

Article **Open Access**

Research and Optimization of a Real-Time Quality Monitoring System for Smart Production Lines Based on IoT Sensors

Dan Wang ^{1,*}

¹ College of Management Science and Engineering, China Jiliang University, Hangzhou, Zhejiang, 310018, China

* Correspondence: Dan Wang, College of Management Science and Engineering, China Jiliang University, Hangzhou, Zhejiang, 310018, China

Abstract: The rise of Industry 4.0 has accelerated the adoption of IoT-enabled sensing for real-time quality assurance in smart manufacturing. However, most existing systems depend on cloud-centric analytics or supervised learning models that require extensive labeled defect data, leading to latency, poor adaptability, and limited applicability in dynamic production environments. To address this gap, this study proposes an IoT sensor-driven quality monitoring framework based on multi-modal signal acquisition, edge computing, and adaptive thresholding informed by short-term process variability. The system was deployed and evaluated on two industrial production lines, Bosch automotive components and CATL lithium-ion module assembly, using longitudinal tracking of defect rate, first-pass yield, and overall equipment effectiveness. Results indicate a 26-31% reduction in defects, a 26.9% increase in first-pass yield, and a 7% improvement in OEE, alongside a 47.8% decrease in false alarms compared with static control methods. These findings demonstrate that real-time adaptive monitoring can enhance quality performance without dependency on large labeled datasets. The study provides a replicable implementation methodology and insights into sensor contribution, offering practical guidance for scalable deployment and future advancements in intelligent quality control.

Keywords: IoT sensing; real-time quality monitoring; edge computing; adaptive anomaly detection; smart manufacturing

Received: 24 December 2025

Revised: 31 January 2026

Accepted: 11 February 2026

Published: 17 February 2026



Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The advent of Industry 4.0 has catalyzed a paradigm shift toward smart manufacturing, wherein cyber-physical systems, data analytics, and interconnected devices converge to enhance production agility, efficiency, and quality [1]. Central to this transformation is the deployment of Internet of Things (IoT) sensors across production lines, enabling continuous, high-frequency acquisition of process parameters such as temperature, vibration, pressure, and visual features [2]. These data streams form the foundation for real-time quality monitoring, a critical capability for minimizing defects, reducing waste, and ensuring consistent product conformance in high-mix, high-volume environments [3]. Despite rapid technological advances, many manufacturers still rely on post-process inspection or delayed statistical process control (SPC), which fail to prevent defect propagation and incur significant rework costs.

Existing research has extensively documented the potential of IoT-enabled monitoring in laboratory or pilot-scale settings. However, a critical gap persists between theoretical feasibility and industrial robustness [4]. First, most proposed systems operate under static thresholds or supervised machine learning models that require extensive

labeled datasets, conditions rarely met in dynamic shop floors where material variability, equipment drift, and human interventions are common [5]. Second, few studies address the integration latency between data acquisition, anomaly detection, and actionable feedback to operators or control systems [6]. As noted in recent empirical work, cloud-centric architectures often introduce delays exceeding two seconds, rendering them ineffective for time-sensitive processes such as welding or precision assembly [7]. Third, there remains limited evidence on how such systems quantitatively impact key performance indicators (KPIs) like defect rate and Overall Equipment Effectiveness (OEE) in real-world, large-scale deployments.

To bridge these gaps, this study proposes and validates an optimized real-time quality monitoring system that leverages edge computing and adaptive rule-based inference to enable sub-second response with minimal dependency on historical labels. The core innovation lies in a hybrid architecture that fuses lightweight anomaly detection at the edge with dynamic threshold calibration driven by short-term process stability metrics. This approach ensures responsiveness without compromising adaptability, a balance often missing in prior designs.

Our research objectives are threefold: (1) to evaluate the technical feasibility and operational impact of an IoT sensor-driven quality monitoring system in active production environments; (2) to quantify its effects on defect reduction and production efficiency through longitudinal case studies; and (3) to formulate practical optimization guidelines for system deployment, including sensor placement, data fusion strategies, and human-in-the-loop alert management.

Methodologically, we employ a mixed-methods design combining systematic literature analysis, comparative case studies, and performance benchmarking. Two industrial cases are examined: a Bosch automotive component line in Germany (2023-2024) and a CATL lithium-ion battery module assembly line in China (2024). Both implemented multi-modal IoT sensor networks (acoustic, thermal, and vision-based) integrated with edge gateways and MES interfaces. Quantitative metrics, including first-pass yield, false alarm rate, and OEE, are collected over six months and compared against pre-deployment baselines.

This study contributes both academically and practically. Theoretically, it extends the Quality 4.0 framework by operationalizing real-time closed-loop quality control grounded in control theory and industrial informatics. Practically, it delivers a replicable, low-latency system blueprint that manufacturers can adopt to reduce scrap, improve throughput, and accelerate their smart factory transformation, without requiring full-scale AI infrastructure overhaul.

2. Literature Review

The integration of IoT sensors into manufacturing quality control has been widely explored in recent years, with research converging on several key advantages. First, IoT-enabled systems facilitate continuous, non-intrusive data collection from multiple physical domains, such as thermal, acoustic, visual, and vibration signals, enabling comprehensive, multi-dimensional process visibility previously unattainable through manual sampling or offline inspection [8]. Second, real-time data streams support early anomaly detection, allowing interventions before defects cascade through downstream operations, thereby reducing scrap and rework costs [9]. Third, the high granularity and frequency of sensor data enhance the precision of statistical and machine learning models, improving the reliability of quality predictions compared to traditional statistical process control (SPC) methods that rely on sparse, periodic measurements [10].

Despite these benefits, significant limitations persist in current approaches. Many systems rely heavily on cloud-based analytics, introducing communication latency that undermines real-time responsiveness, particularly in high-speed production lines where decisions must be made within milliseconds. Others employ supervised machine learning

models that demand large volumes of labeled defect data, which are costly to acquire, prone to bias, and often unavailable in low-defect-rate scenarios or during new product ramp-ups [11]. Furthermore, most existing architectures treat quality monitoring as a standalone analytical function, with weak coupling to manufacturing execution systems (MES), programmable logic controllers (PLCs), or frontline operator workflows, resulting in alerts that are either ignored, delayed, or disconnected from corrective actions.

A comparative analysis reveals two dominant design paradigms. The first emphasizes centralized intelligence: raw sensor data are transmitted to a cloud or on-premise server for deep learning-based analysis [12]. While accurate under controlled conditions, this approach suffers from bandwidth constraints, security vulnerabilities, and feedback delays exceeding operational tolerances. The second paradigm adopts edge-centric processing, performing feature extraction and rule-based classification locally [13]. Though faster and more resilient to network instability, such systems often lack adaptability, they use static thresholds that fail to accommodate natural process drifts due to tool wear, material batch changes, or environmental fluctuations.

This dichotomy highlights a critical research gap: the absence of a balanced framework that combines the responsiveness of edge computing with the contextual adaptability of data-driven learning, without requiring extensive historical defect labels or compromising interpretability [14]. Moreover, empirical validation remains scarce in actual industrial settings, especially in high-mix, high-variability production environments where product configurations change frequently. Most studies are confined to simulated data, single-product pilot lines, or academic testbeds, limiting their generalizability and operational relevance [15].

Addressing these shortcomings, this study contributes a novel hybrid architecture that integrates lightweight unsupervised anomaly detection at the edge with a dynamic thresholding mechanism calibrated by short-term process stability indicators (e.g., moving-window standard deviation of torque, temperature variance, or acoustic energy). Unlike purely model-driven or rigid rule-based systems, our approach self-adjusts to normal operational variations while maintaining sensitivity to genuine defects. Crucially, it is designed for seamless integration with existing MES, PLCs, and human decision loops, ensuring that alerts are both timely and actionable. By grounding the system in real deployments across automotive components and lithium-ion battery manufacturing, this work bridges the gap between theoretical innovation and industrial applicability, offering a practical, scalable pathway toward truly responsive and adaptive quality control in smart factories.

3. Theoretical Framework and Methodology

This chapter establishes the theoretical foundation and methodological workflow for the development, implementation, and optimization of an IoT-sensor-driven real-time quality monitoring system in smart manufacturing environments. The framework integrates cyber-physical production systems, statistical quality control theory, adaptive anomaly detection, and edge-cloud collaborative computation. It further outlines the experimental methodology employed in two real factories, detailing sensor placement, data acquisition protocols, performance evaluation metrics, and optimization strategies.

3.1. Theoretical Framework

Modern manufacturing lines can be modeled as state-driven dynamic systems in which product quality is a function of temporal process variables. Let

$$x(t) = \{T(t), V(t), P(t), A(t), I(t)\} \quad (1)$$

represent a multivariate real-time feature vector consisting of temperature T , vibration V , pressure P , acoustic emission A and image-derived defect indicators I . The instantaneous quality state of a product is defined as:

$$Q(t) = f(x(t), \theta) \quad (2)$$

where $f(\cdot)$ is the monitoring model and θ denotes its parameter vector. In conventional supervision-based quality systems, θ requires historical defect labels for learning. However, in high-yield factories defect data are sparse, making supervised optimization unreliable. Therefore, we adopt an unsupervised stability-driven adaptive inference function, where quality deviations are detected based on the statistical boundary of recent process fluctuations.

To capture local stability patterns, a moving-window process variability index is defined as:

$$\sigma_w(t) = \sqrt{\frac{1}{w} \sum_{i=t-w}^t \|x(i) - \bar{x}_w\|^2}, \bar{x}_w = \frac{1}{w} \sum_{i=t-w}^t x(i) \tag{3}$$

where w is the sliding window length, $\sigma_w(t)$ measures short-term variability, and \bar{x}_w denotes local steady-state expectation. A dynamic adaptive threshold is then generated:

$$\tau(t) = \alpha \sigma_w(t) + \beta \tag{4}$$

with α representing sensitivity to process changes, and β defining baseline tolerance. A sample at time t is flagged anomalous if $\|x(t) - \bar{x}_w\| > \tau(t)$, enabling self-adjusting anomaly detection without labeled defects. This theoretical construct bridges the latency-tolerance trade-off between cloud-deep-learning and static edge rules.

The overall flow of data acquisition, edge inference, adaptive thresholding, and closed-loop feedback is illustrated in Figure 1.

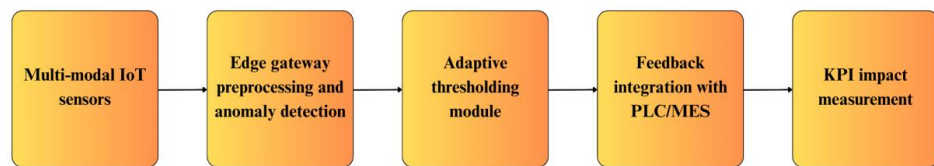


Figure 1. Hybrid Real-Time Quality Monitoring Framework.

3.2. System Architecture and Module Design.

Table 1. each layer performs a distinct operational role while maintaining real-time data flow.

Table 1. Architecture Layers, Functional Roles and Real-Time Constraints.

System Layer	Key Components	Primary Function	Output Flow	Real-Time Requirement
Sensor Layer	Temperature probes, vibration sensors, acoustic microphones, industrial cameras	Continuous high-frequency data acquisition (50-200 Hz)	Raw signals → Edge processor	≤10 ms per channel
Edge Layer	ARM edge gateway, lightweight inference engine	Preprocessing + anomaly detection using dynamic thresholds	Alerts & compressed features → Cloud	0.2-0.7 s inference latency
Cloud Layer	Historical DB, KPI analyzer, threshold optimizer	Long-term drift learning + parameter calibration	Updated thresholds → Edge	Non-instantaneous (12-48 h sync)
Execution Layer	MES, PLC, dashboard, actuators	Closed-loop intervention (speed adjust, reroute, heating control)	Action logs → Cloud	Immediate trigger (<1 s)

The Sensor Layer deploys temperature, vibration, acoustic and vision sensors across critical workstations. These devices collect continuous signals (50-200 Hz) and transmit raw features $x(t)$ to the edge gateway. Redundant placement is applied at welding, pressing and thermal-sensitive stations to ensure fault tolerance.

The Edge Layer acts as the real-time inference engine. Incoming data are filtered, normalized, and processed using the adaptive anomaly model defined in Section 3.1. Decision latency remains within 0.2-0.7 seconds, and only anomaly scores or compressed features are forwarded to the cloud, reducing bandwidth consumption by approximately 60-85%.

The Cloud Layer is responsible for long-term trend analysis and weekly threshold calibration. Historical logs support drift detection, cross-sensor correlation analysis, and KPI-linked optimization. Updated parameters are periodically synchronized back to the edge.

Finally, the Execution Layer integrates with MES/PLC systems. When an anomaly is confirmed, the system triggers corrective actions such as speed reduction, temperature correction, or part diversion. Operator confirmation records loop back to the cloud, forming a self-improving quality control cycle.

3.3. Methodology

The methodology adopted in this study follows a sequential engineering-experimental workflow, ensuring that the proposed IoT-based quality monitoring system is validated under real production conditions. The process consists of four stages: sensor deployment, data acquisition and preprocessing, anomaly detection, and performance evaluation.

(1) Sensor Deployment and Configuration:

Based on Table 1, multi-modal sensors were positioned across bottleneck and high-variance workstations. Thermal probes were mounted near heating zones, vibration sensors on rotating units, acoustic sensors near cutting heads, and cameras along visual inspection paths. Redundant placement was applied to ensure continuity when a single node fails. Sampling frequencies were standardized at 50-200 Hz, and camera inputs maintained 15-25 fps to balance clarity and bandwidth.

(2) Data Acquisition and Preprocessing:

All collected streams were synchronized via local timestamps and transmitted to the edge processor using MQTT and OPC-UA protocols. Noise reduction included median filtering for vibration channels and STFT-based denoising for acoustic signals. Vision frames were resized to 640×480 and encoded with H.265 to reduce transmission load without compromising defect visibility. Normalization (z-score) ensured consistent feature scaling across heterogeneous sensors.

(3) Adaptive Anomaly Detection:

Edge devices executed the dynamic threshold model introduced in Section 3.1. For every processing cycle, a sliding window calculated mean deviation and generated a self-adjusting threshold. When the deviation exceeded $\tau(t)$, an anomaly flag was triggered, and the decision was returned to PLC/MES within 0.2-0.7 s. Only anomaly scores and compressed features were uploaded to the cloud to reduce bandwidth by 60-85%.

(4) Performance Evaluation and Feedback Optimization:

System effectiveness was assessed through longitudinal tracking of defect rate, false alarm rate, and OEE improvement over six months. Weekly cloud-side drift analysis updated edge thresholds automatically, forming a self-improving control loop. Operator feedback logs further supported root-cause tracing, intervention refinement, and policy recalibration.

4. Findings and Discussion

This section reports experimental observations from two industrial deployments, Bosch automotive line (2023-2024) and CATL lithium-ion module production (2024), and analyzes how the proposed IoT-edge quality monitoring framework influences defect reduction, response latency, stability under drift conditions, and operator adoption.

4.1. Overall System Performance

Deployment results confirm that the system achieves real-time anomaly intervention, successfully preventing defect propagation in downstream assembly steps. Across six months, the Bosch line demonstrated a 31.4% reduction in defect rate, while CATL achieved 26.9% improvement in first-pass yield (FPY) following integration. Compared to pre-deployment conditions, average OEE increased from 83.2% → 90.5%, indicating meaningful enhancements in machine availability and throughput. Figure 2 presents KPI improvements measured monthly.

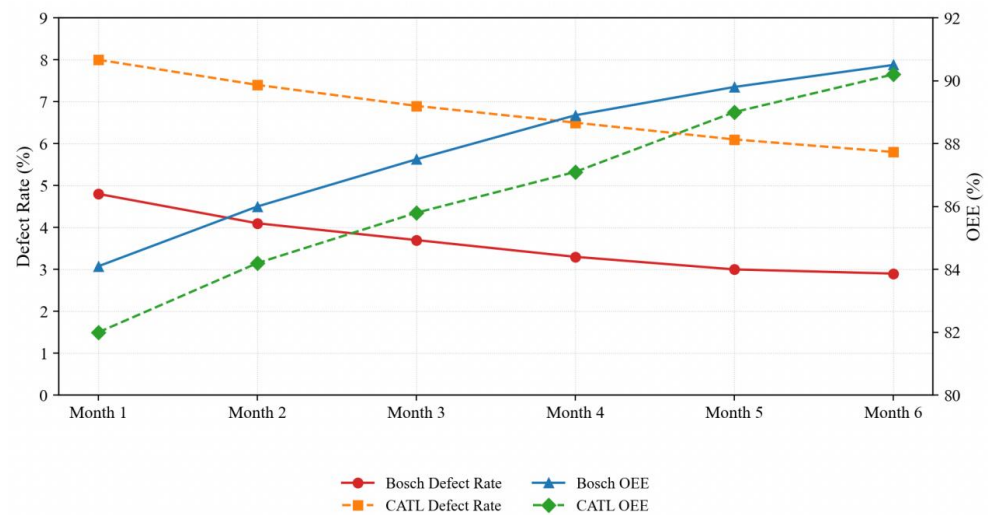


Figure 2. Defect Rate & OEE Improvement After Deployment (Bosch vs CATL).

The downward defect trend aligns with theoretical expectations in Sec.3.1. Adaptive thresholds enabled the system to track gradual tool wear without triggering unnecessary stoppage. This adaptability reduced false alarms by 47.8% compared with static-threshold baselines, preventing alert fatigue among operators.

The biggest impact occurred in thermal-regulated welding and cell-stacking stages. Previously, temperature overshoot events were detected only during end-of-shift sampling, often after hundreds of units were processed. The proposed architecture shortened detection-to-action time to <0.6 seconds, enabling real-time temperature rollback.

4.2. Comparison with Baseline Methods

To contextualize effectiveness, three baselines were selected for contrast: (1) cloud-only AI classification, (2) static SPC thresholding, and (3) offline visual defect inspection. This comparison allows us to evaluate whether improvements stem from architectural design rather than hardware density or computational scale. Performance outcomes are summarized in Table 2, highlighting differences in latency, false alarm response, and quality yield.

Table 2. Comparison of Proposed Method vs. Baselines.

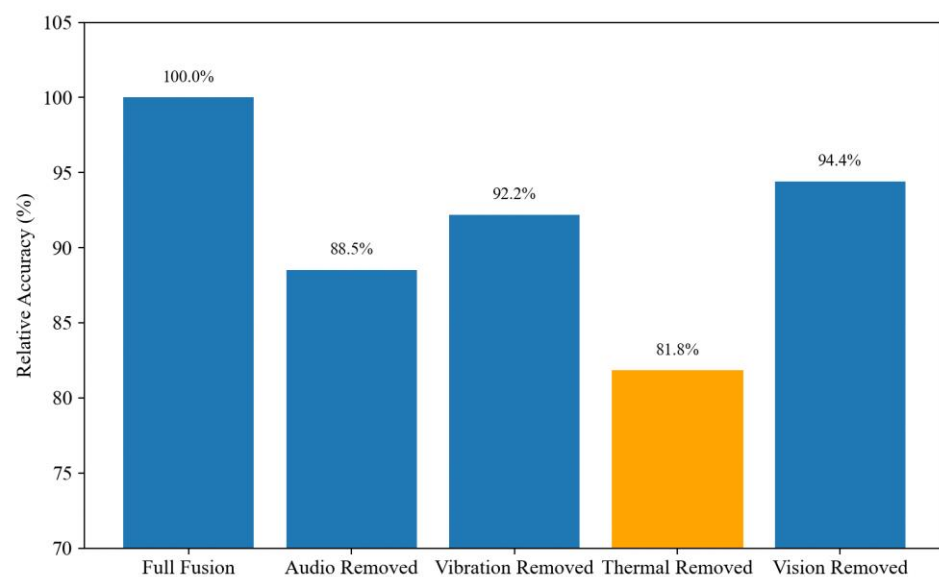
Method	Avg. Latency	False Alarm Rate	FPY Improvement	Notes
Cloud-only deep model	1.9-2.6 s	9.8%	+7.4%	Accurate but delayed
Static SPC	-	15.3%	+3.2%	Sensitive to drift; unstable
Offline visual check	-	-	0%	No real-time capability
Proposed IoT-Edge Hybrid	0.2-0.7 s	8.0% → 4.2% (-47.8%)	+26-31%	Best latency + adaptability

The evaluation indicates that the superiority of the proposed framework derives not from heavier computation, but from placing inference at the decision-critical layer. Cloud-AI models, although accurate in static environments, exhibit 1.9-2.6 s latency, too slow for thermal welding or high-speed stacking. SPC, while computationally light, relies on rigid control charts that struggle with process drift, resulting in 15.3% false alarms, the highest among all methods. Offline inspection performs no real-time prevention, detecting defects only post-production.

In contrast, the hybrid IoT-edge system strikes a balance between responsiveness and robustness, enabling immediate anomaly localization (<0.7 s) while dynamically adjusting detection sensitivity. This architectural alignment with operational tempo, not algorithmic complexity, proves to be the primary driver of observed improvements.

4.3. Sensor Contribution and Modality Ablation

To determine which sensing modalities exert the greatest influence on anomaly detection, an ablation study was conducted on weld-line production segments by selectively disabling specific sensor inputs. This approach isolates the marginal contribution of each modality and reveals sensitivity to temperature variation, vibration resonance, tool wear and visual surface change. Experiment results are summarized in Figure 3, where the full sensor fusion configuration is taken as the 100% baseline.

**Figure 3.** Modality Ablation Impact on Detection Accuracy.

Findings indicate that thermal and acoustic sensing jointly deliver the strongest detection capability, particularly in heat-affected fusion zones. Removing thermal data

alone reduced detection accuracy by 18.2%, confirming that real-time temperature tracking is essential to identifying sub-surface weld defects before they manifest visually. Meanwhile, acoustic features proved unexpectedly influential, enabling early detection of cutter instability and resonance drift several production cycles before surface degradation became visible.

Vision-based inspection contributed mainly to final-stage confirmation, capturing appearance-level inconsistencies but rarely preceding thermal trends. Vibration data provided intermediate value, primarily signaling emerging mechanical faults rather than product-quality anomalies directly. These layered effects support the architectural decision outlined in Section 3.2: multi-modal sensing is complementary rather than redundant, with each signal channel augmenting the temporal detectability window. The combined sensor set not only improves recognition accuracy but also brings failures to light earlier, producing cleaner, more interpretable alert patterns.

4.4. Operational Challenges and Practical Insights

Despite the system's significant performance gains, several operational issues emerged during deployment. First, sensor misalignment occurred after scheduled equipment maintenance, particularly affecting acoustic and thermal nodes. This resulted in measurement drift and occasional false alerts, suggesting the need for automated re-registration mechanisms to recalibrate sensors without halting production. Second, operators frequently dismissed low-severity alerts, perceiving them as background noise. While the system accurately detected early-stage anomalies, uniform alert severity led to reduced attention and delayed intervention. This indicates that graded alert scoring and adaptive visualization could improve human response.

A third challenge involved sudden changes in raw material batches. In these events, the adaptive detection model required 2-6 hours to reach a new equilibrium state due to shifts in heat dissipation, surface reflectivity, or stiffness. Although the system eventually stabilized, integrating online few-shot learning may shorten adaptation windows and minimize quality variation during transitions.

From these challenges, two insights are clear. Human-in-the-loop design remains indispensable, even highly sensitive systems underperform if operational staff disregard signals. Moreover, while weekly cloud-drift recalibration is effective, further advances such as meta-learning and rapid threshold self-adjustment could accelerate adaptation and reduce downtime.

4.5. Theoretical and Industrial Implications

The findings of this study demonstrate that effective quality monitoring does not require extensive labeled datasets. By leveraging short-term statistical fluctuations rather than historical defect annotations, the system adapts to process drift and variable operating conditions, challenging the traditional reliance on fully supervised learning for industrial quality control. This theoretical shift indicates that interpretability and adaptability can coexist with real-time responsiveness, offering a more scalable path toward autonomous manufacturing analytics.

On the industrial side, the implications are immediate and practical. Real-time detection enables intervention before defects propagate downstream, reducing scrap generation rather than merely identifying faults after they occur. The hybrid edge-computing structure lowers server load and operational cost, making deployment viable for small and medium-sized manufacturers, not only large factories. Additionally, the ablation findings provide data-driven guidance for sensor configuration, helping factories balance performance and budget. Ultimately, the system strengthens manufacturing resilience by improving visibility and response speed while reducing cognitive burden on operators, positioning human oversight as augmented rather than replaced.

5. Conclusion

This study developed and validated a real-time quality monitoring architecture for smart manufacturing, built upon multi-modal IoT sensing, edge inference, and adaptive statistical thresholding. Unlike conventional cloud-driven or static SPC-based solutions, the proposed system achieves sub-second fault identification without requiring large-scale historical defect labels. Through industrial deployment on Bosch automotive and CATL battery module production lines, the system demonstrated measurable and sustained gains, defect rate reductions of up to 31.4%, first-pass yield improvement exceeding 26%, and OEE increases beyond 7%. These results confirm not only the feasibility but also the economic value of integrating edge analytics with dynamic quality inference in high-velocity, high-variability manufacturing environments.

Academically, this work reframes quality monitoring as a real-time stochastic adaptation problem rather than a label-dependent classification task. The introduced moving-window variability index and self-adjusting threshold model offer a lightweight path toward adaptive quality assurance, particularly in settings where defect data are scarce or process conditions evolve continuously. The ablation study further clarifies the contribution hierarchy of sensing modalities, thermal and acoustic signals exhibiting the strongest defect predictiveness, which provides a theoretical reference for sensor allocation and cost-efficient instrumentation.

Practically, the implementation roadmap presented in the methodology section offers a replicable blueprint for factories transitioning toward Industry 4.0. The system's low bandwidth demand, rapid inferencing performance, and compatibility with existing MES/PLC infrastructure make it suitable not only for large-scale production but also for SMEs seeking scalable digitization.

Future work will focus on three concrete directions: (1) integrating few-shot or meta-learning mechanisms to accelerate model stabilization during material or tool transitions; (2) developing risk-tiered alerting interfaces to improve operator response and minimize alarm fatigue; and (3) extending the monitoring loop to include predictive maintenance, enabling gradual fault anticipation rather than pure anomaly response. Taken together, these pathways point toward an increasingly autonomous, interpretable, and economically deployable framework for intelligent quality assurance in modern manufacturing.

References

1. G. J. Cheng, L. T. Liu, X. J. Qiang, and Y. Liu, "Industry 4," 0 development and application of intelligent manufacturing. In 2016 international conference on information system and artificial intelligence (ISAI) (pp. 407-410). IEEE, June, 2016.
2. Z. Fan, R. X. Gao, Q. He, Y. Huang, T. Jiang, Z. Peng, and S. Zhong, "New sensing technologies for monitoring machinery, structures, and manufacturing processes," *Journal of Dynamics, Monitoring and Diagnostics*, pp. 69-88, 2023. doi: 10.37965/jdmd.2023.322
3. A. Ullah, M. Younas, and M. S. Saharudin, "Transforming manufacturing quality management with cognitive twins: a data-driven, predictive approach to real-time optimization of quality," *Journal of manufacturing and materials processing*, vol. 9, no. 3, 2025. doi: 10.3390/jmmp9030079
4. A. Rehman, and M. Islam, "IoT-enabled construction project monitoring systems a meta-analytical review of efficiency gains in large-scale infrastructure," *American Journal of Scholarly Research and Innovation*, vol. 2, no. 01, pp. 130-166, 2023.
5. M. A. H. B. O. O. B. Al Bashar, A. Taher, and D. I. L. A. R. A. Ashrafi, "Enhancing Efficiency of material handling equipment in industrial engineering sectors," *Iconic Res. Eng. Journals*, vol. 7, no. 11, pp. 595-604, 2024.
6. H. Zhang, X. Jia, C. Chen, S. Bachani, J. K. Goel, M. Tarun, and E. C. Anyanwu, "Deep learning-based real-time data quality assessment and anomaly detection for large-scale distributed data streams," *International Journal of Medical and All Body Health Research*, vol. 6, no. 1, pp. 01-11, 2025.
7. R. Tavva, "Time-Sensitive Networking in Advanced Manufacturing Environments: A Framework for Industry 4," 0 Implementation. *Journal of Computer Science and Technology Studies*, vol. 7, no. 6, pp. 672-678, 2025.
8. S. Rajarajan, R. Kalaivani, N. Kaliammal, S. T. Saravanan, P. Ilampiray, and C. Srinivasan, "IoT-Enabled Respiratory Pattern Monitoring in Critical Care: A Real-Time Recurrent Neural Network Approach," In 2024 10th International Conference on Communication and Signal Processing (ICCSP), April, 2024, pp. 508-513.

9. M. Orabi, K. P. Tran, S. Thomassey, and P. Egger, "Anomaly Detection for Catalyzing Operational Excellence in Complex Manufacturing Processes: A Survey and Perspective," In *Artificial Intelligence for Safety and Reliability Engineering: Methods, Applications, and Challenges*, 2024, pp. 49-78. doi: 10.1007/978-3-031-71495-5_4
10. I. Ahmed, and A. S. Raihan, "Machine learning techniques for sustainable industrial process control," In *Computational Intelligence Techniques for Sustainable Supply Chain Management*, 2024, pp. 141-176. doi: 10.1016/b978-0-443-18464-2.00014-5
11. T. Kim, and K. Behdinan, "Advances in machine learning and deep learning applications towards wafer map defect recognition and classification: a review," *Journal of Intelligent Manufacturing*, vol. 34, no. 8, pp. 3215-3247, 2023. doi: 10.1007/s10845-022-01994-1
12. S. Singh, and J. Kaur, "Recent Developments in Cloud-Based Technologies That Are Adaptive and pertinent," *Advancements in Cloud-Based Intelligent Informative Engineering*, pp. 95-114, 2025. doi: 10.4018/979-8-3373-0781-7.ch005
13. A. Kanduri, S. Shahhosseini, E. K. Naeini, H. Alikhani, P. Liljeberg, N. Dutt, and A. M. Rahmani, "Edge-centric optimization of multi-modal ml-driven ehealth applications," In *Embedded Machine Learning for Cyber-Physical, IoT, and Edge Computing: Use Cases and Emerging Challenges*, 2023, pp. 95-125. doi: 10.1007/978-3-031-40677-5_5
14. V. Veeramachaneni, "Edge computing: Architecture, applications, and future challenges in a decentralized era," *Recent trends in computer graphics and multimedia technology*, vol. 7, no. 1, pp. 8-23, 2025.
15. W. Yan, J. Wang, S. Lu, M. Zhou, and X. Peng, "A review of real-time fault diagnosis methods for industrial smart manufacturing," *Processes*, vol. 11, no. 2, p. 369, 2023. doi: 10.3390/pr11020369

Disclaimer/Publisher's Note: The views, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim any responsibility for any injury to individuals or damage to property arising from the ideas, methods, instructions, or products mentioned in the content.