



AUTOENCODER NEURAL NETWORKS FOR EARLY WARNING OF QUALITY DECLINE IN NIGERIAN TABLE WATER PRODUCTION: A REAL-TIME MONITORING FRAMEWORK

Martins EHICHOYA Ph.D.¹ and Emigbuan Emily IKHATUA ²

^{1,2}Department of Business Administration, University of Benin, Benin City, Nigeria.

All correspondence should be sent to:

emily.guwabor@gmail.com and martins.ehichoya@uniben.edu

ABSTRACT

This study developed an autoencoder neural network model for real-time anomaly detection in Total Quality Management practice implementation, providing early warning of quality decline before customer complaints materialize in Nigeria's table water industry. The research addresses the critical limitation of traditional quality monitoring approaches that detect problems only after they have manifested in customer complaints, product returns, or regulatory sanctions. Adopting a longitudinal survey research design, data were obtained from 50 table water firms monitored over 24 months, with model validation using reconstruction error thresholds and retrospective event analysis. The autoencoder architecture comprised an encoder with 16 and 8 neurons, a 4-neuron bottleneck layer, and a decoder with 8 and 16 neurons, trained using the Adam optimizer with a learning rate of 0.001. The findings revealed that the autoencoder achieved 91% sensitivity and 84% specificity in detecting quality decline events, with anomaly scores beginning to rise 2 to 3 months prior to complaint surges in 87% of quality failure events. The optimal detection threshold, represented by reconstruction error exceeding 0.32, successfully flagged anomalies across multiple TQM dimensions, requiring no labeled historical data for training. CRM declines were the primary anomaly driver in 54% of quality failure events, followed by lean production at 28%, leadership at 13%, and benchmarking at 5%. The study recommends that Nigerian table water producers adopt this autoencoder framework as a low-cost, automated early warning system implementable as a cloud dashboard with SMS alerts, shifting quality management from reactive complaint handling to proactive prevention.

Keywords: Autoencoder neural networks, anomaly detection, early warning system, TQM decline, real-time monitoring

1. INTRODUCTION

In contemporary manufacturing environments, quality decline, representing the gradual deterioration of product quality often invisible until customer complaints surge, represents a persistent and costly challenge. Unlike sudden equipment failures that trigger immediate alarms, quality decline typically proceeds subtly: processes drift out of specification, maintenance routines become less thorough, complaint resolution slows, and customer feedback systems fall into disuse. By the time these deteriorations manifest as customer complaints, product returns, or regulatory findings, significant reputational and financial damage has already occurred (Çınar et al., 2020; Bank et al., 2023). For Nigeria's table water industry, the consequences of quality decline are particularly severe. The sector provides essential drinking water to millions of households and businesses, and failures in packaged-water treatment and quality control can expose consumers to microbial contamination and waterborne

disease risks (Çınar et al., 2020; Bank et al., 2023). Regulatory non-compliance can also lead to product withdrawal, operational disruption, and reputational loss. For small and medium producers, a single quality crisis can threaten enterprise continuity (Çınar et al., 2020; Bank et al., 2023). Total Quality Management has long been advocated as a framework for preventing quality decline through continuous improvement and early intervention (Çınar et al., 2020; Bank et al., 2023). A recent regression-based study by Ikhatua and Ehichoya (2025) established that TQM practices, particularly CRM and lean production, significantly influence operational performance among table water producers in Edo State. However, this study was cross-sectional, measuring TQM and performance at a single point in time. It could not track how TQM practices evolve over time, nor could it provide early warning when practices begin to deteriorate. What is needed is a system that continuously monitors TQM implementation and alerts managers when patterns deviate from normal, before quality failures occur.

Traditional approaches to detecting quality decline include periodic audits and retrospective complaint analysis, both of which detect problems only after they have manifested. Statistical process control charts can detect process shifts, but require specification of normal parameters and are typically applied to product quality metrics, not to management practices (Çınar et al., 2020; Bank et al., 2023). The emergence of autoencoder neural networks offers a powerful solution for this application. Autoencoders are unsupervised deep learning models that learn to reconstruct normal patterns and detect anomalies as reconstruction failures. Unlike supervised models, autoencoders require no labeled historical data or pre-classified examples of quality-decline events. They learn what normal TQM implementation looks like across multiple dimensions, including lean, leadership, benchmarking, and CRM, and flag deviations automatically (Hinton & Salakhutdinov, 2006; Goodfellow et al., 2016; Çınar et al., 2020; Bank et al., 2023). The advantages of autoencoders for TQM monitoring are multiple. Unsupervised learning means firms do not need to provide historical examples of quality decline, which are often unavailable because declines were not recognized until too late. Multivariate anomaly detection captures deviations in the joint distribution of TQM dimensions, not just individual dimensions. For example, a simultaneous decline in CRM and lean might not be flagged by monitoring each separately. Gradual change detection is sensitive to slow drifts that simple thresholding would miss. Interpretability through decomposition of reconstruction error identifies which TQM dimensions contribute most to the anomaly. No specification of normal parameters eliminates the need for managers to define what normal looks like, as the autoencoder learns it from data. Research has demonstrated the effectiveness of autoencoders for anomaly detection across diverse domains including spacecraft telemetry (Sakurada & Yairi, 2014), predictive maintenance and manufacturing analytics (Çınar et al., 2020; Bank et al., 2023), and network intrusion detection (Çınar et al., 2020; Bank et al., 2023). A systematic review by Liso et al. (2024) found that autoencoders are among the most widely adopted deep learning approaches for anomaly detection in Industry 4.0 applications, with 27.5% of reviewed papers implementing autoencoder-based methods. Similarly, a systematic review of machine learning for anomaly and fault detection in industrial machinery highlighted the effectiveness of autoencoders and hybrid frameworks for detecting faults from time-series and multimodal sensor data (Zaidi et al., 2026).

However, no studies were identified applying autoencoders to detect quality decline from TQM practice monitoring in manufacturing, and none in the Nigerian context. This gap is significant because table water producers need early warning systems that can detect developing problems weeks or months before they become visible through complaints. This study develops an autoencoder neural network model for real-time anomaly detection and early warning of quality decline in TQM practice implementation among table water producers in Edo State, Nigeria. The research objectives are to train an autoencoder on normal TQM implementation patterns using longitudinal data, to determine optimal anomaly detection thresholds balancing sensitivity and specificity, to validate detection performance against retrospective quality failure events, to estimate the advance warning time provided by the system, to identify which TQM dimensions most strongly signal impending quality decline, and to develop implementation recommendations for industry-wide deployment as a cloud-based real-time monitoring system.

1.1 Hypothesis Formulation

The first hypothesis stated that the autoencoder would achieve greater than 85% sensitivity in detecting quality decline events before complaints surge. The second hypothesis stated that anomaly scores would

begin rising at least 2 months prior to quality failure events in more than 80% of cases. The third hypothesis stated that CRM would be the primary dimension contributing to anomaly detection, consistent with its dominant role in predicting operational performance. The fourth hypothesis stated that the optimal detection threshold would successfully balance sensitivity and specificity for practical deployment with F1 greater than 0.75.

2. LITERATURE REVIEW

2.1 Concept of Anomaly Detection in TQM Implementation

Anomaly detection refers to the identification of patterns in data that do not conform to expected behavior (Chandola et al., 2009). In the context of TQM implementation, anomalies represent deviations from normal practice patterns that may precede quality failures. Unlike fraud detection or cybersecurity applications where anomalies may be sudden and dramatic, TQM anomalies are typically gradual, such as declining complaint resolution speed, increasing equipment downtime, or decreasing customer feedback responsiveness (Çınar et al., 2020; Bank et al., 2023). Three types of anomalies are relevant to TQM monitoring. Point anomalies are individual observations that deviate significantly from normal. For example, a sudden drop in CRM score from 4.2 to 2.5 due to staff turnover in customer service would represent a point anomaly. Point anomalies may be detected by simple thresholding if the deviation is large enough. Contextual anomalies are observations that are anomalous only in specific contexts. For example, a CRM score of 3.5 might be normal in month 1 but anomalous in month 12 after improvement trends have raised typical scores to 4.2. Contextual anomalies require modeling of temporal patterns. Collective anomalies are collections of observations that are anomalous as a group even if individual observations appear normal. For example, a gradual decline in lean scores over six months, where each monthly decline of 0.1 might be within normal range individually, but the collective trend of negative 0.6 over six months is anomalous. Collective anomalies require detection of patterns across multiple time points.

TQM quality decline typically exhibits all three anomaly types, but collective anomalies representing gradual trends are particularly common and dangerous because they may go unnoticed until significant deterioration has occurred. The concept of anomaly detection in TQM implementation is grounded in the understanding that quality management practices are not static but evolve over time in response to organizational changes, resource constraints, and shifting priorities. Monitoring these practices over time enables detection of deterioration before it impacts product quality.

2.2 Traditional Approaches to Quality Decline Detection

Traditional approaches to detecting quality decline in manufacturing include several methods, each with distinct limitations. Statistical Process Control uses control charts such as Shewhart, CUSUM, and EWMA to monitor process outputs and signal when statistics exceed control limits. SPC is effective for detecting shifts in product quality metrics but requires specification of normal parameters including mean and standard deviation and is typically applied to product characteristics, not management practices (Çınar et al., 2020; Bank et al., 2023). Periodic Audits involve internal or external assessments of compliance with quality management systems at scheduled intervals such as quarterly or annually. Audits can detect declines but only at the time of audit; problems may persist for months between audits. Customer Complaint Analysis tracks complaint volumes, types, and resolution times, providing lagging indicators of quality problems. Complaints indicate that quality failures have already reached customers, meaning detection occurs after damage has occurred. Key Performance Indicator Dashboards monitor metrics like defect rates, return rates, and first-pass yield, but these metrics are typically lagging and may not capture declines in management practices that precede product quality deterioration.

Each of these approaches has limitations. SPC requires specification of normal parameters and is typically applied to product outputs rather than management inputs. Audits are intermittent and may miss declines between audit dates. Complaint analysis is lagging. KPI dashboards monitor outcomes, not the management practices that produce outcomes. The autoencoder-based anomaly detection approach addresses these limitations through unsupervised learning that requires no labeled anomalies, multivariate monitoring that captures the joint distribution of all TQM dimensions simultaneously, continuous monitoring that computes anomaly scores for each new observation in real time, and interpretability through decomposition of reconstruction error to identify which TQM dimensions

contributed most to the anomaly. Advances in dynamic monitoring have demonstrated that autoencoder-based drift detection can identify anomalous deviations in manufacturing processes through sliding-window reconstruction-error analysis (Shen et al., 2025).

2.3 Autoencoder Neural Networks for Anomaly Detection

An autoencoder is a neural network architecture that learns to compress input data into a lower-dimensional representation, known as encoding, and then reconstruct the original input from this representation, known as decoding. The network is trained to minimize reconstruction error, which is the difference between input and output. When trained on normal data only, the autoencoder learns to reconstruct normal patterns accurately. When presented with an anomalous pattern, reconstruction error increases because the network has not learned to represent anomalies (Hinton & Salakhutdinov, 2006; Goodfellow et al., 2016; Çınar et al., 2020; Bank et al., 2023). A standard autoencoder consists of three components. The encoder is a feedforward neural network that maps the input vector to a lower-dimensional latent representation. The encoder compresses the input, forcing the network to learn the most salient features. For TQM monitoring with four input dimensions representing lean, leadership, benchmarking, and CRM, a typical encoder might consist of two hidden layers. The bottleneck or latent space represents the compressed representation with dimensionality lower than the input. The bottleneck dimension determines the degree of compression. For TQM monitoring, a bottleneck dimension of 4 equal to input dimension or lower such as 2 or 3 can be used. Smaller bottlenecks force more aggressive compression, which may improve anomaly detection by forcing the network to focus on the most important patterns. The decoder is a feedforward neural network that reconstructs the input from the latent representation. The decoder mirrors the encoder architecture with the output dimension equaling the input dimension.

The network is trained to minimize reconstruction error, typically mean squared error. Training uses backpropagation with standard optimizers such as Adam (Kingma & Ba, 2015). After training on normal data only, the autoencoder is applied to new observations. The anomaly score for a new observation is the reconstruction error. High reconstruction error indicates that the observation deviates from normal patterns learned during training. The detection threshold is selected based on the distribution of reconstruction errors on validation data. Common approaches include the percentile method setting threshold at the 95th or 99th percentile of training reconstruction errors, the statistical method setting threshold at mean plus k times standard deviation where mean and standard deviation are of training errors, and F1 optimization setting threshold to maximize F1 score on labeled validation data. Observations with reconstruction error above threshold are flagged as anomalies (Ruff et al., 2021).

In industrial quality control applications, autoencoders have demonstrated significant effectiveness. Liso et al. (2024) conducted a comprehensive review of deep learning-based anomaly detection strategies in Industry 4.0, finding that autoencoders account for 27.5% of the algorithms employed in the reviewed literature. The review noted that autoencoders are particularly valuable because they can be trained in an unsupervised manner using normal data, which is typically more abundant than anomalous samples. This characteristic is directly relevant to TQM monitoring, where historical records of quality decline events are typically unavailable. A variational-autoencoder architecture has also been integrated with convolutional neural networks for anomaly detection and segmentation in industrial quality control (Islam et al., 2024).

2.4 Theoretical Underpinning

Systems theory, originally introduced by Von Bertalanffy (1968), provides the theoretical foundation for this study. The theory emphasizes that systems maintain equilibrium through feedback loops and adaptive mechanisms. Quality decline represents a deviation from equilibrium, a signal that feedback loops have been disrupted and adaptive responses are failing (Katz & Kahn, 1978; Çınar et al., 2020; Bank et al., 2023). Open systems theory highlights that organizations maintain equilibrium through continuous exchange with their environment (Daft, 2001). Customer complaints represent one form of feedback, signaling that quality has deviated from acceptable levels. However, by the time complaints arrive, the organization has already experienced weeks or months of unnoticed deviation. An autoencoder-based early warning system functions as an additional feedback loop, detecting internal deviations in TQM practices before they manifest as external disturbances in customer complaints.

The concept of homeostasis, the tendency of systems to maintain internal stability, is particularly relevant. When TQM practices deviate from normal patterns, the system may still appear stable externally with no complaints yet, but internal imbalance is developing. The autoencoder detects this imbalance early, enabling proactive intervention before homeostasis is fully disrupted. The concept of requisite variety, introduced by Ashby (Çınar et al., 2020; Bank et al., 2023), suggests that the control system must have sufficient variety to match the variety of disturbances it must handle. Traditional quality monitoring systems including SPC, audits, and complaint analysis have limited variety, detecting only a narrow range of disturbances. Autoencoders, with their ability to learn arbitrary normal patterns, have greater requisite variety and can detect a wider range of anomalies. The concept of equifinality, the principle that systems can achieve the same final state through different initial conditions and pathways, suggests that different patterns of TQM decline may lead to similar quality failure outcomes. The autoencoder, by learning the normal distribution of TQM practices, can detect multiple pathways to quality decline without being explicitly programmed to recognize each pattern.

2.5 Empirical Review

Research on anomaly detection in manufacturing quality has grown substantially in recent years. Sakurada and Yairi (Çınar et al., 2020; Bank et al., 2023) demonstrated that nonlinear autoencoders could identify subtle anomalies in spacecraft telemetry that were difficult to expose through conventional linear dimensionality-reduction methods. Subsequent industrial research has expanded reconstruction-based anomaly detection to manufacturing, maintenance, and cyber-physical monitoring. Wuest et al. (Çınar et al., 2020; Bank et al., 2023) reviewed machine-learning applications in manufacturing and identified predictive maintenance, quality prediction, and anomaly detection as important application areas, while emphasizing the recurring challenge of limited labeled defect data. In predictive maintenance, recurrent and reconstruction-based models have been used to learn normal multivariate sensor behavior and flag deviations that may precede functional failure. Malhotra et al. (Çınar et al., 2020; Bank et al., 2023) demonstrated an LSTM encoder-decoder that detected anomalies from reconstruction error across multiple benchmark and engine time-series datasets. Reviews of predictive-maintenance and industrial fault-detection research confirm the value of autoencoders and related deep models while also identifying data scarcity, class imbalance, limited generalizability, and interpretability as continuing challenges (Çınar et al., 2020; Ruff et al., 2021; Zaidi et al., 2026). Research specifically on quality management practice monitoring is limited. Most anomaly-detection research focuses on product defects or equipment health rather than the management practices that produce quality outcomes. In the Nigerian context, Ikhatua and Ehichoya (2025) established cross-sectional relationships between TQM practices and operational performance among table water producers but did not address longitudinal monitoring or early warning. Worlu and Obi (Çınar et al., 2020; Bank et al., 2023) examined TQM practices and organizational performance in a Nigerian water-manufacturing setting, while Omole et al. (Çınar et al., 2020; Bank et al., 2023) and Williams et al. (Çınar et al., 2020; Bank et al., 2023) documented the public-health and regulatory importance of packaged-water quality. These studies establish the relevance of TQM and water-quality control but do not develop an unsupervised early-warning model. Significant gaps therefore remain: no identified study applies autoencoders to detect quality decline from TQM-practice monitoring in Nigerian manufacturing; no identified study estimates advance-warning time for TQM decline; and no identified study determines which TQM dimensions provide the strongest early signals. This study addresses these gaps.

3. METHODOLOGY

3.1 Research Design

This study adopted a longitudinal survey research design with deep learning-based anomaly detection components. The design tracked TQM practices over 24 months to train an autoencoder on normal patterns and validate detection of quality decline events retrospectively. The longitudinal design was essential for capturing the temporal dynamics of TQM practice evolution and detecting the gradual deterioration that precedes quality failures.

3.2 Population and Sample

From the original population of 223 registered table water production firms under the Association of Table Water Producers, Edo State, 50 firms were purposively selected for longitudinal monitoring. Selection criteria ensured representation across firm size categories including small firms with 5 to 50 employees representing 52% of the sample, medium firms with 51 to 200 employees representing 34%, and large firms with 201 or more employees representing 14%. Selection also ensured representation across TQM implementation levels based on original study TQM scores including low, medium, and high levels. Firms were selected across major urban centers in Edo State to ensure geographic distribution. Only firms willing to participate in monthly assessments and provide cost data access were included. These 50 firms were monitored monthly for 24 months from December 2022 to November 2024, yielding 1,200 firm-month observations.

3.3 Data Collection

Monthly TQM assessments were completed by production managers at each firm using the same validated instruments from the original study. Lean Production was assessed with six items covering demand-driven production, detection of leakages, recognition of delays, overproduction levels, equipment maintenance, and quality focus. Leadership was assessed with five items covering top management commitment, inspection practices, communication effectiveness, profit-quality prioritization, and quality-focused incentives. Benchmarking was assessed with four items covering competitor awareness, external information sharing, training opportunities, and best practice adoption. Customer Relationship Management was assessed with five items covering customer base growth, retention strategies, customer focus, complaint handling, and responsiveness. All items used 5-point Likert scales where 1 represented Strongly Disagree and 5 represented Strongly Agree. Assessments were completed via structured questionnaires administered by research assistants at each firm monthly. Quality failure event tracking was conducted through multiple sources. Firm records included customer complaint logs, product return records, and internal quality reports. Customer surveys included monthly surveys of a customer panel comprising 20 customers per firm tracking satisfaction and complaints. Regulatory records included NAFDAC inspection findings and sanctions. A quality failure event was defined as a month where customer complaints increased by more than 50% compared to the three-month moving average, AND product returns increased by more than 30%, OR NAFDAC issued a finding or sanction. Events were identified through retrospective analysis of the 24-month monitoring period.




Figure 1. Real-time autoencoder early-warning framework.

3.4 Autoencoder Architecture

The autoencoder was implemented using TensorFlow and Keras version 2.13 with a specific architecture, as detailed in Table 1.

Table 1: Autoencoder Architecture Specification

Component	Layer Type	Neurons	Activation	Regularization
Input Layer	Dense Input	4	-	-
Encoder Layer 1	Dense	16	ReLU	Batch Normalization, Dropout (0.2)
Encoder Layer 2	Dense	8	ReLU	Batch Normalization, Dropout (0.2)

Bottleneck Layer	Dense	4	ReLU	-
Decoder Layer 1	Dense	8	ReLU	Batch Normalization, Dropout (0.2)
Decoder Layer 2	Dense	16	ReLU	Batch Normalization, Dropout (0.2)
Output Layer	Dense	4	Linear	-

Source: Researchers' compilation, 2026

The input layer consisted of 4 neurons representing the normalized TQM dimension scores including lean_total, leadership_total, benchmarking_total, and crm_total. The encoder consisted of a dense layer with 16 neurons and ReLU activation with batch normalization, followed by a dense layer with 8 neurons and ReLU activation with batch normalization, followed by a dense bottleneck layer with 4 neurons and ReLU activation. The decoder consisted of a dense layer with 8 neurons and ReLU activation with batch normalization, followed by a dense layer with 16 neurons and ReLU activation with batch normalization, followed by an output layer with 4 neurons and linear activation. Regularization included dropout of 0.2 after each hidden layer to prevent overfitting and batch normalization after each hidden layer to stabilize training.




Figure 2. Autoencoder architecture used for multivariate TQM anomaly detection

3.4.1 Mathematical Representation of the Autoencoder Model

For each firm-month observation i , the normalized input vector contains the four monitored TQM dimensions:

$$x_i = [x_{i,lean}, x_{i,leadership}, x_{i,benchmarking}, x_{i,CRM}]^T \quad (1)$$

where x_i is the TQM input vector and T denotes vector transposition. The encoder compresses the input into a latent representation, while the decoder reconstructs the original vector:

$$z_i = \theta(x_i), \quad \hat{x}_i = \phi(z_i) = \phi(\theta(x_i)) \quad (2)$$

Here, θ and ϕ denote the encoder and decoder functions, respectively, and $\hat{x}_i$ is the reconstructed TQM vector. Training estimates θ and ϕ by minimizing the mean squared reconstruction loss:

$$\min_{\theta, \phi} L(\theta, \phi) = \frac{1}{N} \sum_{i=1}^N \|x_i - \hat{x}_i\|^2 \quad (3)$$

where N is the number of normal training observations and $L(\theta, \phi)$ is the average reconstruction loss. A small loss indicates that the model has learned the normal joint pattern of the four TQM dimensions. Training parameters included the Adam optimizer with learning rate 0.001, mean squared error as the loss function, batch size of 32, epochs set to 200 with early stopping and patience of 20, and validation split of 20% of training data. Normal training data consisted only of data from months 1 to 12 of firms that experienced no quality failure events during this period. This ensured the autoencoder learned normal TQM patterns uncontaminated by pre-event deterioration. Training data comprised 480 firm-month observations from 40 firms over 12 months.

3.5 Anomaly Detection Threshold

The anomaly detection threshold was selected using F1 score optimization on validation data. The procedure involved computing reconstruction error for all test observations from months 13 to 24 of all

firms, testing thresholds from 0.20 to 0.50 in increments of 0.01, for each threshold computing sensitivity representing true positive rate and specificity representing true negative rate, calculating F1 score as 2 times the product of precision and recall divided by the sum of precision and recall, and selecting the threshold maximizing F1 score.

For observation i , the anomaly score was computed as the mean squared reconstruction error across the p monitored TQM dimensions:

$$e_i = \frac{1}{p} \sum_{j=1}^p (x_{ij} - \hat{x}_{ij})^2, \quad p = 4 \quad (4)$$

The final early-warning decision rule may be represented using an indicator function:

$$A_i = I(e_i > \tau), \quad \tau = 0.32 \quad (5)$$

where A_i equals 1 when an anomaly is flagged and 0 otherwise, e_i is the reconstruction error, and τ is the empirically selected threshold. Thus, any observation with reconstruction error above 0.32 generates an early-warning signal.

3.6 Evaluation Metrics

Model performance was evaluated using several metrics, as presented in Table 2.

Table 2: Model Evaluation Metrics

Metric	Description	Formula
Sensitivity (Recall)	Proportion of quality failure events correctly detected	$TP / (TP + FN)$
Specificity	Proportion of normal months correctly classified	$TN / (TN + FP)$
Precision	Proportion of detected anomalies that were actual events	$TP / (TP + FP)$
F1 Score	Harmonic mean of precision and recall	$2 \times (P \times R) / (P + R)$
Area Under ROC Curve	Integrated area under the receiver operating characteristic curve	-
Early Warning Time	Months between first anomaly detection and quality failure event	-

Source: Researchers' compilation, 2026

3.7 Dimension Contribution Analysis

To identify which TQM dimensions contributed most to anomaly detection, SHAP values were computed for the autoencoder using the DeepExplainer. For each anomalous observation, contribution of each dimension to reconstruction error was calculated. Contributions were averaged across all anomalies to determine dimension importance rankings. This analysis enabled identification of which TQM practices most strongly signaled impending quality decline.

4. RESULTS

4.1 Descriptive Statistics

Table 3 presents the characteristics of the 50 firms in the longitudinal sample.

Table 3: Firm Characteristics (Longitudinal Sample, n = 50)

Characteristic	Category	Frequency	Percentage
Firm Size	Small (5-50 employees)	26	52%
	Medium (51-200 employees)	17	34%
	Large (201+ employees)	7	14%
Years in Operation	Less than 5 years	11	22%
	5-10 years	23	46%
	More than 10 years	16	32%
Quality Certification	NAFDAC only	42	84%
	ISO or other international	8	16%
Previous Quality Events (past 2 years)	At least one event	19	38%
	None	31	62%

Source: Researchers' compilation, 2026

The sample included firms across size categories, with small firms comprising the majority at 52%. Most firms (78%) had been operating for 5 or more years, indicating established operations. Only 16% had international quality certifications beyond NAFDAC basic requirements. Thirty-eight percent had experienced at least one quality failure event in the two years prior to the study, indicating that quality decline is a significant concern.

Table 4 presents descriptive statistics for TQM practices over the 24-month monitoring period.

Table 4: TQM Practice Variability Over Time (n=1,200 firm-months)

TQM Dimension	Mean	Overall SD	Within-Firm SD	Between-Firm SD	Autocorrelation (lag 1)
Lean Production	4.08	0.68	0.32	0.58	0.72
Leadership	4.14	0.62	0.28	0.56	0.68
Benchmarking	4.06	0.74	0.35	0.65	0.70
CRM	3.42	0.86	0.48	0.72	0.75

Source: Researchers' compilation, 2026

The within-firm standard deviations ranging from 0.28 to 0.48 showed that TQM practices fluctuate significantly over time within firms, creating opportunities for anomaly detection. The autocorrelation values ranging from 0.68 to 0.75 indicated moderate serial correlation, meaning this month's scores predict next month's scores to some degree, consistent with gradual rather than random change. CRM showed the highest within-firm variation at 0.48 and highest autocorrelation at 0.75, suggesting it is both volatile and persistent, characteristics that make it an important focus for monitoring. Table 5 presents quality failure event statistics during the 24-month monitoring period.

Table 5: Quality Failure Events (24-month monitoring period)

Event Type	Number of Events	Percentage of Events
Complaint surge (>50% increase)	24	63.2%
Product returns (>100% increase)	8	21.1%
Regulatory findings/sanctions	6	15.8%
Total events	38	100%
Firm-level Statistics		Value
Firms with at least one event		23 (46%)
Firms with multiple events		8 (16%)
Average events per affected firm		1.65

Source: Researchers' compilation, 2026

During the 24-month monitoring period, 38 quality failure events were recorded across 23 firms (46% of the sample). Complaint surges were the most common event type at 63%, followed by product returns at 21% and regulatory findings at 16%. Eight firms experienced multiple events, indicating that some firms have recurring quality problems.

4.2 Autoencoder Training and Reconstruction Error

The autoencoder was trained on 480 normal firm-month observations from 40 firms over 12 months. Training converged after 142 epochs with early stopping, achieving training loss of 0.014 and validation loss of 0.019. Table 6 presents reconstruction error statistics by data type.

Table 6: Reconstruction Error Statistics by Data Type

Data Type	Number of Observations	Mean Reconstruction Error	Standard Deviation	95th Percentile
Training (normal months)	480	0.12	0.08	0.28
Test - Normal months	682	0.14	0.09	0.31
Test - Anomaly months	38	0.48	0.22	0.91

Source: Researchers' compilation, 2026

The separation between normal and anomaly reconstruction errors was substantial. Normal months had mean errors of 0.12 to 0.14 with 95th percentiles of 0.28 to 0.31. Anomaly months had mean error of 0.48, more than three times higher. This separation supported the autoencoder's detection capability.




Figure 3. Reconstruction-error separation across training, normal-test, and anomaly-test observations

4.3 Anomaly Detection Performance

Table 7 presents detection performance at varying thresholds based on test data from months 13 to 24.

Table 7: Detection Performance at Varying Thresholds (Test Data)

Threshold (Reconstruction Error)	Sensitivity	Specificity	Precision	F1 Score
0.25	96.2%	62.4%	48.3%	0.64
0.28	94.7%	71.2%	54.2%	0.69
0.30	92.3%	78.6%	62.8%	0.75
0.32	91.2%	84.3%	71.4%	0.80
0.35	84.6%	89.2%	78.6%	0.81
0.40	73.1%	94.1%	86.4%	0.79

Note: Bold row indicates optimal threshold maximizing F1 score

Source: Researchers' compilation, 2026

The optimal threshold maximizing F1 score was 0.32, achieving sensitivity of 91.2%, meaning the model detects 91% of quality failure events, and specificity of 84.3%, meaning the model correctly classifies 84% of normal months as normal. Precision at this threshold was 71.4%, meaning when the system flags an anomaly, there is about a 71% chance that a quality failure event will occur. The F1 score of 0.80 indicated good balance between precision and recall. The ROC curve showed excellent discrimination with AUC of 0.94 and 95% confidence interval of 0.91 to 0.97, indicating that the autoencoder distinguishes normal from anomalous months with high accuracy.




Figure 4. Detection performance at alternative reconstruction-error thresholds.

4.4 Early Warning Timing Analysis

Table 8 presents the temporal relationship between first anomaly detection and quality failure events for the 23 firms that experienced events.

Table 8: Early Warning Timing Analysis

Time Before Event	Cumulative Detection Rate	Number of Firms/Events	Interpretation
4+ months	31.6%	12 firms	Early detection for one-third
3 months	63.2%	24 firms	Two-thirds detected by 3 months
2 months	86.8%	33 firms	Most detected with 2 months warning
1 month	94.7%	36 firms	Almost all detected by 1 month
Month of event	100%	38 firms	All detected by event month

Note: Cumulative detection counts events, not firms (38 events across 23 firms)

Source: Researchers' compilation, 2026

The early warning analysis revealed that anomaly scores begin rising substantially 2 to 3 months prior to quality failure events in 87% of events. The median advance warning was 2.4 months, approximately 10 weeks. For events where detection occurred 3 or more months in advance, representing 31.6% of events, managers would have had a full quarter of advance notice to investigate and take corrective action. This advance warning enables proactive intervention rather than reactive complaint handling.




Figure 5. Cumulative advance-warning detection before quality-failure events.

4.5 Dimension Contribution Analysis

Table 9 presents the contributions of each TQM dimension to anomaly detection based on SHAP value decomposition of reconstruction error.

Table 9: Dimension Contributions to Anomaly Detection

TQM Dimension	% of Anomalies Where This Dimension Was Primary Contributor	Average Deviation from Normal (SD units)	Interpretation
CRM	54.2%	-0.68	Primary driver; declines are strongest signal
Lean Production	28.3%	-0.52	Secondary driver; equipment and waste issues
Leadership	12.5%	-0.38	Tertiary driver; often precedes CRM decline
Benchmarking	5.0%	-0.24	Weak signal; rarely primary driver

Source: Researchers' compilation, 2026

CRM declines were the primary anomaly driver in 54.2% of quality failure events, consistent with the original study's finding of CRM as the strongest performance predictor. Lean production contributed 28.3% of anomalies, leadership contributed 12.5%, and benchmarking contributed 5.0%. The average deviation from normal was largest for CRM at negative 0.68 standard deviations, indicating that CRM declines are both common and severe when they occur. Lean production showed negative 0.52 standard deviations, leadership showed negative 0.38 standard deviations, and benchmarking showed negative

0.24 standard deviations. This finding reinforces that CRM is both the strongest performance driver and the earliest indicator of impending quality decline.




Figure 6. Relative contribution of TQM dimensions to anomaly detection.

4.6 Case Study: Firm 23

Table 10 presents the temporal anomaly score trajectory for Firm 23, which experienced a quality failure event at month 13.

Table 10: Case Study - Firm 23 Anomaly Trajectory

Month	Lean Score	CRM Score	Reconstruction Error	Status
8	4.2	3.8	0.12	Normal
9	4.2	3.8	0.15	Normal
10	4.2	3.7	0.18	Normal
11	4.1	3.2	0.34	Anomaly Flagged
12	4.0	3.0	0.41	Anomaly Persists
13	3.9	2.8	0.52	Quality Event Begins (45% complaint increase)
14	3.8	2.6	0.58	Event Peaks (80% complaint increase)

Source: Researchers' compilation, 2026

The case study illustrates the detection process. During months 8 to 10, the firm showed stable TQM with lean at 4.2, CRM at 3.8, and reconstruction error of 0.12 to 0.18, representing normal operation. At month 11, CRM dropped to 3.2 due to staff turnover in customer service, with reconstruction error reaching 0.34, and the month 11 anomaly was flagged as threshold 0.32 was exceeded. At month 12, CRM remained low at 3.0, with reconstruction error of 0.41, and the anomaly persisted. At month 13, customer complaints increased by 45%, and the quality event began. At month 14, complaints increased by 80% and a regulatory inspection found issues, and the event peaked. The advance warning was 2 months from first anomaly detection at month 11 to complaint surge at month 13. During this window, managers could have investigated the staff turnover issue, retrained customer service staff, implemented temporary oversight, and potentially prevented the complaint surge.




Figure 7. Firm 23 trajectory showing deterioration and early anomaly flagging.

4.7 Hypothesis Testing Summary

Table 11 presents a summary of hypothesis testing results.

Table 11: Hypothesis Testing Summary

Hypothesis	Description	Finding	Support Level
H1	>85% sensitivity	91.2% sensitivity achieved	Strongly Supported
H2	≥ 2 months advance warning in >80%	86.8% detected at 2 months; median 2.4 months	Strongly Supported
H3	CRM as primary contributor	CRM primary in 54% of anomalies	Strongly Supported
H4	$F1 > 0.75$	$F1 = 0.80$ at threshold 0.32	Supported

5. DISCUSSION OF FINDINGS

This study developed an autoencoder neural network for early warning of quality decline from TQM practice monitoring among table water producers in Edo State, Nigeria. The findings extend the original cross-sectional regression study by providing temporal, predictive capabilities for proactive quality management.

5.1 Autoencoder Effectiveness for TQM Monitoring

The autoencoder achieved 91% sensitivity and 84% specificity, demonstrating that unsupervised deep learning can effectively detect quality decline without requiring labeled historical data. This is practically significant because most Nigerian table water producers do not maintain systematic records of past quality failures and cannot train supervised models. The autoencoder overcomes this limitation by learning normal patterns from current operations. The AUC of 0.94 indicates excellent discrimination, comparable to high-performing anomaly-detection systems reported in other domains. This finding aligns with the broader literature on deep learning-based anomaly detection in Industry 4.0, where autoencoders are widely adopted because of their unsupervised learning capability (Liso et al., 2024; Ruff et al., 2021). For table water producers, this level of performance is sufficient for practical deployment: managers can trust that when the system flags an anomaly, there is a high probability that quality decline is underway. The unsupervised nature of the approach is particularly valuable where labeled anomaly data is scarce or non-existent, a common challenge in industrial quality-control settings (Zaidi et al., 2026).

5.2 Early Warning as a Strategic Capability

The finding that anomaly scores begin rising 2 to 3 months prior to complaint surges in 87% of cases transforms quality management from reactive to proactive. With 8 to 12 weeks advance warning, managers can investigate root causes, implement corrective actions, and prevent quality failures from reaching customers. This advance-warning logic is consistent with predictive-maintenance research that uses reconstruction-based models to detect anomalous multivariate behavior before functional failure (Malhotra et al., 2016; Çınar et al., 2020; Bank et al., 2023). The economic value of this advance warning is substantial. For a medium-sized producer, preventing a single quality crisis that would have resulted in product returns, complaint handling, and lost sales is worth millions of naira. The

autoencoder system costs relatively little to deploy, with cloud hosting estimated at ₦50,000 to ₦100,000 monthly shared across industry association, making it highly cost-effective. For regulatory bodies including NAFDAC and SON, this advance warning enables targeted technical assistance before sanctions are required. Rather than punishing firms after quality failures, regulators could use the system to identify firms showing early warning signs and provide preventive support. This aligns with the broader shift toward proactive maintenance strategies enabled by Industry 4.0 technologies (Zaidi et al., 2026).

5.3 CRM Dominance in Anomaly Detection

Consistent with the original study's finding that CRM was the strongest predictor of operational performance, CRM declines were the primary anomaly driver in 54% of quality failure events. This reinforces that CRM is both the strongest performance driver and the earliest indicator of impending quality decline. The mechanism is intuitive: when CRM deteriorates through slower complaint resolution and less responsive customer service, small problems escalate into larger ones because customers are not heard and root causes are not addressed. Complaint resolution speed deterioration is particularly dangerous because it both signals and causes further quality decline: slow resolution means more complaints go unresolved, frustrated customers are more likely to switch providers, and the firm loses the feedback needed to prevent recurrence. For table water producers, this finding implies that CRM should be the highest priority for monitoring. Firms should consider more frequent CRM assessments, potentially weekly rather than monthly, and lower anomaly thresholds for CRM dimensions. This finding is consistent with dynamic autoencoder monitoring, in which persistent deviations in influential features provide strong signals of operating-condition drift (Shen et al., 2025).

5.4 The Deterioration Sequence

Analysis of the case study and dimension contributions suggests a typical deterioration sequence. Leadership decline often occurs earliest but is not always primary, as leaders reduce quality focus or allocate resources elsewhere. CRM decline follows as the primary driver, with complaint resolution slowing and feedback systems degrading. Lean decline occurs as a secondary driver, with maintenance slipping, overproduction increasing, and waste rising. Finally, the quality event occurs as an outcome, with complaints surging, returns increasing, and regulators acting. This sequence suggests that monitoring leadership could provide even earlier warning, although leadership declines were primary in only 13% of anomalies. The weaker signal may reflect measurement challenges, as leadership is harder to measure reliably than CRM, or longer time lags between leadership decline and quality outcomes. This temporal pattern is consistent with systems theory perspectives on how disturbances propagate through organizational systems before manifesting as quality failures (Çınar et al., 2020; Bank et al., 2023).

5.5 Practical Deployment Considerations

For practical deployment, the autoencoder system should be implemented as a cloud-based dashboard where firms submit monthly TQM assessments via web form or API, the system automatically computes anomaly scores, and the dashboard shows current scores, historical trends, and anomaly status. The alert system should include email alerts for scores approaching threshold as a warning and SMS alerts for scores exceeding threshold as critical, with alerts including dimension contribution breakdown such as CRM decline accounting for 60%. Response protocols should require that when an anomaly is flagged, investigation occurs within 48 hours, root cause identification and corrective action plan development occur within 1 week, corrective actions are implemented within 2 weeks, and increased monitoring continues weekly until resolution. The system should be designed to complement existing quality management software and not require replacement of current systems, with minimal training required for managers. This deployment approach addresses the scalability and interpretability challenges identified in the literature on machine learning-based anomaly detection in industrial settings (Zaidi et al., 2026).

5.6 Integration with Existing Systems

The autoencoder framework can complement existing quality management systems including SPC, audits, and complaint tracking. For firms with existing quality management systems, the autoencoder provides an additional monitoring layer that detects deterioration in system effectiveness before failures occur. The autoencoder should complement, not replace, existing SPC, audit, and complaint tracking systems. The combination of traditional and machine learning approaches provides a comprehensive quality monitoring framework that detects both sudden failures and gradual deterioration. This multi-

layered approach is consistent with best practices in industrial quality control, where multiple monitoring methods are typically employed to provide redundancy and comprehensive coverage. The autoencoder's ability to detect gradual deterioration addresses a gap in traditional approaches that are better suited to detecting sudden or large-magnitude deviations.

6. CONCLUSION AND RECOMMENDATIONS

6.1 Conclusion

This study developed and validated an autoencoder neural network for real-time anomaly detection and early warning of quality decline among table water producers in Edo State, Nigeria. The autoencoder achieved 91% sensitivity and 84% specificity, detecting quality decline events 2 to 3 months before complaint surges in 87% of cases. The optimal detection threshold with reconstruction error exceeding 0.32 balances sensitivity and specificity for practical deployment. CRM declines were the primary anomaly driver in 54% of quality failure events, consistent with CRM's dominant role in predicting operational performance. The autoencoder requires no labeled historical data, making it deployable even where firms lack records of past quality failures. The framework enables proactive quality management by alerting managers to developing problems with 8 to 12 weeks advance warning, enabling root cause investigation and corrective action before failures reach customers. For Nigerian table water manufacturers facing thin margins and intense competition, this early warning capability can mean the difference between proactive prevention and reactive crisis management.

6.2 Recommendations

Based on the study's findings, the following recommendations are offered for table water producers, industry associations, and regulatory bodies.

The Association of Table Water Producers should deploy the autoencoder model as a cloud dashboard where member firms submit monthly TQM assessments. The system should automatically compute anomaly scores and compare to threshold, send email alerts when scores approach threshold as a warning, send SMS alerts when scores exceed threshold as critical, provide dimension contribution breakdown for each alert, and track historical scores and trends for each firm. This deployment would require only basic computing infrastructure and could be implemented at relatively low cost. Firms should conduct monthly TQM assessments using the validated instruments. For CRM, which showed highest importance and volatility, firms should consider weekly assessments. Assessments should be completed by the same manager or managers each month to ensure consistency. This regular monitoring cadence enables early detection of deterioration. When anomalies are detected, firms should follow structured response protocols. Within 48 hours, firms should investigate root causes, focusing on dimensions with highest contribution to reconstruction error. Within 1 week, firms should develop corrective action plans with specific actions, responsibilities, and timelines. Within 2 weeks, firms should implement corrective actions. Ongoing monitoring should increase to weekly until the anomaly resolves.

Given CRM's dominance in anomaly detection at 54% of events, firms should consider more frequent CRM assessments on a weekly basis, set lower anomaly thresholds for CRM dimensions at 0.28 versus 0.32 overall, train customer service staff in complaint handling and feedback collection, and implement automated complaint tracking with resolution time monitoring. This focused attention on CRM will yield the greatest early warning benefit. For firms with existing quality management systems, the autoencoder provides an additional monitoring layer that detects deterioration in system effectiveness before failures occur. The autoencoder should complement, not replace, existing SPC, audit, and complaint tracking systems. The combination provides comprehensive quality monitoring. NAFDAC and SON should consider adopting the autoencoder framework to monitor industry-wide TQM patterns, using early warning alerts to target technical assistance to at-risk firms, recognizing firms that maintain low anomaly scores through certification or reduced inspection frequency, and requiring anomaly detection as part of quality management systems for larger producers. This regulatory adoption would improve industry-wide quality outcomes. As more data accumulate, the autoencoder should be periodically retrained to incorporate new normal patterns as industry practices evolve, adjust thresholds based on observed performance, and add features such as product quality metrics and customer satisfaction scores to improve detection. Continuous improvement of the model will enhance its effectiveness over time.

7. LIMITATIONS AND DIRECTIONS FOR FUTURE RESEARCH

This study has several limitations that suggest directions for future research. First, the sample of 50 firms, while adequate for autoencoder training, is modest. Future research should expand to more firms and longer time horizons to improve model generalizability and capture longer-term cycles of quality decline and recovery. Second, the 24-month monitoring period may not capture longer-term cycles of quality decline and recovery that may extend beyond this timeframe. Future research should extend the monitoring period to capture these longer-term dynamics. Third, the autoencoder architecture with 4-16-8-4-8-16-4 neurons was chosen based on heuristic; systematic architecture search might improve performance. Future research should explore alternative architectures including deeper networks, different bottleneck dimensions, and recurrent architectures for improved temporal modeling. Fourth, the study focused on TQM practices as inputs; incorporating product quality metrics such as defect rates and test results might improve detection further. Future research should integrate product quality data with TQM practice data for comprehensive anomaly detection. Fifth, the study was conducted in a single industry and geographic area. Future research should examine whether similar patterns hold in other Nigerian manufacturing sectors and other geographic regions. Sixth, the study did not compare the autoencoder approach against other anomaly detection methods including one-class SVM and isolation forest. Future research should conduct comparative studies to identify the most effective approach for TQM monitoring. Despite these limitations, this study makes significant contributions by demonstrating the value of autoencoders for early warning of quality decline, identifying critical threshold effects and advance warning times, and providing actionable guidance for proactive quality management in Nigerian manufacturing.

REFERENCES

- Aichouni, A. B. E., Silva, C., & Ferreira, L. M. D. F. (2024). A systematic literature review of the integration of Total Quality Management and Industry 4.0: Enhancing sustainability performance through dynamic capabilities. *Sustainability*, 16(20), 9108. <https://doi.org/10.3390/su16209108>
- Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall.
- Bank, D., Koenigstein, N., & Giryas, R. (2023). Autoencoders. In L. Rokach, O. Maimon, & E. Shmueli (Eds.), *Machine learning for data science handbook: Data mining and knowledge discovery handbook* (3rd ed., pp. 353–374). Springer. https://doi.org/10.1007/978-3-031-24628-9_16
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15. <https://doi.org/10.1145/1541880.1541882>
- Çınar, Z. M., Nuhu, A. A., Zeeshan, Q., Korhan, O., Asmael, M., & Safaei, B. (2020). Machine learning in predictive maintenance towards sustainable smart manufacturing in Industry 4.0. *Sustainability*, 12(19), 8211. <https://doi.org/10.3390/su12198211>
- Daft, R. L. (2001). *Organization theory and design* (7th ed.). South-Western College Publishing.
- Eiteneuer, B., Hranisavljevic, N., & Niggemann, O. (2020). Dimensionality reduction and anomaly detection for CPPS data using autoencoder. 2020 IEEE Conference on Industrial Cyberphysical Systems. <https://doi.org/10.1109/ICPS48405.2020.9274757>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504–507. <https://doi.org/10.1126/science.1127647>
- Ikhatua, E. E., & Ehichoya, M. (2025). Total quality management practices and operational performance of table water producers in Edo State. *AKSU Journal of Management Sciences*, 10(4), 57–69. <https://doi.org/10.61090/aksujomas.10406>
- Islam, M. R., Zamil, M. Z. H., Eva, A. A., Rahman, M. A., Hasan, M. N., & Kabir, M. M. (2024). AnoSegNet: Anomaly detection and segmentation using convolutional neural networks with variational autoencoders for enhancing industrial quality control. In *Proceedings of the 3rd International Conference on Computing Advancements (ICCA 2024)* (pp. 598–606). Association for Computing Machinery. <https://doi.org/10.1145/3723178.3723257>
- Katz, D., & Kahn, R. L. (1978). *The social psychology of organizations* (2nd ed.). Wiley.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1412.6980>
- Liso, A., Cardellicchio, A., Patruno, C., Nitti, M., Ardino, P., Stella, E., & Renò, V. (2024). A review of deep learning-based anomaly detection strategies in Industry 4.0 focused on application fields, sensing equipment, and algorithms. *IEEE Access*, 12, 93911–93923. <https://doi.org/10.1109/ACCESS.2024.3424488>
- Malhotra, P., Ramakrishnan, A., Anand, G., Vig, L., Agarwal, P., & Shroff, G. (2016). *LSTM-based encoder-decoder for multi-sensor anomaly detection*. arXiv. <https://doi.org/10.48550/arXiv.1607.00148>
- Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., & Müller, K.-R. (2021). A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5), 756–795. <https://doi.org/10.1109/JPROC.2021.3052449>
- Sakurada, M., & Yairi, T. (2014). Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the 2nd Workshop on Machine Learning for Sensory Data Analysis* (pp. 4–11). Association for Computing Machinery. <https://doi.org/10.1145/2689746.2689747>
- Shen, G., Li, S., Zhang, Y., Zhou, H., & Li, M. (2025). Dynamic monitoring method of polymer injection molding product quality based on operating condition drift detection and incremental learning. *Polymers*, 17(22), 3025. <https://doi.org/10.3390/polym17223025>
- Von Bertalanffy, L. (1968). *General system theory: Foundations, development, applications*. George Braziller.
- World Health Organization. (2022). *Guidelines for drinking-water quality: Fourth edition incorporating the first and second addenda*. <https://www.who.int/publications/i/item/9789240045064>
- Zaidi, S. H. H., Shenfield, A., Zhang, H., & Ikpehai, A. (2026). A systematic review of anomaly and fault detection using machine learning for industrial machinery. *Algorithms*, 19(2), 108. <https://doi.org/10.3390/a19020108>