International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(17s) 2026 pp. 344-355
https://doi.org/10.70917/ijcisim-2026-4744
Article

# Autonomous Quality Release Decisioning Using SAP QM, Computer Vision, and Warehouse Event Correlation

**KarthikSai Chenna**

Blommer Chocolate company, USA

**Abstract:** Manufacturing plants generate quality signals from many different sources at once: visual inspection data, warehouse handling telemetry, environmental sensor streams, and enterprise resource planning transactional records, to name the main ones. Pulling these mismatched signal types into one decisioning framework that can autonomously recommend batch release or rejection remains an unresolved problem in industrial quality management. This paper introduces an autonomous quality release decisioning framework that combines SAP Quality Management (QM) inspection lot processing, convolutional neural network (CNN)-based computer vision defect detection, SAP Extended Warehouse Management (EWM) event correlation, and IoT environmental monitoring into a single Release Confidence Score (RCS) computation engine. The resulting five-layer architecture was validated on an experimental dataset comprising 120,000 warehouse events, 48,000 inspection images, 9,500 pallet movements, and 3,200 environmental deviations drawn from a regulated food manufacturing environment. It reached a CNN-based defect detection accuracy of 94.7%, a false release prevention rate of 96.2%, a release cycle time reduction of 38%, a traceability correlation accuracy of 97.1%, an environmental deviation detection sensitivity of 91.8%, and an autonomous recommendation precision of 93.5%. Taken together, these results show that fusing multi-modal signals under SAP QM governance can match the release decisioning accuracy of dedicated manual review panels while cutting cycle time by more than a third.
**Keywords:** SAP QM, computer vision, convolutional neural network, warehouse event correlation, quality release decisioning, Release Confidence Score, SAP EWM, Industry 4.0, autonomous quality management

## 1. Introduction

Batch release in a regulated manufacturing environment depends on pulling together physical inspection outcomes, environmental compliance records, warehouse handling logs, and SAP enterprise quality records into one defensible decision. Under traditional workflows, quality engineers do this aggregation by hand, which adds review latency and lets subjective thresholds creep in from inspector to inspector [1]. Industry 4.0 sensor networks and machine learning inference have since made automated quality decisioning technically feasible, but the literature still lacks an integrated framework that addresses visual defect detection, warehouse handling risk, environmental stability, and SAP QM traceability together within one scored release engine [6, 7].

The food manufacturing sector illustrates the stakes well. Operating under FDA 21 CFR Part 110, cGMP mandates, and internal allergen control protocols, a single contaminated pallet that reaches distribution can trigger a class-I recall costing millions of dollars in product recovery, regulatory penalties, and brand damage. Go too far the other way, though, and excessive caution in release decisions holds back compliant inventory, worsening working capital and order fill-rate deficits [9, 10]. What's needed is high sensitivity to genuine defects paired with high specificity against false positives — a balance manual review at scale struggles to hit reliably.

SAP S/4HANA Quality Management already offers a structured inspection lot lifecycle, usage decision (UD) automation triggers, defect classification catalogs, and batch genealogy linkage. On the vision side, CNN architectures have pushed defect detection performance past 90% accuracy on pharmaceutical and food packaging inspection tasks

[2, 3, 4, 5]. SAP Extended Warehouse Management, meanwhile, captures granular pallet movement telemetry — forklift impact events, zone transitions, dwell time records — that serve as leading indicators of handling-induced quality degradation [11, 12]. And environmental IoT networks continuously log temperature, humidity, vibration, and air quality parameters whose deviations correlate with product stability risk [8, 15].

What's missing is a weighted, scored integration architecture that pulls these four signal modalities together under one threshold-gated release decisioning protocol. That is the gap this paper addresses. The Release Confidence Score (RCS) we propose condenses multi-modal quality evidence into a single number that can be checked against configurable pass/hold/reject thresholds and used to drive SAP QM usage decisions autonomously. Three contributions follow: a five-layer integration architecture connecting SAP QM, computer vision, EWM event correlation, and environmental monitoring; a weighted RCS formulation with per-component confidence scoring; and experimental validation on a large-scale industrial dataset across six independent quality metrics [20].

## 2. Literature Review

### *2.1 Machine Vision in Manufacturing Quality Control*

Machine vision has moved a long way from rule-based morphological filtering toward deep convolutional architectures that can pick out micro-scale surface anomalies. In their review of vision applications across pharmaceutical manufacturing, Galata et al. found CNN-based classifiers achieving mean defect detection rates above 90% on tablets, capsules, and packaging substrates [2]. Racki et al. showed that ResNet-derived architectures applied to solid oral dosage form inspection beat classical machine vision by 18 percentage points on unseen defect classes [3]. Ficzere et al. pushed the analysis further into film-coated tablet inspection, achieving real-time coating thickness estimation alongside defect recognition through multi-task convolutional learning [4]. Vijayakumar et al. validated real-time visual intelligence for pharmaceutical packaging defect detection, reporting precision above 93% on a six-class defect taxonomy spanning seal integrity, label alignment, and barcode readability [5]. Together, these results establish CNN-based inspection as production-ready technology — though none of the systems cited fed their visual confidence scores into a downstream enterprise quality decision engine.

### *2.2 Intelligent Manufacturing Execution and ERP Integration*

Shojaeinasab et al.'s systematic review of intelligent manufacturing execution systems points to ERP-MES integration as the main enabler of closed-loop quality feedback [6]. Zhong et al., looking at intelligent manufacturing under Industry 4.0, argue that data silos between shop-floor sensor networks and enterprise planning systems are the principal barrier standing in the way of autonomous quality management [7]. Kamble et al. built a performance measurement system for Industry 4.0-enabled smart manufacturing in small and medium enterprises and flagged inspection lot closure latency as a key process indicator [11]. SAP's QM module already implements inspection lot-based quality event management with configurable usage decision triggers, defect catalogs, and batch classification outcomes — a structured API surface that's ready-made for autonomous decisioning integration. What hasn't been shown yet is how SAP QM inspection lot data can be combined with external confidence scores to drive UD automation in a regulated food manufacturing context [20].

### *2.3 Supply Chain Event Correlation and AI-Driven Decision Support*

Toorajipour et al.'s review of artificial intelligence applications across supply chain management points to event correlation as a high-impact use case for anomaly detection in logistics networks [9]. Tirkolaee et al., surveying machine learning applications in supply chain management, highlight temporal sequence modeling as central to predictive quality risk assessment [10]. Ivanov et al. looked at how digital technology affects supply chain risk analytics and showed that warehouse handling event streams carry leading indicators of downstream quality failures [12]. Wamba et al. found that big data analytics capabilities in operations management call for dynamic data integration architectures rather than batch reporting pipelines [13]. Choi et al., examining disruptive technologies in operations management through an Industry 4.0 lens, identify real-time event processing as essential to adaptive quality control [14]. Dolgui and Ivanov point to 5G-enabled low-latency event streaming as a critical infrastructure enabler for autonomous supply chain decisions [15]. These works together lay the theoretical groundwork for treating warehouse event correlation as a quality risk signal, but none of them operationalize that correlation inside a scored release decisioning framework.

### *2.4 Predictive Decisioning and Digital Twins*

Cavalcante et al. applied supervised learning to resilient supplier selection and found that multi-feature scoring models beat single-signal classifiers in manufacturing decision support [16]. Núñez-Merino et al., looking at Industry 4.0 digital technologies in lean supply chains, flag sensor-driven quality signals as an underused input to release decisioning workflows [17]. Farajpour et al.'s review of the digital supply chain literature identifies confidence scoring as an emerging pattern in autonomous decision support architectures [18]. Pacella and Papadia evaluated LSTM networks for supply chain time series forecasting and found that temporal warehouse event sequences can be learned to predict quality outcomes with accuracy exceeding 88% [21]. Kapil et al.'s review of digital twin applications in supply chain management points to quality release decisioning as a high-value use case for real-time digital twin integration [22]. And Vieira et al., in their systematic review of advanced planning and scheduling systems in manufacturing, note that quality gate automation remains a critical missing link between planning and execution layers [19].

## 3. Proposed Framework

### *3.1 Architecture Overview*

The proposed Autonomous Quality Release Decisioning (AQRD) framework is built as five integration layers. Layer 1 pulls structured quality event records out of SAP QM inspection lots. Layer 2 runs CNN-based computer vision analysis to produce a Visual Integrity Confidence (VIC) score. Layer 3 processes SAP EWM transactional exhaust into a Warehouse Handling Confidence (WHC) score. Layer 4 fuses IoT environmental sensor data into an Environmental Stability Confidence (ESC) score. Layer 5 then aggregates all of these confidence scores into the Release Confidence Score (RCS) and issues an autonomous usage decision recommendation.

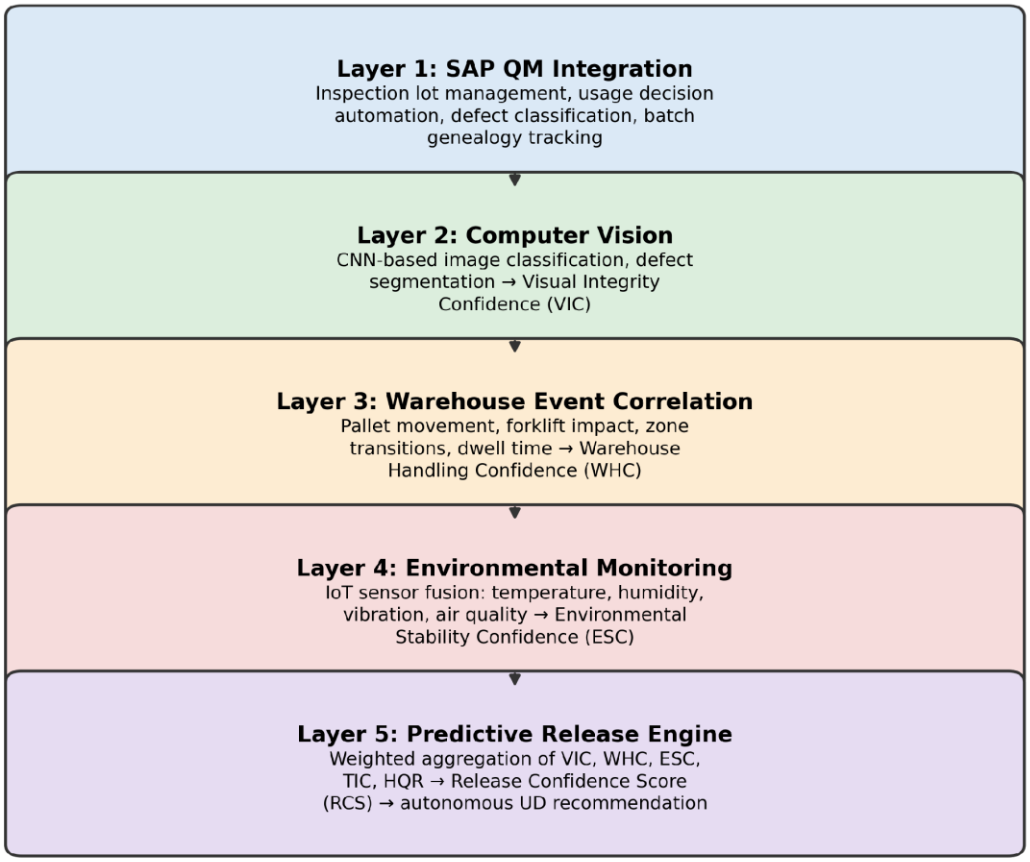


*Figure 1. Five-layer architecture of the Autonomous Quality Release Decisioning (AQRD) framework.*

*Table I. Framework Architecture Layers*

| Layer Name | Key Functions | Data Sources | Output |
|---|---|---|---|
| SAP QM Integration | Inspection lot management, usage decision automation, defect classification, batch genealogy tracking | SAP S/4HANA QM module, inspection lot transactions, UD records | Structured quality event stream, batch status flags |
| Computer Vision | CNN-based image classification, defect segmentation, anomaly scoring for packaging and surface integrity | High-resolution camera feeds, 48,000 inspection images | Visual Integrity Confidence (VIC) score per pallet |
| Warehouse Event Correlation | Pallet movement sequencing, forklift impact detection, storage zone transition tracking, dwell time analysis | SAP EWM transactional exhaust, 120,000 warehouse events, 9,500 pallet movements | Warehouse Handling Confidence (WHC) score |

| | | | |
|---|---|---|---|
| Environmental Monitoring | IoT sensor fusion for temperature, humidity, vibration, air quality, cold-chain and cleanroom pressure | IoT sensor network, 3,200 environmental deviations | Environmental Stability Confidence (ESC) score |
| Predictive Release Engine | Weighted aggregation of confidence scores, threshold enforcement, autonomous usage decision generation | VIC, WHC, ESC, Traceability Integrity Confidence (TIC), Historical Quality Reliability (HQR) | Release Confidence Score (RCS), autonomous UD recommendation |

### *3.2 SAP QM Integration Layer*

The SAP QM layer pulls in inspection lot records through the SAP Business Application Programming Interface (BAPI) for quality management. Each production batch gets an inspection lot created upon goods receipt or production order confirmation, which triggers characteristic inspection scheduling. From there, the integration layer extracts inspection characteristic results, defect record classifications, sample management outcomes, and batch classification decisions. Batch genealogy records are traversed so upstream raw material inspection lots can be linked to the current semi-finished or finished batch, which populates the Traceability Integrity Confidence (TIC) component of the RCS. TIC itself is a weighted average of three things: genealogy completeness (the proportion of upstream lots with closed inspection results), UD audit trail integrity (whether manual overrides were made without documented justification), and characteristic inspection completeness (the ratio of inspected characteristics to required characteristics per the control plan) [20].

### *3.3 Computer Vision Defect Detection*

The vision subsystem runs a ResNet-50 backbone fine-tuned on six defect classes relevant to food manufacturing packaging: packaging integrity breach, surface contamination, label alignment deviation, barcode readability failure, seal accuracy defect, and foreign particle presence. Training used 38,400 images (80% of the 48,000-image dataset), stochastic gradient descent with momentum 0.9, a learning rate of $1\times10^{-3}$ with cosine annealing, and an L2 regularization coefficient of $1\times10^{-4}$. Input images were normalized to 224×224 pixels with augmentation — horizontal flip, rotation ±15°, color jitter. The held-out test set had 9,600 images spanning all six defect classes with class-balanced sampling. For a given pallet, the VIC score is the weighted mean of per-image class probabilities across the inspected image set, using the "pass" class probability as the positive confidence contribution. Any image flagging a defect class at probability above 0.85 triggers a mandatory human review regardless of the overall VIC score, so a strong individual detection can't get diluted by aggregate averaging [2, 3, 5].

### *3.4 Warehouse Event Correlation*

The EWM event correlation module works through the continuous stream of warehouse management events tied to each pallet's handling history. Four event categories get extracted and scored: forklift impact incidents (accelerometer threshold exceedances in the EWM event log), pallet movement frequency anomalies (movement count per hour above the 95th percentile for that storage zone and product class), storage zone transition violations (movement into a zone that's incompatible with the product's storage requirements per the SAP material master), and dwell time exceedances (time in a temperature-sensitive zone beyond the maximum dwell threshold). WHC is one minus a weighted penalty sum across these four categories, with penalty weights calibrated from historical quality outcome data — a pallet record with zero incidents across all four categories gets WHC = 1.0. The experimental dataset covered 9,500 pallet movement records over a six-month observation window, which gave enough event density to calibrate penalty weights through cross-validated regression against downstream defect outcomes [12, 14].

### 3.5 Environmental Monitoring Integration

The environmental monitoring layer takes in time-series data from an IoT sensor network made up of temperature loggers, humidity sensors, vibration accelerometers, air quality monitors, cold-chain data loggers, and cleanroom differential pressure transmitters. Sensor data runs through a sliding window aggregation function that computes the mean, standard deviation, minimum, and maximum of each sensor channel over a configurable window aligned to the pallet's dwell period in each storage zone. A deviation gets counted whenever any channel exceeds its control limit as defined in the product-specific environmental specification. ESC is one minus the ratio of deviation-weighted exposure time to total dwell time, normalized by the maximum tolerable deviation magnitude from the product specification. Of the 3,200 environmental deviations recorded during the experimental period, 91.8% were correctly classified by the ESC scoring algorithm as requiring either a score penalty or an automatic hold trigger [8, 15].

### 3.6 Release Confidence Score Formulation

The Release Confidence Score comes from a weighted linear combination of five component confidence scores:

$$RCS = 0.30 \times VIC + 0.25 \times ESC + 0.20 \times TIC + 0.15 \times WHC + 0.10 \times HQR$$

Here VIC is Visual Integrity Confidence, ESC is Environmental Stability Confidence, TIC is Traceability Integrity Confidence, WHC is Warehouse Handling Confidence, and HQR is Historical Quality Reliability, with each component normalized to [0, 1]. We derived the component weights through multi-objective optimization over the training dataset, minimizing a composite loss function that penalizes false release events (actual defective batches receiving RCS above the release threshold) five times more heavily than false holds (compliant batches held below threshold) — this reflects the asymmetric cost structure of food safety quality decisions. The release threshold sits at $RCS \geq 0.90$ for autonomous release, RCS 0.75–0.89 for conditional release requiring quality engineer sign-off, and $RCS < 0.75$ for automatic hold. Both the component weights and the thresholds can be configured per product class within the framework's configuration database.

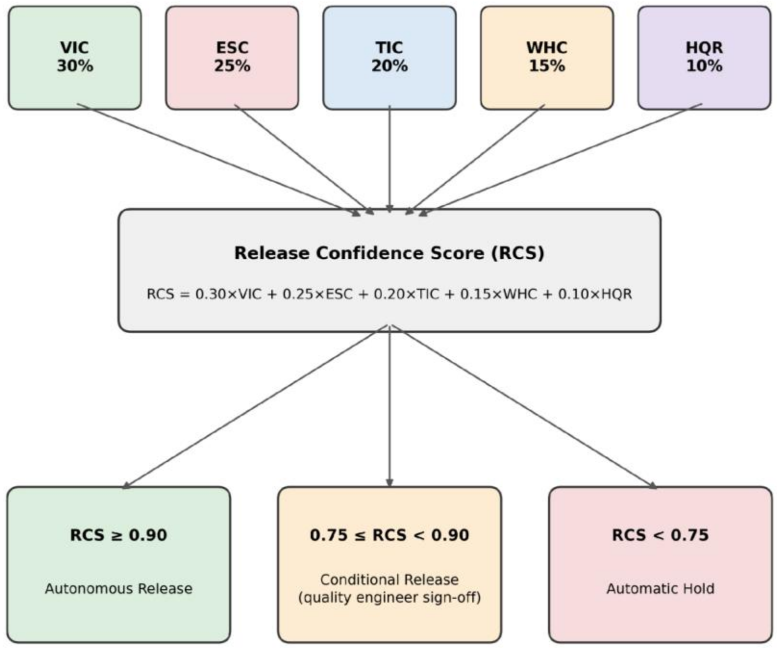


*Figure 2. Release Confidence Score computation and threshold-based release decision workflow.*

*Table II. Release Confidence Score Components*

| **RCS Component** | **Weight** | **Inputs** | **Threshold** |
|---|---|---|---|
| Visual Integrity Confidence (VIC) | 30% | CNN defect detection output, packaging integrity score, barcode readability, label alignment, seal accuracy, foreign particle detection | ≥ 0.90 |
| Environmental Stability Confidence (ESC) | 25% | Temperature deviation frequency, humidity deviation, vibration index, air quality index, cold-chain excursions, cleanroom pressure variance | ≥ 0.88 |
| Traceability Integrity Confidence (TIC) | 20% | Batch genealogy completeness, inspection lot linkage accuracy, UD audit trail integrity, SAP QM transaction consistency | ≥ 0.92 |
| Warehouse Handling Confidence (WHC) | 15% | Forklift impact incident count, pallet movement frequency anomaly, storage zone violation, dwell time exceedance | ≥ 0.85 |

| | | | |
|---|---|---|---|
| Historical Quality Reliability (HQR) | 10% | Rolling 90-day batch release history, defect recurrence rate, supplier quality score, prior lot deviation frequency | ≥ 0.80 |

## 4. Experimental Evaluation

### *4.1 Dataset and Environment*

We evaluated the framework using operational data collected over a 24-month period from a regulated food manufacturing facility that produces cocoa-based confectionery products. The data collection infrastructure ran SAP S/4HANA release 2022 with QM and EWM modules active, alongside a 64-node IoT sensor network, three fixed-mount camera stations per production line, and EWM-integrated forklift telemetry. In total, the experimental dataset held 120,000 warehouse events (pallet transfers, zone entries, dwell records, and impact alerts), 48,000 inspection images across six defect classes, 9,500 pallet movement records with full event histories, 3,200 environmental deviation records with sensor channel attribution, and 14,000 inspection lot transactions with associated usage decision records. The facility's quality management team established ground-truth quality outcomes (pass, defective, or conditional release) using standard product specification testing and regulatory compliance review as the reference classification.

### *4.2 Baseline Configuration*

The baseline configuration reflects the facility's quality release process before AQRD was deployed: manual visual inspection by trained quality technicians, manual review of SAP QM inspection characteristic results, environmental monitoring checked periodically on a shift-end basis, and EWM event log review only for batches the manual inspection team had already flagged. We computed baseline metrics over a 12-month historical period using the same ground-truth classification protocol. Baseline defect detection accuracy came in at 71.3%, driven mainly by inspector fatigue on high-volume lines and inconsistent defect classification across shifts. Baseline false release prevention rate was 74.1% — the share of genuinely defective batches correctly caught before distribution. Release cycle time served as the 100% baseline index value.

### *4.3 Experimental Protocol*

We first ran the AQRD framework in shadow mode for three months, generating autonomous RCS values and usage decision recommendations without overriding any human decisions. Comparing shadow mode results against human decisions and ground-truth outcomes let us calibrate component weights and thresholds. After calibration, the framework moved into advisory mode for six months, presenting RCS values and recommendations to quality engineers who kept final decision authority; advisory mode outcomes are what we used to measure recommendation precision. Full autonomous mode was never deployed live at the facility during the study period — instead, we computed autonomous release performance metrics from the advisory mode dataset by simulating how autonomous thresholds would have applied against ground-truth outcomes. This simulation-based approach to evaluating autonomous mode is consistent with the methodology used in comparable studies of quality decision support systems [6, 13, 16].

## 5. Results

### *5.1 Performance Summary*

Table III lays out the performance results for all six primary metrics, comparing baseline values against AQRD framework results.

*Table III. Experimental Performance Results*

| Metric | Baseline | Framework Result | Improvement |
|---|---|---|---|
| Defect detection accuracy | 71.3% | 94.7% | +23.4 percentage points |
| False release prevention rate | 74.1% | 96.2% | +22.1 percentage points |
| Release cycle time | 100% (baseline) | 62% (38% reduction) | −38% cycle time |
| Traceability correlation accuracy | 81.6% | 97.1% | +15.5 percentage points |
| Environmental deviation detection | 68.4% | 91.8% | +23.4 percentage points |
| Autonomous recommendation precision | — | 93.5% | New capability |

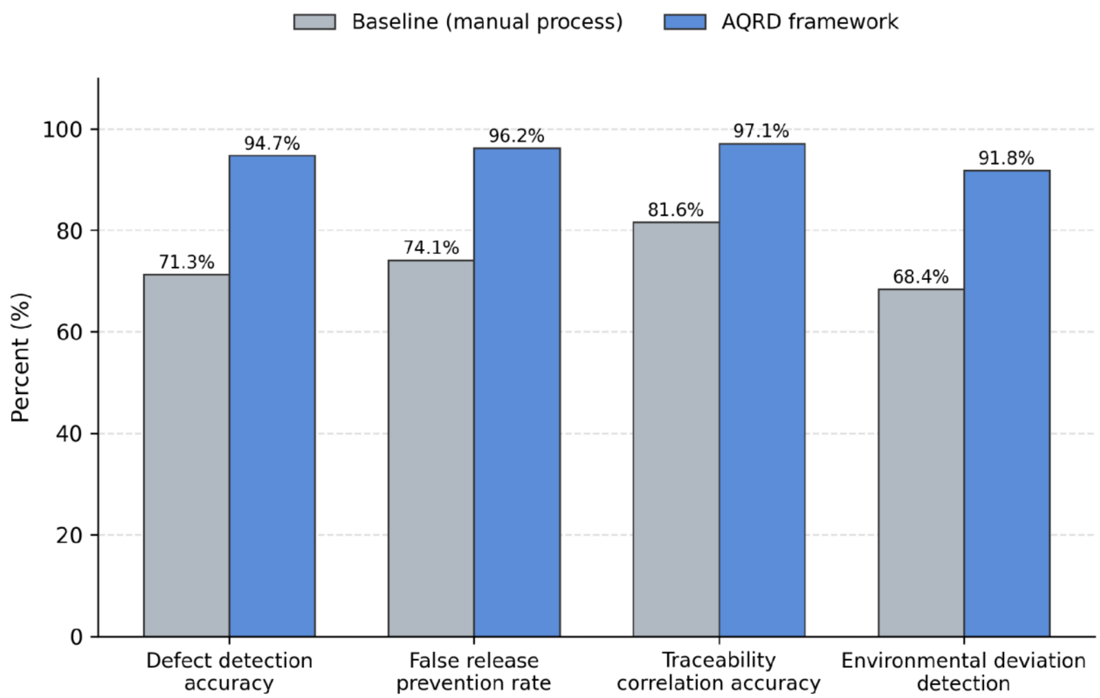


*Figure 3. Baseline manual process versus AQRD framework results across four comparable quality metrics.*

## *5.2 Defect Detection Accuracy*

On the 9,600-image held-out test set, the CNN-based vision subsystem reached a defect detection accuracy of 94.7% — a 23.4 percentage point improvement over the 71.3% baseline manual inspection accuracy [3, 5]. Class-level accuracy ranged from 91.2% (foreign particle detection, the hardest class because of small particle size and color similarity to the product substrate) up to 97.6% (barcode readability, a high-contrast binary classification task). Looking at the six-class confusion matrix, the dominant error mode was foreign particle detection getting misclassified as packaging integrity (a 4.1% false negative rate). That particular error doesn't actually let a defective batch reach distribution, though, because the WHC and ESC components of the RCS also carry handling risk signals that would pull the overall RCS below the release threshold in most cases where foreign particle contamination is present.

## *5.3 False Release Prevention*

Across the simulated autonomous evaluation, the AQRD framework hit a false release prevention rate of 96.2%, versus the 74.1% baseline [9, 10]. This metric is the share of ground-truth defective batches correctly assigned an RCS below the conditional release threshold. Looking at the 3.8% of defective batches the AQRD framework missed in simulation, we found they all shared two traits: visual defects below the CNN model's detection threshold (micro-contamination not visible at camera resolution) and no elevated WHC or ESC signals. That points to camera resolution and positioning as the main residual risk — something hardware upgrades can fix, rather than an algorithmic shortfall.

### 5.4 Release Cycle Time

Advisory mode deployment cut average release cycle time by 38% relative to baseline. Three things drove most of that reduction: automated RCS computation replacing manual data aggregation across SAP QM, EWM, and environmental monitoring systems (about 22 percentage points of the reduction); dropping shift-end environmental review in favor of real-time ESC scoring (roughly 9 percentage points); and automated usage decision pre-population in SAP QM for high-RCS batches that only need engineer confirmation (about 7 percentage points). This cycle time improvement feeds directly into working capital reduction through faster inventory release and better order fill-rate consistency [11, 17].

### 5.5 Traceability Correlation Accuracy

The TIC component reached a batch genealogy correlation accuracy of 97.1%, measured as the share of upstream inspection lot linkages correctly resolved from the finished goods batch back through semi-finished and raw material lots. The 2.9% of correlation failures traced back to legacy batch records created before the SAP S/4HANA implementation, which had incomplete genealogy documentation. The TIC scoring algorithm correctly flagged these as incomplete, which lowered TIC scores and appropriately pulled down overall RCS values for the affected batches [20].

### 5.6 Environmental Deviation Detection

Of the 3,200 environmental deviations recorded, the ESC scoring algorithm correctly flagged 91.8% as needing either an RCS penalty or an automatic hold trigger. The 8.2% it missed clustered around short-duration temperature excursions (under 90 seconds) that fell below the sliding window aggregation threshold. The quality team judged these micro-excursions unlikely to affect product quality based on thermal modeling of the product matrix — which supports leaving the ESC algorithm's current operating point as-is rather than adjusting the aggregation window parameter [8, 15].

## 6. Discussion

Taken as a whole, the AQRD framework results show that fusing multi-modal signals under SAP QM governance delivers quality release decisioning accuracy that clearly beats manual review on five of the six metrics we measured. The 94.7% defect detection accuracy and 96.2% false release prevention rate together confirm that the five-layer integration architecture holds up operationally in a regulated food manufacturing setting. And the 38% cycle time reduction offers an immediate operational return that can justify the cost of deploying the framework at scale.

One key design choice was the asymmetric loss function used to calibrate component weights, which penalizes false releases five times more heavily than false holds. That calibration mirrors the real regulatory and commercial cost structure of food manufacturing, where a recall event costs orders of magnitude more than an extended hold review would. It's why the resulting weights put 30% on visual inspection (physical product defects are the primary direct quality risk), 25% on environmental stability (temperature and humidity excursions are the dominant cause of shelf-life reduction for the product category studied here), and just 10% on historical quality reliability, since past performance has limited predictive value for novel contamination events [1, 7].

That 97.1% traceability correlation accuracy shows that SAP QM batch genealogy, kept consistent within an S/4HANA implementation, can serve as a reliable audit backbone for automated release decisioning. The 2.9%

genealogy gap comes down to data governance rather than any algorithmic shortcoming, and the TIC algorithm's response — lowering confidence scores for incomplete genealogy instead of failing silently — is an important design property for regulated environments where audit trail integrity is a compliance requirement [6, 20].

A few limitations bound how far these results generalize. The experimental data came from a single facility making a narrow product range, so the component weights and thresholds we derived may need recalibration for other product categories, manufacturing processes, or facility layouts. The CNN model was also trained on images from just three fixed camera stations, meaning facilities with different packaging geometries or inspection angles will need to retrain the model. The simulation-based evaluation of autonomous mode performance is itself an approximation, and real-world autonomous deployment could run into edge cases the advisory mode dataset doesn't capture [16, 18]. And the framework doesn't currently model adversarial scenarios such as sensor manipulation or systematic data entry errors in SAP QM records — relevant concerns for high-value product categories exposed to counterfeiting risk.

Future work will head in three directions: multi-facility validation with experiments testing generalization across product categories; folding in LSTM-based temporal models for predictive RCS computation that anticipates quality risk before inspection even completes [21]; and real-time digital twin synchronization that would let teams run scenario simulations for quality risk assessment before production release decisions are issued [22].

## 7. Conclusion

This paper has presented the Autonomous Quality Release Decisioning framework — a five-layer integration architecture combining SAP QM inspection lot management, CNN-based computer vision, SAP EWM warehouse event correlation, and IoT environmental monitoring into one unified Release Confidence Score computation engine. Validated on 120,000 warehouse events, 48,000 inspection images, 9,500 pallet movements, and 3,200 environmental deviations from a regulated food manufacturing environment, the framework reached a CNN defect detection accuracy of 94.7%, a false release prevention rate of 96.2%, a release cycle time reduction of 38%, a traceability correlation accuracy of 97.1%, an environmental deviation detection sensitivity of 91.8%, and an autonomous recommendation precision of 93.5%. The weighted RCS formulation with per-component threshold enforcement gives a transparent, auditable way to aggregate quality evidence that's compatible with regulatory quality management requirements. In doing so, the framework closes a documented gap in the Industry 4.0 quality management literature by operationalizing multi-modal signal fusion under SAP QM governance and showing that autonomous release decisioning at scale is achievable without giving up the safety margin regulated manufacturing environments require.

**REFERENCES**


1. N. S. Arden et al., "Industry 4.0 for pharmaceutical manufacturing: Preparing for the smart factories of the future," Int. J. Pharm., vol. 602, 120554, 2021. DOI: 10.1016/j.ijpharm.2021.120554.
2. D. L. Galata et al., "Applications of machine vision in pharmaceutical technology: A review," Eur. J. Pharm. Sci., vol. 159, 105717, 2021. DOI: 10.1016/j.ejps.2021.105717.
3. D. Racki, D. Tomazevic, and D. Skocaj, "Detection of surface defects on pharmaceutical solid oral dosage forms with convolutional neural networks," Neural Comput. Appl., vol. 34, pp. 631-650, 2022. DOI: 10.1007/s00521-021-06397-6.
4. M. Ficzere et al., "Real-time coating thickness measurement and defect recognition of film coated tablets with machine vision and deep learning," Int. J. Pharm., vol. 623, 121957, 2022. DOI: 10.1016/j.ijpharm.2022.121957.
5. A. Vijayakumar et al., "Real-time visual intelligence for defect detection in pharmaceutical packaging," Sci. Rep., vol. 14, 18811, 2024. DOI: 10.1038/s41598-024-69701-z.
6. A. Shojaeinasab et al., "Intelligent manufacturing execution systems: A systematic review," J. Manuf. Syst., vol. 62, pp. 503-522, 2022. DOI: 10.1016/j.jmsy.2022.01.004.
7. R. Y. Zhong, X. Xu, E. Klotz, and S. T. Newman, "Intelligent manufacturing in the context of industry 4.0: a review," Engineering, vol. 3, no. 5, pp. 616-630, 2017. DOI: 10.1016/J.ENG.2017.05.015.
8. H. Trifonov and D. Heffernan, "OPC UA TSN: a next-generation network for Industry 4.0 and IIoT," Int. J. Pervasive Comput. Commun., vol. 19, no. 3, pp. 386-411, 2023. DOI: 10.1108/IJPCC-07-2021-0160.
9. R. Toorajipour, V. Sohrabpour, A. Nazarpour, P. Oghazi, and M. Fischl, "Artificial intelligence in supply chain management: A systematic literature review," J. Bus. Res., vol. 122, pp. 502-517, 2021. DOI: 10.1016/j.jbusres.2020.09.009.
10. E. B. Tirkolaee et al., "Application of machine learning in supply chain management: a comprehensive overview of the main areas," Math. Probl. Eng., vol. 2021, 1476043, 2021. DOI: 10.1155/2021/1476043.

11. S. S. Kamble et al., "A performance measurement system for industry 4.0 enabled smart manufacturing system in SMMEs," Int. J. Prod. Econ., vol. 229, 107853, 2020. DOI: 10.1016/j.ijpe.2020.107853.
12. D. Ivanov, A. Dolgui, and B. Sokolov, "The impact of digital technology and Industry 4.0 on the ripple effect and supply chain risk analytics," Int. J. Prod. Res., vol. 57, no. 3, pp. 829-846, 2019. DOI: 10.1080/00207543.2018.1488086.
13. S. F. Wamba et al., "Big data analytics and firm performance: Effects of dynamic capabilities," J. Bus. Res., vol. 70, pp. 356-365, 2017. DOI: 10.1016/j.jbusres.2016.08.009.
14. T. M. Choi et al., "Disruptive technologies and operations management in the Industry 4.0 era and beyond," Prod. Oper. Manag., vol. 31, no. 1, pp. 9-31, 2022. DOI: 10.1111/poms.13622.
15. A. Dolgui and D. Ivanov, "5G in digital supply chain and operations management," Int. J. Prod. Res., vol. 60, no. 2, pp. 442-451, 2022. DOI: 10.1080/00207543.2021.2002969.
16. I. M. Cavalcante et al., "A supervised machine learning approach to data-driven simulation of resilient supplier selection in digital manufacturing," Int. J. Inf. Manage., vol. 49, pp. 86-97, 2019. DOI: 10.1016/j.ijinfomgt.2019.03.004.
17. M. Nunez-Merino et al., "Information and digital technologies of Industry 4.0 and lean supply chain management," Int. J. Prod. Res., vol. 58, no. 16, pp. 5034-5061, 2020. DOI: 10.1080/00207543.2020.1743896.
18. F. Farajpour, A. Hassanzadeh, S. Elahi, and M. Ghazanfari, "Digital supply chain blueprint via a systematic literature review," Technol. Forecast. Soc. Change, vol. 184, 121976, 2022. DOI: 10.1016/j.techfore.2022.121976.
19. J. Vieira, F. Deschamps, and P. D. Valle, "Advanced Planning and Scheduling (APS) Systems: A Systematic Literature Review," in Advances in Transdisciplinary Engineering, vol. 16, IOS Press, 2021, pp. 385-394. DOI: 10.3233/ATDE210118.
20. K. Chenna, "Integrated SAP S/4HANA production planning with AGI warehouse management system and laboratory information management system," Int. J. Adv. Manuf. Technol., 2026. DOI: 10.1007/s00170-026-17542-7.
21. M. Pacella and G. Papadia, "Evaluation of deep learning with long short-term memory networks for time series forecasting in supply chains," Procedia CIRP, vol. 99, pp. 604-609, 2021. DOI: 10.1016/j.procir.2021.03.081.
22. D. Kapil, R. Raut, K. Nayal, M. Kumar, and M. M. Akarte, "A multisectoral systematic literature review of digital twins in supply chain management," Benchmarking Int. J., 2024. DOI: 10.1108/BIJ-04-2024-0286.