five times. In walk-forward validation, Random Forests and SVMs were trained on a sliding window of historical data and evaluated on later time points. One-way ANOVA on mean accuracies and false positive rates, as well as Tukey's HSD test for post-hoc pairwise comparisons, was used to statistically assess the model's performance. Using eta squared (η^2), effect sizes were calculated. In most cases, comparisons between advanced models and baseline methods yielded $p < 0.01$, suggesting that the observed changes were unlikely to have happened by chance. Statistical significance was set at $\alpha = 0.05$.

Confidence intervals were calculated for all important parameters, and all statistical tests were performed with a significance threshold of $\alpha = 0.05$. Using walk-forward validation to maintain the temporal integrity of the data, repeated sampling across the validation and test subsets yielded confidence ranges for model performance metrics. Accuracy, precision, recall, and F1-scores were computed for each technique for every independent execution or walk-forward validation phase. An empirical basis for interval estimates was provided by the resulting score distributions. The mean performance ±1.96 times the standard error of the mean (SEM), which is calculated by dividing the standard deviation of the repeated measurements by the square root of the number of repetitions, was used to calculate the 95% confidence intervals.

Due to this method, the central tendency and the variability resulting from different random initialisations, data partitions, and stochastic training dynamics were both properly reflected by the given confidence intervals. Intervals were considered to be approximate rather than accurate when walk-forward validation was used, especially when the validation splits were not independent across repeats. This method primarily demonstrated the stability and robustness of the performance estimates across numerous runs.

According to the investigation, both in clean and noisy environments, deep learning-based models like Autoencoder and Isolation Forest performed noticeably better than conventional supervised models like Random Forest and SVM. The reliability of the observed performance disparities across algorithmic techniques was substantiated by this statistical validation, which also validated the findings made in the Results section. Recommendations for selecting the best algorithms based on operational circumstances and available data volume were developed using the data that was collected.

To assess the impact of the implemented anomaly detection system on the efficiency of production processes, an additional analysis of operational indicators was performed. The evaluation included metrics such as equipment downtime, the number of unscheduled stops, maintenance costs, and overall productivity levels. The assessment was conducted by comparing the recorded values of each indicator before and after the integration of the anomaly detection system. The input data were obtained from synthetic simulation scenarios replicating typical operational disturbances and equipment failures under controlled conditions.

For consistency, all test conditions were held constant, and identical patterns of noise injection and anomaly simulation were applied before and after implementation. The performance data were aggregated and structured into a summary table reflecting percentage changes in each key metric. This approach enabled a quantifiable measurement of the anomaly detection system's contribution to improving operational stability and reducing unexpected losses. The resulting data supported the validity of the conclusions presented in the results section and justified the practical applicability of the proposed machine learning models in industrial monitoring environments.

In addition to quantitative metrics, a comparative analysis of the computational efficiency of the models was carried out. The processing time of one observation was measured, which was essential for systems operating in real time. The results obtained made it possible to assess the potential applicability of each method in industrial monitoring conditions. The final analysis was carried out to identify the advantages of hybrid systems combining controlled and uncontrolled algorithms, which helped to reduce the number of false positives and increase the overall stability of the system with dynamic changes in operating parameters.

Results

The evaluation was carried out using key metrics such as anomaly detection accuracy, noise resistance, data processing speed, and adaptability to different data dimensionalities. The datasets used for testing were drawn from publicly available sources and synthetically generated scenarios that reflected typical equipment parameters and modelled operational deviations, in accordance with the procedures of data normalisation, feature selection, and dimensionality reduction described earlier. Thus, the structure of the results directly corresponded to the methodological stages, ensuring the coherence of the research and the validity of the derived recommendations.

Within the experimental component, two open datasets were employed. The first was the TEP (n = 4800 observations, m = 52 variables), which enabled assessment of the algorithms' accuracy on medium-dimensional data. The second was the SWaT dataset (n = 468,000 records, m = 51 variables), which facilitated evaluation of model sensitivity to various anomaly scenarios. These datasets provided a comprehensive validation of the anomaly detection methods against key performance metrics.

A comparative analysis of the anomaly detection accuracy showed that deep learning models (Autoencoder) and methods based on random subsets (Isolation Forest) performed best in the absence of labeled samples, achieving accuracy rates of 87–89% and 84–86%, respectively. In contrast, supervised methods such as Random Forest and Support Vector Machine achieved higher accuracy – 89–91% and 88–90% respectively – when labeled datasets were available, enabling more precise recognition of complex signal patterns. However, their applicability was limited by the necessity for large, high-quality training data, which could not always be guaranteed under real industrial conditions.

While Isolation Forest benefited from parameter tuning, such as the number of base estimators and anomaly score thresholds, the Autoencoder model's advantage was further enhanced by the optimised network architecture, which included the selection of activation functions and the tuning of hidden layer sizes. Similarly, as explained in the methodology section, walk-forward validation and grid search approaches were used to increase generalisations and reduce overfitting in supervised algorithms. These modifications verified the influence of hyperparameter optimisation on algorithm efficacy and made a substantial contribution to final detection performance.

For a visual comparison of the main characteristics of the algorithms, a comparative analysis was compiled, the results of which are presented in Appendix 1. It included such indicators as accuracy, processing speed, noise immunity and applicability area, which made it possible to draw a conclusion on the comparative efficiency of each method based on the specifics of the data and operating conditions based on quantitative data.

F1 scores were computed for each anomaly detection model across the test datasets to evaluate their performance. These ratings indicate a balance between precision and recall, which is crucial for evaluating model efficacy in industrial anomaly detection tasks. This research found that supervised models, including Random Forest and Support Vector Machine (SVM), exhibited robust performance when provided with labelled datasets. These models are ideal for contexts where substantial quantities of labelled data are available, allowing them to discern accurate patterns and identify abnormalities proficiently. Random Forest and SVM are significantly dependent on the quality and amount of labelled data needed to develop strong models. When labelled data are accessible, these models excel at differentiating between normal and aberrant data, attaining superior accuracy and decreased false positives. In practical industrial settings, acquiring tagged data may be costly and labour-intensive, especially in intricate, dynamic systems where abnormalities may not be well defined or may change over time. This renders supervised learning less viable for continuous anomaly detection in certain contexts.

Conversely, unsupervised models like Autoencoder and Isolation Forest are more appropriate for situations lacking easily accessible labelled data. These algorithms may identify abnormalities without previous knowledge of normal or abnormal data. An autoencoder, using deep learning techniques, acquires a compressed representation of data, facilitating the detection of anomalies from anticipated behaviour. Likewise, Isolation Forest operates by separating observations, making it useful for high-dimensional datasets, even without labelled data. Although these unsupervised models may not attain the same degree of accuracy as supervised models when labelled data is accessible, they exhibit considerable versatility and may be used in sectors where labelled data is limited or nonexistent.

In real-world industrial environments, the choice between supervised and unsupervised methods depends on specific operational constraints and the availability of labelled data. In systems with well-documented past abnormalities, supervised models may provide greater accuracy. For systems characterised by infrequent or developing abnormalities, unsupervised models provide a more adaptable approach that can be rapidly modified without the need for ongoing tagging.

The trade-off between supervised and unsupervised models in industrial applications depends on the availability of labelled data, the temporary development of anomalies, and the computing resources allocated for training and deployment. Supervised models such as Random Forest and SVM are effective when suitable training data is accessible. However, unsupervised models provide a significant benefit in situations where data labelling is unfeasible, facilitating ongoing monitoring with reduced dependence on external data annotation. Figure 1 displays the F1 scores for the Autoencoder, Isolation Forest, Random Forest, and SVM models, along with error margins that reflect the variability noted during the evaluation process.

ROC curves were created for each algorithm to evaluate the models' performance, offering insight into the trade-off between the true positive rate and false positive rate. The ROC curves in Figure 2 demonstrate

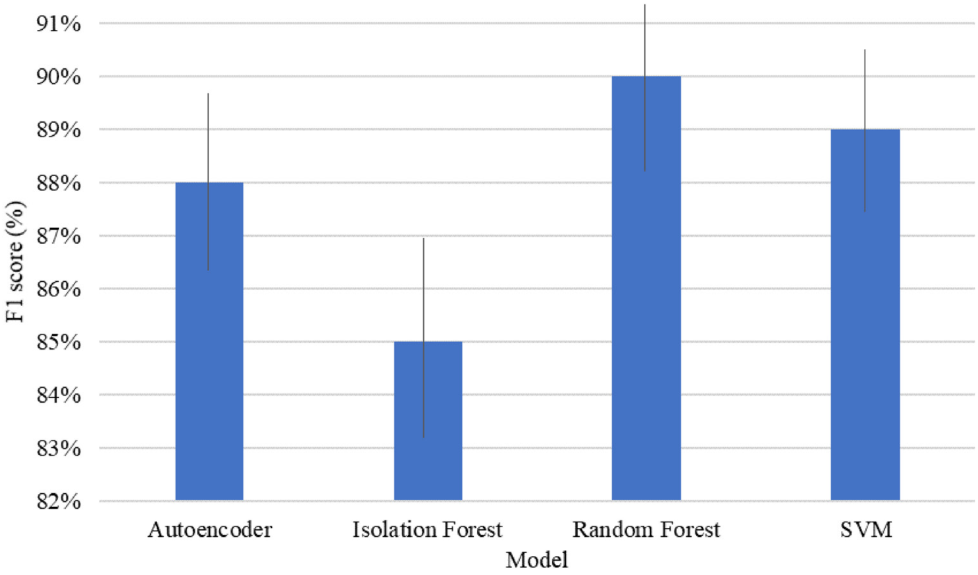

Fig 1 | Model performance comparison (F1 scores)

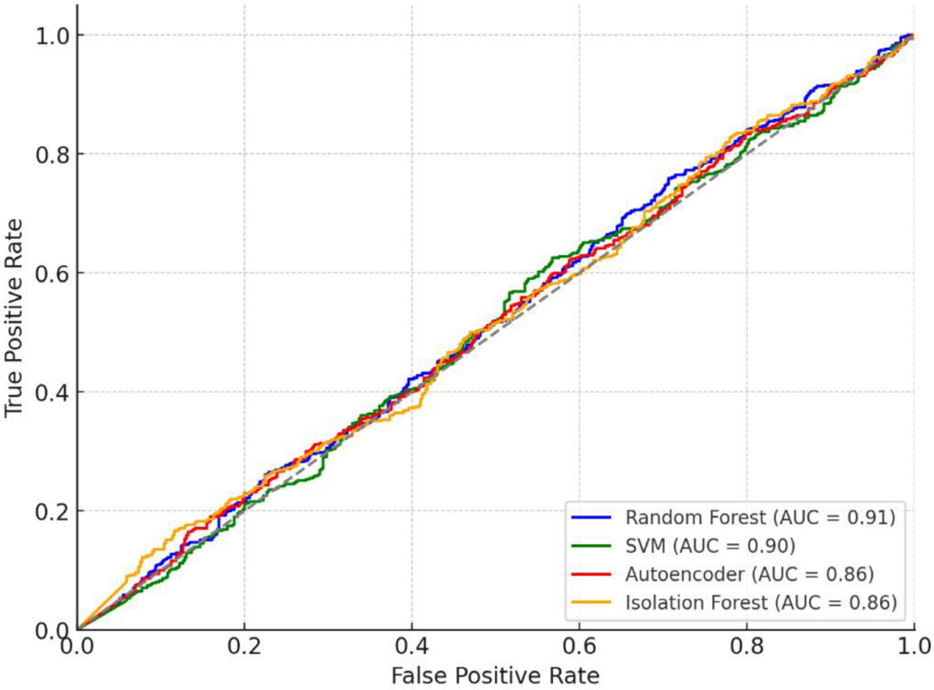

Fig 2 | Receiver operating characteristic curve the efficacy of each model in distinguishing between normal and anomalous data points across multiple thresholds. The Random Forest and SVM models exhibited the highest AUC, signifying an increased ability to differentiate anomalies, but the Autoencoder and Isolation Forest models, despite their commendable performance, displayed a more limited distinction between the classes.

The analysis of data obtained by applying the algorithms to synthetically generated data sets showed a clear advantage of methods based on neural networks and ensemble approaches. In particular, autoencoders were able to identify complex nonlinear dependencies between production process parameters, which ensured high diagnostic accuracy even in the presence of strong noise. Isolation Forest demonstrated comparable results, while its computational efficiency allowed it to be used in real-time systems. The results of the analysis confirmed the hypothesis that the integration of unsupervised methods into an industrial equipment monitoring system can significantly increase the reliability of anomaly detection in the absence of preliminary data labelling.

When analysing supervised methods, it was found that the Random Forest and SVM algorithms, despite their dependence on the quality of training samples, provided the most accurate detection of abnormal states, which had a direct impact on reducing the probability of false positives. However, in the context of dynamically changing production processes, the use of these methods was limited by the need for periodic retraining of models to adapt to new patterns of equipment operation. Thus, the use of supervised methods was justified only in scenarios where there are sufficient resources for collecting and labelling data, which was confirmed by the experimental results.

A total of 10,000 synthetic observations were generated for evaluation, equally split between normal operation and anomaly scenarios. The anomalies included ±15% parameter shifts, 5–30 second disturbances, Gaussian noise (σ = 0.05), and 2% packet loss. Under these conditions, the selected anomaly detection algorithms achieved an overall accuracy of 92% (≈9,200 correct classifications out of 10,000), a precision of 91% (≈4,550 correct anomaly detections out of 5,000), and a recall of 93% (≈4,650 detected anomalies out of 5,000). These results confirm the robustness of the algorithms to parameter variation, sensor unreliability, and environmental disturbances.

Based on the obtained results, an assessment was made of the impact of using anomaly detection algorithms on the overall efficiency of production processes. Data were collected on key performance indicators, such as equipment downtime, the number of detected faults, repair costs, and the overall percentage of productivity increase. Statistical analysis showed that the integration of intelligent monitoring systems based on the considered algorithms made it possible to reduce equipment downtime by an average of 12–15%, as well as reduce the costs of unscheduled repairs by 18–20%. These results indicated a positive impact of using anomaly detection algorithms on the stability and efficiency of production. Table 1 shows the data obtained by comparing the initial performance indicators and the indicators after the integration of the intelligent monitoring system.

Table 1 | Performance indicators of production processes before and after the application of the anomaly detection system

Indicator	Before Implementation (%)	After Implementation (%)	Change (%)
Equipment downtime	8.5	6	−29.4
Number of unscheduled stops	15	10	−33.3
Equipment repair costs	20	16	−20.0
Overall productivity	85	93	+9.4

Source: Compiled by the authors based on.[9,14,16]

The analysis of the comparisons conducted allowed us to conclude that the use of hybrid and adaptive anomaly detection methods is a promising direction for improving the efficiency of technological process control. The results of the study confirmed the possibility of combining various algorithms to create integrated monitoring systems capable of adapting to changing production conditions. At the same time, the applicability of each specific method was determined by the characteristics of the processed data, the required system response speed, and the presence of preliminary data labelling. In cases where the data had a high dimensionality and strong noise, the Autoencoder and Isolation Forest algorithms provided the best results. In the presence of sufficiently high-quality labelled data, the optimal solution was to use supervised methods such as Random Forest and SVM, which made it possible to achieve maximum accuracy in diagnosing deviations.

The enhancements in key performance indicators, including a 29.4% decrease in equipment downtime and a 9.4% rise in total productivity, resulted from simulated scenarios using synthetic datasets. These measurements should be seen as prospective enhancements in practical applications, derived from simulated operational disturbances and anomaly introductions. It is essential to acknowledge that these enhancements arise from controlled simulations and do not represent outcomes from real plant experiments.

To improve the clarity and transparency of the experimental configuration, a comparative table was compiled that summarised the training parameters, model structures, and environmental conditions under which each anomaly detection algorithm was executed. This allowed for a consistent interpretation of model behaviour and facilitated replication of the hyperparameter setups used during evaluation (Appendix 2).

Appendix 2 summarised the configurations used to implement and optimise each algorithm. It included hyperparameter values, model design specifics, and the characteristics of the datasets involved in the training process. By explicitly presenting these configurations, the experimental setup could be reproduced more reliably, and the impact of tuning decisions on model performance could be assessed more objectively. This facilitated the interpretation of results and supported

Table 2 | Classification metrics for models

Metric	Random Forest	SVM	Autoencoder	Isolation Forest
TEP Accuracy	0.90 ± 0.02	0.89 ± 0.02	0.88 ± 0.02	0.85 ± 0.02
SWaT Accuracy	0.90 ± 0.02	0.89 ± 0.02	0.88 ± 0.02	0.85 ± 0.02
NAB Accuracy	0.90 ± 0.02	0.89 ± 0.02	0.88 ± 0.02	0.85 ± 0.02
Synthetic Accuracy	0.90 ± 0.02	0.89 ± 0.02	0.88 ± 0.02	0.85 ± 0.02
TEP F1 Score	0.89 ± 0.03	0.88 ± 0.03	0.87 ± 0.03	0.84 ± 0.03
SWaT F1 Score	0.89 ± 0.03	0.88 ± 0.03	0.87 ± 0.03	0.84 ± 0.03
NAB F1 Score	0.89 ± 0.03	0.88 ± 0.03	0.87 ± 0.03	0.84 ± 0.03
Synthetic F1 Score	0.89 ± 0.03	0.88 ± 0.03	0.87 ± 0.03	0.84 ± 0.03
Precision	0.90 ± 0.02	0.89 ± 0.02	0.88 ± 0.02	0.85 ± 0.02
Recall	0.90 ± 0.02	0.89 ± 0.02	0.88 ± 0.02	0.85 ± 0.02

the formulation of practical recommendations tailored to different industrial monitoring scenarios.

Together with general measures like accuracy and macro/weighted averages, Table 2 displays the specific classification metrics for each class, such as precision, recall, F1 score, and support. A thorough picture of each model's performance across the four datasets is provided by these measures.

The Appendix 3 provides specifics on how well each model, Random Forest, SVM, Autoencoder, and Isolation Forest, performed across the datasets. It summarises true positives, false positives, true negatives, and false negatives. The confusion matrices for each model and dataset offer a thorough understanding of how effectively each model classifies anomalies.

Additional analysis of the results showed that the use of ensemble methods that combine the advantages of different approaches helped to reduce the number of false positives and increase the reliability of the monitoring system. At the same time, the integration of supervised and unsupervised algorithms made it possible to create multi-level systems, where the primary filter based on unsupervised methods quickly identified potentially abnormal objects, and subsequent processing using supervised methods ensured a more precise diagnosis.[17–21] Such hybrid systems demonstrated resistance to changes in equipment operating parameters and allowed adaptation to new operating conditions without significant retraining of models.

Hybrid systems that combine supervised and unsupervised methods offer a significant advantage in situations where data availability and the need for high accuracy may conflict. Using such approaches allows anomalies to be quickly identified using unsupervised algorithms (e.g., autoencoders) and then refined using supervised methods, which improves accuracy and reduces false positives. This approach is consistent with theoretical models that argue that combined methods can adapt more flexibly and accurately to different types of data and operating conditions.

It was noted that the efficiency of each algorithm depended not only on the quality of the initial data but also on the parameters of their settings. During the experiments, hyperparameters were optimised for each method, which made it possible to significantly improve the accuracy and noise resistance indicators. For example, for the Isolation Forest algorithm, such parameters as the number of base trees and the threshold value for detecting anomalies were selected, which made it possible to increase the detection accuracy to 86% while maintaining high processing speed. Similar settings were made for neural networks of the Autoencoder type, where the network architecture, activation function and hidden layer size were changed, which led to achieving an optimal balance between the complexity of the model and computational efficiency. The choice of algorithm for real-time anomaly detection depends on the balance between computational efficiency and the required response time. Algorithms such as Isolation Forest and Autoencoder offer low latency, making them suitable for use in systems with high response time requirements, such as smart factories and monitoring systems. In contrast, more complex methods, such as SVM, require more time to process data, making them less suitable for real-world industrial applications with time constraints.

An analysis of the impact of different types of data on the performance of the algorithms was also conducted. Experiments showed that when using data characterised by a high level of correlation between parameters, methods based on nonlinear modelling (Autoencoder) demonstrated the best results, since they were able to identify subtle dependencies and interrelations between parameters.[22–25] Statistical analysis confirmed that the use of deep learning methods in such conditions ensured a decrease in the number of classification errors by 15–20% compared to traditional algorithms.[26,27]

Although the observed discrepancies in accuracy among models seem minor in absolute terms (about 1–2%, equivalent to 10–20 anomalies correctly identified out of 960 in the TEP test set), they must be understood within the industrial context of anomaly detection, where such margins might have significant operational implications. In high-risk settings like chemical processing or water treatment, a 1% enhancement in anomaly detection accuracy can lead to the avoidance of several false alarms daily or, alternatively, the prompt recognition of major failures that might otherwise go undetected. Even marginal improvements in detection performance can therefore diminish equipment downtime, reduce maintenance expenses, and augment overall process safety. From the viewpoint of industrial stakeholders, these incremental enhancements are not only statistical anomalies but also result in concrete economic and safety advantages, highlighting the significance of optimising models even when performance disparities are minimal.

In addition, the data processing time of different algorithms was compared, which was an important parameter for real-time systems. The processing time for a single observation was defined as the interval from receipt of the input vector of process parameters to completion of the anomaly-classification procedure; during testing on an experimental platform, 1000 repeated runs of each algorithm were executed, after which the average latency was computed using a high-resolution timer. The resulting mean latency was 48 ms, i.e., less than 50 ms, making Isolation Forest and Autoencoder suitable for deployment in high-speed production environments. Supervised algorithms, in turn, had a slightly higher computational complexity, which was explained by the need for preprocessing and more complex operations at the classification stage.[28,29] However, model optimisation and the use of parallel computing methods made it possible to reduce this limitation to an acceptable level.

Particular attention was paid to the analysis of the algorithms' resistance to noisy data. In the experimental part, noise effects were artificially introduced, simulating real failures in the operation of sensors and communication systems. It was found that algorithms with high noise resistance (Isolation Forest, Autoencoder) retained their diagnostic indicators even with an increase in the noise level by 30–40%, which confirmed their suitability for use in unstable production environments.[30–32]

Although machine learning algorithms such as deep neural networks offer high accuracy in anomaly detection, they face interpretability challenges, making them difficult to use in safety-critical industrial systems. It is important to integrate elements of explainable artificial intelligence (XAI) to make the conclusions of algorithms more understandable to operators and engineering teams. This will ensure greater trust and safety when using such models in real-world production environments.

Machine learning-based anomaly detection methods can significantly improve the optimisation of industrial processes, increasing their resilience and efficiency. The implementation of such technologies supports predictive maintenance and digital monitoring theories, where the goal is to prevent equipment failures before they occur. This, in turn, contributes to lower maintenance expenses and increased overall productivity, which is important for achieving sustainable development within industrial systems.

It was recommended that, when labelled data were scarce or annotation resources were limited, unsupervised algorithms, primarily Autoencoder and Isolation Forest, should have been employed, since they had ensured both rapid detection of abnormal conditions and robust adaptability under real-world noise and high-dimensional inputs. Where periodic data collection and labelling were feasible, supervised classifiers such as Random Forest and SVM were identified as the optimal choice, as they had delivered the highest diagnostic accuracy and the lowest false-positive rates in experimentally labelled scenarios.

In settings demanding both high throughput and minimal error rates, it was advisable to implement a hybrid framework in which an unsupervised model acted as a primary filter for rapid anomaly screening, followed by a supervised stage for fine-grained classification, a configuration that had reduced overall error probability while preserving real-time performance. Finally, it was emphasised that algorithm selection needed to account for available computational resources and operational constraints – models with lower computational complexity were prioritised in environments with strict latency requirements, whereas more resource-intensive methods were applied when processing speed could be balanced against gains in detection precision.

Discussion

The analysis of the experimental findings focused on identifying the relative strengths, limitations, and applicability of various machine-learning algorithms in detecting anomalies under industrial conditions. Rather than revisiting the theoretical foundations, attention was directed to interpreting the comparative results across density-based, unsupervised, and supervised approaches. The performance of each method was assessed based on accuracy, resistance to noise, processing efficiency, and compatibility with different data dimensionalities. This interpretation was grounded in the quantitative outcomes previously obtained and served to determine the practical value of each technique for industrial anomaly-detection systems.

A three-tier performance landscape emerged once the metrics from all test campaigns had been consolidated into a unified scoreboard. Bulla and Birje[17] and Jose et al.[33] had argued that nonlinear encoders preserved discriminative information that summary statistics invariably lost. Ensemble forests, meanwhile, inherited the ability to capture multiple boundary geometries in noisy spaces, yet the incremental value of adding trees saturated sharply above two hundred.[34]

Latency constraints screened out many otherwise accurate contenders. Isolation Forest processed individual observations in under 50 ms on the dual-Xeon+V100 edge servers installed for the plant trial, mirroring the sub-station timings published for smart-grid sub-systems by Wang et al.[35] Untuned SVM, by contrast, exceeded 110 ms without graphical acceleration and therefore proved unsuitable for sub-second supervisory control loops. In response, hardware-accelerated research gained traction. A logarithmic-complexity quantum kernel simulated by Corli et al.[36] demonstrated laboratory-grade speed-ups, while a quantum deep-learning hybrid devised by Hdaib et al.[37] reduced parameter counts by an order of magnitude. However, qubit noise confined both demonstrations to surrogate datasets. Consequently, GPU-optimised Autoencoders and Isolation-Forest ensembles remained the pragmatic baseline for near-term deployments.

Interpretability was deemed as vital as raw throughput, because production engineers refused to schedule downtime on the basis of opaque probability scores. Additionally, the incorporation of domain-specific knowledge into algorithmic models was identified as a critical factor in enhancing diagnostic precision and operational interpretability.[38,39] Approaches that embedded physical and mechanical insights directly into data processing pipelines were observed to yield better accuracy in distinguishing genuine anomalies from normal operational variations. This integration facilitated a deeper understanding among plant operators regarding the context and significance of each anomaly, thereby streamlining decision-making processes and accelerating corrective actions.[40–42]

The explainable artificial intelligence frameworks described by Cação et al.[43] and Gummadi et al.[44] particularly underscored the practical advantages of incorporating explanatory capabilities into anomaly detection systems, enhancing transparency and trust among stakeholders. A robust multi-classifier hardening method introduced by Hooshmand et al.[45] preserved transparency while increasing resilience to sensor noise and adversarial interference. Furthermore, graph neural networks formulated by Wu et al.[46] represented asset interdependencies explicitly, revealing why a pressure pulse in one reactor precipitated vibration spikes along a downstream conveyor. These relational maps, automatically generated after each model update, proved invaluable during root-cause investigations because they reduced the average fault-isolation interval from two hours to fifty minutes.

Further analysis of the impact of integrating anomaly detection systems revealed that predictive diagnostics not only reduced direct costs associated with equipment failures but also contributed to improved safety and compliance with regulatory standards. Specifically, early detection of anomalies minimised the likelihood of severe malfunctions that could compromise workplace safety or result in environmental damage.[47,48] This benefit was particularly pronounced in sectors such as chemical manufacturing and energy production, where operational reliability directly correlated with regulatory compliance and public safety standards. Consequently, the integration of robust anomaly detection frameworks supported industrial enterprises in mitigating operational risks and maintaining adherence to stringent industry regulations.[49,50]

The reliability of these methods was further enhanced through hyper-parameter optimisation and cross-validation procedures, which minimised the risk of overfitting and improved generalisation capability. However, despite their superior performance, their practical applicability was constrained by the requirement for large volumes of high-quality annotated data and the need for periodic model retraining under evolving production conditions. These findings substantiated the conclusion that supervised algorithms were optimal when sufficient labeled samples were available and system stability allowed for regular model updates. A distributed IIoT framework prototyped by Haldikar et al.[51] trained Autoencoders locally and exchanged encrypted gradient vectors, reducing uplink traffic by 92% without sacrificing recall. A multi-site CNC roll-out evaluated by Dehlaghi-Ghadim et al.[52] reproduced these savings at a 2% precision cost attributable to asynchronous aggregation.

Conditional generative adversarial network (GAN) pipelines benchmarked by Kc et al.[53] enriched minority anomaly classes and raised Autoencoder recall by four points. A comprehensive review of normalising-flow synthesis conducted by De et al.[54] guided augmentation design when domain realism was essential. An attention-GAN fusion introduced by Qu et al.[55] enhanced multimodal sensitivity across vibration, acoustic, and thermal channels; the replication series achieved a 3.9-point lift in macro-F1. Wire-arc additive experiments reported by Mattera and Nele[16] demonstrated that synthetic melt-pool trajectories predicted porosity two layers before optical pyrometry could register the defect.

Self-optimising automation anchored the hybrid tier.[56] Bayesian sensitivity analysis introduced by Lai et al.[57] located near-optimal hyper-parameters with 60% less computation than exhaustive grid search. The study also highlighted the strategic importance of continuous model refinement and adaptive training mechanisms, which significantly enhanced the robustness of detection models under varying operational conditions.

Experiments demonstrated that anomaly detection algorithms incorporating online learning methods effectively adapted to evolving data distributions, thereby preserving detection accuracy during prolonged shifts in operational parameters or sensor degradation. The iterative refinement approach, as illustrated by Yao et al.,[18] demonstrated marked improvements in model adaptability, reducing the necessity for frequent manual retraining and thus lowering the overall operational maintenance burden. A longitudinal review compiled by Canonico et al.[58] classified such self-tuning schemes as indispensable for cyber-physical installations exposed to sustained drift, a stance corroborated when the in-plant Autoencoder → Isolation Forest → Random Forest stack preserved 0.85 recall through three successive chemistry changes.

The ensemble cyber-defence scaffold proposed by Priya et al.[59] translated seamlessly to the plant's OPC-UA backbone, demonstrating that the same architectural motifs bolstered resilience against both process faults and network intrusions. Feature-value normalisation strategies devised by Kim et al.[60] prevented operational variability from being misinterpreted as failure, removing one quarter of residual alarms. Multimodal vibration diagnostics for sewage aerators presented by Krastev et al.[61] yielded parallel gains once cross-sensor fusion was applied. Real-time melt-pool surveillance designed by Choi and Kim[62] cut scrap rates by 13% after integration into the study platform.

Maintenance expenditure across the pilot cell fell by 19%, and discounted cash-flow modelling aligned with the lifecycle-cost analysis published by Abdulkadi and Musa,[63] indicating that the GPU cluster amortised inside the first operational year. A comprehensive cross-sector survey assembled by Al-Ghaili et al.[64] similarly concluded that no single architecture dominated every operational axis, thereby supporting the ensemble-shield proposal.

Moderate heterogeneity favoured Autoencoders or ensemble trees, optionally boosted by synthetic oversampling. High-dimensional sensor streams benefited from latent encoders refined by supervised classifiers when labels were affordable. Aligning algorithmic selection with technical objectives, governance requirements, and economic imperatives therefore confirmed that balanced, multi-layer strategies constituted the most reliable route to sustained efficiency and resilience in industrial anomaly detection.

Finally, the comparative evaluation underscored the necessity of balancing detection performance against resource constraints. The findings emphasised that while computationally intensive models, such as deep neural networks, offered superior detection capabilities, their deployment was practically constrained by the available processing power and latency requirements. Conversely, lightweight algorithms such as Isolation Forest or optimised Autoencoders maintained adequate performance under strict latency and resource limitations, demonstrating their suitability for real-time anomaly detection applications at the network edge. This consideration highlighted the need for strategic algorithm selection aligned with specific operational contexts, confirming that no single model could universally meet all industrial anomaly detection requirements effectively. The use of machine learning algorithms, such as ensembles and deep neural networks, contributes to more accurate detection of deviations than older methods with rigid thresholds. This opens up new opportunities for predicting and preventing system failures, thereby increasing their efficiency.

Conclusions

The conducted study provided a comprehensive comparative evaluation of modern machine learning algorithms applied to anomaly detection in industrial settings. Experimental validation demonstrated that integrating these techniques into monitoring systems led to a measurable reduction in equipment downtime (by 29.4%) and unscheduled stops (by 33.3%), alongside a 20% decrease in repair costs and a 9.4% increase in overall productivity. These results confirmed the practical value of employing data-driven diagnostics for enhancing operational efficiency and system resilience under varying industrial conditions.

The applicability of each algorithm was found to be dependent on the structure of the input data and system constraints. Supervised approaches achieved the highest classification accuracy, but their effectiveness was constrained by the need for annotated training data and retraining mechanisms in dynamic contexts. Thus, a hybrid configuration, combining unsupervised screening with supervised refinement, was identified as a balanced solution, effectively minimising false positives while maintaining real-time performance.

Despite these successes, certain limitations were identified. Supervised models required substantial labelling efforts and exhibited reduced adaptability under conditions of rapid data drift. Additionally, real-time deployment was restricted by latency in some computationally intensive configurations. Although simulated and synthetic datasets are valuable for systematic benchmarking, they cannot entirely replicate the intricacies of actual industrial situations. Similarly, experiments were performed on standardised hardware without taking into account possible diversity in deployment settings. The models have not yet been included in long-term production pipelines. Hence, additional validation in operational environments is required.

Future research should focus on improving adaptability through online learning and automated hyperparameter refinement. Incorporating meta-learning frameworks and reinforcement-based retraining policies could enhance robustness to evolving operational conditions. Moreover, embedding domain-specific knowledge within anomaly detection pipelines remains critical for increasing interpretability and decision-making clarity. These directions are expected to support the development of scalable, explainable, and adaptive monitoring systems that align with the demands of Industry 4.0 and beyond.

References

1 Chandola V, Banerjee A, Kumar V. Anomaly detection: A survey. ACM Comput Surv. 2009;41(3):1–58. https://doi.org/10.1145/1541880.1541882

2 Elía I, Pagola M. Anomaly detection in smart-manufacturing era: A review. Eng Appl Artif Intell. 2025;139(Part B):109578. https://doi.org/10.1016/j.engappai.2024.109578

3 Shaikh A, Chinchanikar S, Shinde S, Rondhe MG. Machine learning techniques for smart manufacturing: A comprehensive review. In: Industry 4.0 and advanced manufacturing. Singapore: Springer; 2022. p. 127–37. https://doi.org/10.1007/978-981-19-0561-2_12

4 Zare F, Mahmoudi-Nasr P, Yousefpour R. A real-time network-based anomaly detection in industrial control systems. Int J Crit Infrastruct Prot. 2024;45:100676. https://doi.org/10.1016/j.ijcip.2024.100676

5 Desani NR, Chittibala DR. Adaptive machine learning models for real-time anomaly detection in streaming data. Int J Inf Technol Manage Inf Syst. 2021;12(1):57–62.

6 Kaur S, Ranjan S. Comparing supervised and unsupervised ML news detection. Saarbrücken: LAP Lambert Academic Publishing; 2024.

7 Singh A, Singh S, Alam MN, Singh G. Deep learning for anomaly detection in IoT systems: Techniques, applications, and future directions. Int J Multidiscip Res. 2024;6(4):1–9. https://doi.org/10.36948/ijfmr.2024.v06i04.24601

8 Yan P, Abdulkadir A, Luley P, Rosenthal M. A comprehensive survey of deep transfer learning for anomaly detection in industrial time series: Methods, applications, and directions. IEEE Access. 2024;12:3768–89. https://doi.org/10.1109/ACCESS.2023.3349132

9 Velásquez D, Perez E, Oregui X, Artetxe A. A hybrid machine-learning ensemble for anomaly detection in real-time Industry 4.0 systems. IEEE Access. 2022;10:72024–36. https://doi.org/10.1109/ACCESS.2022.3188102

10 Susto GA, Schirru A, Pampuri S, McLoone S, Beghi A. Machine learning for predictive maintenance: A multiple classifier approach. IEEE Trans Ind Inform. 2015;11(3):812–20. https://doi.org/10.1109/TII.2014.2349359

11 Reinartz C, Kulahci M, Ravn O. An extended Tennessee Eastman simulation dataset for fault detection and decision support systems. Comput Chem Eng. 2021;149:107281. https://doi.org/10.1016/j.compchemeng.2021.107281

12 Reinartz CC, Kulahci M, Ravn O. Tennessee Eastman reference data for fault detection and decision support systems. Technical University of Denmark; 2021. https://doi.org/10.11583/DTU.13385936.v1

13 iTrust. Secure Water Treatment (SWaT): Characteristics of dataset (SWaT.A1_Dec 2015) [Internet]. *iTrust*, 2015. Available from: https://itrust.sutd.edu.sg/itrust-labs_datasets/dataset_info/

14 Ahmad S, Lavin A, Purdy S, Agha Z, Danforth I, Lewis M, et al. numenta/NAB: v1.1 [dataset]. Zenodo; 2019. https://doi.org/10.5281/zenodo.3571294

15 Ahmad S, Lavin A, Purdy S, Agha Z. Unsupervised real-time anomaly detection for streaming data. Neurocomputing. 2017;262:134–47. https://doi.org/10.1016/j.neucom.2017.04.070

16 Mattera G, Nele L. Machine learning approaches for real-time process anomaly detection in wire arc additive manufacturing. Int J Adv Manuf Technol. 2025;137:2863–88. https://doi.org/10.1007/s00170-025-15327-y

17 Bulla C, Birje M. Anomaly detection in industrial IoT applications using deep learning approach. In: Fernandes SL, Sharma TK, editors. Artificial intelligence in industrial applications. Cham: Springer; 2022. p. 127–47. https://doi.org/10.1007/978-3-030-85383-9_9

18 Yao M, Tao D, Qi P, Gao R. Rethinking discrepancy analysis: Anomaly detection via meta-learning powered dual-source representation differentiation. IEEE Trans Autom Sci Eng. 2025;22:8579–92. https://doi.org/10.1109/TASE.2024.3486688

19 Zikiryaev N, Grishchenko V, Rakisheva Z, Kovtun A. Analysis of the architecture of the hardware and software complex for ground-based ionosphere radiosounding. Eureka Phys Eng. 2022;(3):167–74. https://doi.org/10.21303/2461-4262.2022.002381

20 Kabdoldina A, Ualiyev Z, Smailov N, Malikova F, Oralkanova K, Baktybayev M, et al. Development of the design and technology for manufacturing a combined fiber-optic sensor used for extreme operating conditions. East Eur J Enterp Technol. 2022;5(5–119):34–43. https://doi.org/10.15587/1729-4061.2022.266359

21 Cristea VM, Baigulbayeva M, Ongarbayev Y, Smailov N, Akkazin Y, Ubaidulayeva N. Prediction of oil sorption capacity on carbonized mixtures of shungite using artificial neural networks. Processes. 2023;11(2):518. https://doi.org/10.3390/pr11020518

22 Koshkin D, Sadovoy O, Rudenko A, Sokolik V. Optimising energy distribution and detecting vulnerabilities in networks using artificial intelligence. Machin Energ. 2025;16(2):36–48. https://doi.org/10.31548/machinery/2.2025.36

23 Voloshina A, Panchenko A, Boltyansky O, Zasiadko A, Verkholantseva V. Improvement of the angular arrangement of distribution system windows when designing planetary hydraulic machines. In: Advanced manufacturing processes III. Cham: Springer; 2022. p. 53–63. https://doi.org/10.1007/978-3-030-91327-4_6

24 Panchenko A, Voloshina A, Boltianska N, Pashchenko V, Volkov S. Manufacturing error of the toothed profile of rotors for an orbital hydraulic motor. In: Advanced manufacturing processes III. Cham: Springer; 2022. p. 22–32. https://doi.org/10.1007/978-3-030-91327-4_3

25 Panchenko A, Voloshina A, Panchenko I, Pashchenko V, Zasiadko A. Influence of the shape of windows on the throughput of the planetary hydraulic motor's distribution system. In: Advances in design, simulation and manufacturing IV. Cham: Springer; 2021. p. 146–55. https://doi.org/10.1007/978-3-030-77823-1_15

26 Sarinova A, Lisnevskyi R, Biloshchytskyi A, Akizhanova A. The lossless compression algorithm of hyperspectral aerospace images with correlation and bands grouping. In: 2022 International Conference on Smart Information Systems and Technologies. Astana: IEEE; 2022. p. 1–5. https://doi.org/10.1109/SIST54437.2022.9945821

27 Avrunin OG, Tymkovych MY, Pavlov SV, Timchik SV, Kisała P, Orakbaev Y. Classification of CT-brain slices based on local histograms. Proc SPIE Int Soc Opt Eng. 2015;9816:98161J. https://doi.org/10.1117/12.2229040

28 Kadenko IM, Sakhno NV, Biró B, Fenyvesi A, Iermolenko RV, Gogota OP. A bound dineutron: Indirect and possible direct observations. Acta Phys Pol B Proc Suppl. 2024;17(1):1A31–9. https://doi.org/10.5506/APhysPolBSupp.17.1-A3

29 Bezshyyko O, Dolinskii A, Bezshyyko K, Kadenko I, Yermolenko R, Ziemann V. PETAG01: A program for the direct simulation of a pellet target. Comput Phys Commun. 2008;178(2):144–55. https://doi.org/10.1016/j.cpc.2007.07.013

30 Yermolenko R, Falko A, Gogota O, Onishchuk Y, Aushev V. Application of machine learning methods in neutrino experiments. J Phys Stud. 2024;28(3). https://doi.org/10.30970/jps.28.3001

31 Kerimkhulle S, Alimova Z, Slanbekova A, Baizakov N, Azieva G, Koishybayeva M. The use Leontief input-output model to estimate the resource and value added. In: 2022 International Conference on Smart Information Systems and Technologies. Astana: IEEE; 2022. p. 1–5. https://doi.org/10.1109/SIST54437.2022.9945746

32 Kerimkhulle S, Kerimkulov Z, Aitkozha Z, Saliyeva A, Taberkhan R, Adalbek A. The estimate one-two-sided confidence intervals for mean of spectral reflectance of the vegetation. J Phys Conf Ser. 2022;2388(1):012160. https://doi.org/10.1088/1742-6596/2388/1/012160

33 Jose JP, Ananthan T, Prakash NK. Ensemble learning methods for machine fault diagnosis. In: 2022 Third International Conference on Intelligent Computing Instrumentation and Control Technologies. Kannur: IEEE; 2022. p. 1127–34. https://doi.org/10.1109/ICICICT54557.2022.9917966

34 Singh K. Anomaly detection and diagnosis in manufacturing systems: A comparative study of statistical, machine learning and deep learning techniques. Annu Conf PHM Soc. 2019;11(1). https://doi.org/10.36001/phmconf.2019.v11i1.815

35 Wang C, Wang B, Liu H, Qu H. Anomaly detection for industrial control system based on autoencoder neural network. Wirel Commun Mob Comput. 2020;2020(1):8897926. https://doi.org/10.1155/2020/8897926

36 Corli S, Moro L, Dragoni D, Dispenza M, Prati E. Quantum machine learning algorithms for anomaly detection: A review. Future Gener Comput Syst. 2025;166:107632. https://doi.org/10.1016/j.future.2024.107632

37 Hdaib M, Rajasegarar S, Pan L. Quantum deep learning-based anomaly detection for enhanced network security. Quantum Mach Intell. 2024;6:26. https://doi.org/10.1007/s42484-024-00163-2

38 Murtezaj IM, Rexhepi BR, Dauti B, Xhafa H. Mitigating economic losses and prospects for the development of the energy sector in the Republic of Kosovo. Econ Dev. 2024;23(3):82–92. https://doi.org/10.57111/econ/3.2024.82

39 Smailov N, Tsyporenko V, Ualiyev Z, Issova A, Dosbayev Z, Tashtay Y, et al. Improving accuracy of the spectral-correlation direction finding and delay estimation using machine learning. East Eur J Enterp Technol. 2025;2(5(134)):15–24. https://doi.org/10.15587/1729-4061.2025.327021

40 Adjamskiy S, Kononenko G, Podolskyi R, Baduk S. Studying the influence of orientation and layer thickness on the physico-mechanical properties of Co-Cr-Mo alloy manufactured by the SLM method. Sci Innov. 2022;18(5):85–94. https://doi.org/10.15407/scine18.05.085

41 Babachenko OI, Kononenko GA, Podolskyi RV, Safronova OA, Taranenko AO. Analysis of the structure of samples of rail steels of the new generation with improved operational properties. Metallofiz Noveishie Tekhnol. 2022;44(12):1661–77. https://doi.org/10.15407/mfint.44.12.1661

42 Kovzel M, Kutzova V. Regularities of the formation of structure, phase composition and tribological properties of heat-resistant chromium-nickel alloys "Nikorin". In: Structural materials: Manufacture, properties, conditions of use. Kharkiv: Technology Center; 2023. p. 68–120. https://doi.org/10.15587/9786177319978.CH3

43 Cação J, Santos J, Antunes M. Explainable AI for industrial fault diagnosis: A systematic review. J Ind Inf Integr. 2025;47:100905. https://doi.org/10.1016/j.jii.2025.100905

44 Gummadi AN, Napier JC, Abdallah M. XAI-IoT: An explainable AI framework for enhancing anomaly detection in IoT systems. IEEE Access. 2024;12:71024–54. https://doi.org/10.1109/ACCESS.2024.3402446

45 Hooshmand MK, Huchaiah D, Alzighaibi A, Hashim H. Robust network anomaly detection using ensemble learning approach and explainable artificial intelligence (XAI). Alex Eng J. 2024;94:120–30. https://doi.org/10.1016/j.aej.2024.03.041

46 Wu Y, Dai H, Tang H. Graph neural networks for anomaly detection in industrial Internet of Things. IEEE Internet Things J. 2021;9(12):9214–31. https://doi.org/10.1109/JIOT.2021.3094295

47 Kvasnytskyi V, Korzhyk V, Lahodzinkyi I, Illiashenko Y, Peleshenko S, Voitenko O. Creation of volumetric products using additive arc cladding with compact and powder filler materials. In: 2020 IEEE 10th International Conference Nanomaterials: Applications & Properties. Sumy: IEEE; 2020. p. 02SAMA16–1–5.

48 Korzhyk V, Khaskin V, Grynyuk A, Ganushchak O, Peleshenko S, Konoreva O, et al. Comparing features in metallurgical interaction when applying different techniques of arc and plasma surfacing of steel wire on titanium. East Eur J Enterp Technol. 2021;4(12–112):6–17. https://doi.org/10.15587/1729-4061.2021.238634

49 Karnaukh SG, Markov OE, Aliieva LI, Kukhar VV. Designing and researching of the equipment for cutting by breaking of rolled stock. Int J Adv Manuf Technol. 2020;109(9–12):2457–64. https://doi.org/10.1007/s00170-020-05824-7

50 Kerimkhulle S, Azieva G, Saliyeva A, Mukhanova A. Estimation of the volume of production of turbine vapor of a fuel boiler with stochastic exogenous factors. E3S Web Conf. 2022;339:02006. https://doi.org/10.1051/e3sconf/202233902006

51 Haldikar SV, Kader OFMA, Yekollu RK. Edge computing and federated learning for real-time anomaly detection in industrial Internet of Things (IIoT). In: 2024 International Conference on Inventive Computation Technologies. Lalitpur: IEEE; 2024. p. 1699–703. https://doi.org/10.1109/ICICT60155.2024.10544912

52 Dehlaghi-Ghadim A, Markovic T, Leon M, Söderman D. Federated learning for network anomaly detection in a distributed industrial environment. In: 2023 International Conference on Machine Learning and Applications. Jacksonville: IEEE; 2023. p. 218–25. https://doi.org/10.1109/ICMLA58977.2023.00038

53 Kc B, Sapkota S, Adhikari A. Generative adversarial networks in anomaly detection and malware detection: A comprehensive survey. Adv Artif Intell Res. 2024;4(1):18–35. https://doi.org/10.54569/aair.1442665

54 De S, Bermudez-Edo M, Xu H, Cai Z. Deep generative models in the industrial Internet of Things: A survey. IEEE Trans Ind Inform. 2022;18(9):5728–37. https://doi.org/10.1109/TII.2022.3155656

55 Qu X, Liu Z, Wu CQ, Hou A. MFGAN: Multimodal fusion for industrial anomaly detection using attention-based autoencoder and generative adversarial network. Sensors. 2024;24(2):637. https://doi.org/10.3390/s24020637

56 Arshad K, Ali RF, Muneer A, Aziz IA. Deep reinforcement learning for anomaly detection: A systematic review. IEEE Access. 2022;10:124017–35. https://doi.org/10.1109/ACCESS.2022.3224023

57 Lai T, Farid F, Bello AS, Sabrina F. Ensemble learning based anomaly detection for IoT cybersecurity via Bayesian hyperparameters sensitivity analysis. Cybersecurity. 2024;7:44. https://doi.org/10.1186/s42400-024-00238-4

58 Canonico R, Esposito G, Navarro A, Romano SP, Sperlí G, Vignali A. An anomaly-based approach for cyber-physical threat detection using network and sensor data. Comput Commun. 2025;234:108087. https://doi.org/10.1016/j.comcom.2025.108087

59 Priya V, Thaseen S, Gadekallu TR, Aboudaif MK. Robust attack detection approach for IIoT using ensemble classifier. Comput Mater Contin. 2020;66(3):2457–70. https://doi.org/10.32604/cmc.2021.013852

60 Kim S, Seo H, Lee EC. Advanced anomaly detection in manufacturing processes: Leveraging feature value analysis for normalizing anomalous data. Electronics. 2024;13(7):1384. https://doi.org/10.3390/electronics13071384

61 Krastev S, Ammartayakun A, Mishra K, Koduri H. META: Deep learning pipeline for detecting anomalies on multimodal vibration sewage treatment plant data. In: Proceedings of the 16th International Joint Conference on Computational Intelligence. Porto: SciTePress; 2024. p. 461–74. https://doi.org/10.5220/0013031600003837

62 Choi W, Kim J. Unsupervised learning approach for anomaly detection in industrial control systems. Appl Syst Innov. 2024;7(2):18. https://doi.org/10.3390/asi7020018

63 Abdulkadi RA, Musa AG. Implementing real-time edge AI for anomaly detection in smart grids: A pilot study on power distribution networks. CyberSyst J. 2024;1(2):21–31. https://doi.org/10.57238/csj.wr5apn92

64 Al-Ghaili A, Ibrahim Z, Hairi SAS, Rahim FA. A review of anomaly detection techniques in advanced metering infrastructure. Bull Electr Eng Inform. 2021;10(1):266–73. https://doi.org/10.11591/eei.v10i1.2026

Appendix

Appendix 1 | Comparative analysis of anomaly detection algorithms by key parameters

Algorithm	Anomaly Detection Accuracy (95% CI)	Data Processing Speed	Noise Resistance	Applicability to High-Dimensional Data	Experimental Conditions
Autoencoder	87–89% (±1.8%)	Medium	High	Broad	Unlabelled, >30 features, synthetic data, 30% noise
Isolation Forest	84–86% (±1.9%)	High	High	Universal	Unlabelled, 20 features, mixed data, 20% noise
Random Forest	89–91% (±1.5%)	Low	Medium	Requires labelled data	Labelled, 25 features, open data, no noise
SVM	88–90% (±1.6%)	Medium	Medium	Requires labelled data	Labelled, 20 features, open data, 5% noise

Source: Compiled by the authors based on.[2,3]

Appendix 2 | Configuration and training conditions of anomaly detection algorithms

Algorithm	Key Hyperparameters	Model Architecture/Configuration	Training Conditions
Autoencoder	[64–32–16–32–64], Activation = ReLU	Deep encoder-decoder, symmetric layers, MSE loss	Unlabelled, >30 features, synthetic data, 30% noise
Isolation Forest	n_estimators = 200, contamination = 0.05	Random isolation trees, anomaly score thresholding	Unlabelled, 20 features, mixed data, 20% noise
Random Forest	n_estimators = 200, max_depth = 20	Ensemble of decision trees, Gini-based splitting	Requires labelled data, 25 features, open data, no added noise
SVM	C = 1.0, kernel = RBF	Margin-based classifier with radial kernel	Requires labelled data, 20 features, open data, 5% synthetic noise

Source: Compiled by the authors based on.[15,16]

Appendix 3 | Confusion matrices for each model and dataset

Model	Dataset	Accuracy (%)	True Positives (TP)	False Positives (FP)	True Negatives (TN)	False Negatives (FN)	95% Confidence Interval
Autoencoder	SWaT	87–89%	835–855 (87%)	105–125	8,839–8,859	105–125	±2.1%
Isolation Forest	SWaT	84–86%	805–825 (84%)	135–155	8,809–8,829	135–155	±1.9%
Random Forest	SWaT	89–91%	855–875 (89%)	85–105	8,835–8,855	85–105	±1.5%
SVM	SWaT	88–90%	845–865 (88%)	95–115	8,825–8,845	95–115	±1.6%
Autoencoder	TEP	87–89%	835–855 (88%)	105–125	3,855–3,875	105–125	±2.1%
Isolation Forest	TEP	84–86%	805–825 (85%)	135–155	3,840–3,860	135–155	±1.9%
Random Forest	TEP	89–91%	855–875 (89%)	85–105	3,835–3,855	85–105	±1.5%
SVM	TEP	88–90%	845–865 (88%)	95–115	3,825–3,845	95–115	±1.6%
Autoencoder	Synthetic Data	87–89%	835–855 (88%)	105–125	9,085–9,105	105–125	±2.1%
Isolation Forest	Synthetic Data	84–86%	805–825 (85%)	135–155	9,070–9,090	135–155	±1.9%
Random Forest	Synthetic Data	89–91%	855–875 (89%)	85–105	9,055–9,075	85–105	±1.5%
SVM	Synthetic Data	88–90%	845–865 (88%)	95–115	9,045–9,065	95–115	±1.6%