

Article - e005

**MULTI-AGENT SYSTEMS FOR ANOMALY DETECTION AND
QUALITY CONTROL IN INDUSTRIAL ENVIRONMENTS: A
COMPREHENSIVE REVIEW**

Vedansh Mishra¹  

¹Department of Advance Computing Specialization, SAGE University, Indore

Received: 23/05/2026

Revision Received: 10/06/2026

Accepted: 28/06/2026

ABSTRACT

The rapid adoption of Industry 4.0 technologies has transformed traditional manufacturing systems into highly interconnected and data-driven environments. Modern production facilities generate large volumes of sensor, machine, and operational data, creating opportunities for intelligent quality control and predictive decision-making.[8],[10],[39] However, conventional quality assurance techniques often struggle to process such heterogeneous data streams in real time. Consequently, anomaly detection has emerged as a critical research area for identifying defects, process deviations, and equipment failures before they impact production outcomes. Simultaneously, Multi-Agent Systems (MAS) have gained significant attention due to their decentralized architecture, scalability, adaptability, and fault-tolerant characteristics. [3],[5], [38]

This literature review examines the evolution of anomaly detection techniques and their integration with MAS for industrial quality control applications. The review covers statistical methods, machine learning algorithms, deep learning approaches, sensor fusion strategies, edge-cloud computing architectures, and contemporary agent frameworks. Furthermore, the paper analyzes current challenges related to real-time performance, interoperability, scalability, and deployment in industrial environments. Existing research demonstrates that while anomaly detection models have achieved remarkable improvements in accuracy, their integration into distributed industrial systems remains limited. Multi-agent architectures offer a promising solution by enabling autonomous monitoring, localized decision-making, and collaborative intelligence across manufacturing environments. Finally, research gaps and future directions are identified to support the development of intelligent, resilient, and scalable industrial quality-control systems.

KEYWORDS: Multi-Agent Systems, Industrial Quality Control, Anomaly Detection, Industry 4.0, Artificial Intelligence, Edge Computing, Deep Learning

1. INTRODUCTION

Manufacturing industries have undergone significant technological transformation over the last two decades. The integration of advanced sensing technologies, Industrial Internet of Things (IIoT) platforms, cloud computing infrastructures, and artificial intelligence has enabled unprecedented levels of automation and operational intelligence.[8],[10],[39] Modern production facilities generate enormous volumes of data through interconnected machines, sensors, cameras, programmable logic controllers, and manufacturing execution systems. These developments have created opportunities for improving productivity, reducing operational costs, and enhancing product quality. However, they have simultaneously increased the complexity of quality assurance processes and defect management systems.

Quality control remains one of the most important aspects of manufacturing operations because even minor defects can result in substantial economic losses. Poor-quality products lead to material wastage, production rework, customer dissatisfaction, warranty claims, and reputational damage. The concept of Cost of Poor Quality (COPQ) highlights the financial consequences associated with manufacturing defects and inefficient quality-management processes. Research has consistently shown that quality-related losses can represent a significant proportion of manufacturing expenditures [1],[2], emphasizing the importance of robust quality-control mechanisms.

Traditionally, industrial quality control relied heavily on manual inspection and statistical process monitoring techniques. Human inspectors visually examined products for defects, while statistical quality-control methods monitored process parameters using predefined thresholds and control charts. Although these techniques have been widely adopted across manufacturing sectors, they suffer from several limitations. Manual inspection is labor-intensive, inconsistent, and prone to fatigue-induced errors. Statistical methods often struggle to detect complex nonlinear relationships among process variables and may fail to identify subtle defects hidden within large volumes of heterogeneous data. As manufacturing systems become increasingly automated and data-rich, these traditional approaches are no longer sufficient to meet industrial requirements.[11],[13]

To address these limitations, anomaly detection has emerged as a crucial research domain within industrial quality control. Anomaly detection refers to the identification of observations or patterns that significantly deviate from normal operational behavior. In manufacturing environments, anomalies may include defective products, machine malfunctions, process instabilities, sensor failures, abnormal vibrations, thermal irregularities, or unexpected operational conditions. Early identification of such anomalies enables organizations to implement corrective actions before defects propagate through the production chain. [13], [14], [15]

In parallel, Multi-Agent Systems (MAS) have evolved as a promising architectural paradigm for managing distributed industrial environments. A Multi-Agent System consists of autonomous computational entities capable of sensing, reasoning, communication, and cooperative decision-making. Unlike centralized systems that rely on a single control mechanism, MAS distribute intelligence across multiple agents, enabling greater scalability, adaptability, and fault tolerance. These characteristics align closely with the requirements of

modern smart manufacturing environments[3], [5], [7] where large numbers of interconnected devices operate simultaneously across geographically distributed production facilities.

The convergence of anomaly detection technologies and Multi-Agent Systems presents a unique opportunity to create intelligent quality-control frameworks capable of autonomous monitoring, distributed decision-making, and adaptive process optimization. This review examines the current state of research in these areas and identifies future directions for integrating advanced anomaly-detection techniques with agent-based industrial architectures.

Recent advances in foundation models, vision transformers, and agentic artificial intelligence have further accelerated the development of intelligent manufacturing systems capable of autonomous perception, reasoning, and adaptive decision-making. These technologies are expected to play a significant role in the next generation of Industry 5.0-enabled production environments [46], [52], [55].

2. INDUSTRIAL QUALITY CONTROL AND SMART MANUFACTURING

Quality control has historically been a cornerstone of industrial manufacturing. Its primary objective is to ensure that products meet predefined specifications and customer requirements while minimizing waste, defects, and process variability. Over time, quality management has evolved from simple inspection-based approaches toward comprehensive organizational frameworks emphasizing continuous improvement, process optimization, and preventive action.

One of the most influential developments in quality management has been the establishment of international standards such as ISO 9001 and industry-specific frameworks including IATF 16949 for automotive manufacturing. These standards provide structured methodologies for quality planning, documentation, risk management, corrective actions, and continuous process improvement. [1], [2] Rather than focusing solely on final-product inspection, modern quality-management systems emphasize proactive monitoring throughout the production lifecycle.

The emergence of Industry 4.0 has further transformed quality-control practices by enabling real-time visibility across manufacturing operations. Smart factories employ interconnected cyber-physical systems that continuously collect information regarding machine performance, environmental conditions, process parameters, and product quality characteristics. Manufacturing Execution Systems (MES), Supervisory Control and Data Acquisition (SCADA) platforms, Enterprise Resource Planning (ERP) systems, and industrial sensors work together to create highly integrated operational ecosystems. The layered architecture commonly adopted in industrial environments connects physical devices at the shop-floor level with enterprise-level decision-making systems, facilitating comprehensive quality monitoring and process optimization. [8], [10], [39], [40]

Recent advances in computer vision and artificial intelligence have accelerated the automation of quality inspection tasks. High-resolution industrial cameras combined with machine learning algorithms can detect defects with levels of consistency often exceeding those of human inspectors. Automated inspection systems are increasingly used in industries such as semiconductor manufacturing, automotive production, textile processing, electronics assembly, and pharmaceutical packaging. These systems improve inspection speed, reduce labor costs, and enhance defect-detection accuracy. [21], [24], [36]

Despite these advancements, significant challenges remain. Industrial environments generate heterogeneous data streams originating from multiple sensor modalities, including visual, acoustic, thermal, pressure, vibration, and electrical measurements. Managing and interpreting such diverse information requires sophisticated analytical methods capable of extracting meaningful patterns while operating under strict latency constraints. Consequently, anomaly detection has become an essential component of modern quality-control strategies.

3. EVOLUTION OF ANOMALY DETECTION TECHNIQUES

3.1 Statistical Approaches

The earliest approaches to anomaly detection were based on statistical process-control techniques. Methods such as Shewhart control charts, Cumulative Sum (CUSUM) charts, and Exponentially Weighted Moving Average (EWMA) charts were designed to identify deviations from expected process behavior by monitoring critical operational variables. These techniques remain widely used because of their simplicity, interpretability, and low computational requirements.

Multivariate statistical techniques further expanded anomaly-detection capabilities by analyzing relationships among multiple process variables simultaneously. Principal Component Analysis (PCA), for example, became a popular method for reducing dimensionality and identifying abnormal process states. PCA-based monitoring systems have demonstrated effectiveness in detecting faults within chemical processes, manufacturing operations, and industrial control systems. [11], [13]

Despite their advantages, statistical approaches exhibit several limitations. Most methods assume linear relationships among variables and rely heavily on predefined assumptions regarding data distributions. In highly dynamic industrial environments characterized by nonlinear interactions and high-dimensional datasets, statistical methods often experience degraded performance. Additionally, they typically require manual feature engineering and may struggle to adapt to changing operating conditions. [11], [13]

3.2 Machine Learning Approaches

Machine learning introduced a more flexible framework for anomaly detection by allowing models to learn patterns directly from historical data. [14], [15], [16] Algorithms such as One-Class Support Vector Machines (OC-SVM), Isolation Forests, K-Means Clustering, and Density-Based Spatial Clustering of Applications with Noise (DBSCAN) have been widely applied in industrial anomaly detection. These approaches can identify abnormal observations without requiring extensive domain-specific rules.

Isolation Forests have become particularly popular because of their computational efficiency and ability to identify rare events within large datasets. Unlike traditional distance-based methods, Isolation Forests isolate anomalies through recursive partitioning, enabling rapid detection of outliers. [14] Similarly, One-Class Support Vector Machines construct decision boundaries around normal operational data, classifying observations outside the boundary as anomalies. [15], [16]

Machine-learning methods offer greater flexibility than statistical techniques and can model nonlinear relationships among process variables. However, they still face challenges related to feature engineering, parameter optimization, data imbalance, and interpretability.

Furthermore, many industrial datasets contain limited examples of defective conditions, making supervised training difficult.

4. DEEP LEARNING FOR INDUSTRIAL ANOMALY DETECTION

The limitations of conventional statistical and machine-learning techniques have driven researchers toward deep learning methodologies capable of automatically extracting complex representations from raw industrial data. Deep learning has fundamentally changed the field of anomaly detection by eliminating the need for extensive manual feature engineering and enabling systems to learn hierarchical patterns directly from large datasets. These capabilities are particularly valuable in industrial environments where anomalies often exhibit subtle visual characteristics that are difficult to describe using handcrafted features.

One of the earliest deep learning approaches employed Autoencoders (AEs), which learn compact representations of normal operational data. During inference, anomalies are identified through reconstruction errors because the network struggles to accurately reproduce patterns that differ significantly from those encountered during training. [18] Variational Autoencoders (VAEs) further enhanced this capability by modeling probabilistic latent representations, improving robustness under uncertain operating conditions.

Generative Adversarial Networks (GANs) introduced another significant advancement in anomaly detection. GAN-based models consist of a generator and a discriminator that compete during training, enabling the system to learn complex data distributions. When applied to industrial inspection tasks, GANs can generate realistic representations of defect-free products and identify anomalies through deviations from these learned distributions. Such approaches have demonstrated promising results in semiconductor manufacturing, textile inspection, and surface-defect analysis. [17], [19]

Time-series anomaly detection has also benefited substantially from deep learning innovations. Long Short-Term Memory (LSTM) networks and Transformer architectures have become increasingly popular for monitoring sequential industrial data. Unlike traditional machine-learning algorithms, these models effectively capture temporal dependencies and long-range relationships among sensor measurements. Applications include predictive maintenance, fault diagnosis, energy monitoring, and process optimization. [28], [29]

Recent studies have demonstrated that transformer-based architectures provide superior capability for modeling long-range spatial dependencies within complex industrial images. Vision Transformers (ViTs) employ self-attention mechanisms that enable global feature extraction without relying solely on local convolution operations. Consequently, transformer-based models have shown improved performance for detecting subtle structural defects, surface scratches, texture inconsistencies, and manufacturing irregularities across several benchmark datasets including MVTec AD and VisA [52], [53].

Furthermore, hybrid CNN–Transformer architectures combine the efficient feature extraction capability of convolutional networks with the contextual reasoning ability of transformers. These hybrid approaches have demonstrated improved robustness under varying illumination conditions and complex industrial environments, making them increasingly attractive for next-generation automated quality inspection systems [53].

The success of deep learning in industrial quality control can be attributed to several factors. First, deep neural networks automatically learn robust feature representations from raw

inputs, reducing dependence on domain expertise. Second, these models can process highly complex and nonlinear relationships among variables. Third, advances in Graphics Processing Units (GPUs) and distributed computing have made large-scale model training feasible for industrial organizations. [21], [22], [23]

Despite these advantages, deep learning models present several challenges. Training requires substantial computational resources, large datasets, and careful hyperparameter optimization. Furthermore, many deep learning systems function as black boxes, making it difficult for engineers to interpret decisions and establish trust in safety-critical environments. Consequently, explainability and transparency remain important research priorities.

5. COMPUTER VISION AND YOLO-BASED INDUSTRIAL INSPECTION

Among all deep learning approaches applied to industrial quality control, computer vision systems have achieved the most widespread adoption. Manufacturing industries increasingly rely on visual inspection because many defects manifest as observable surface irregularities, structural inconsistencies, assembly errors, or dimensional deviations. Traditional image-processing techniques often fail under varying lighting conditions, complex backgrounds, and diverse product geometries. Deep learning-based object detection models have addressed many of these limitations by learning visual representations directly from data.

The Convolutional Neural Network (CNN) serves as the foundation of most industrial computer-vision systems. CNNs utilize hierarchical convolutional filters to extract increasingly complex visual features, enabling the recognition of patterns ranging from simple edges to sophisticated defect structures. Their success has led to widespread deployment across sectors such as automotive manufacturing, electronics assembly, textile production, food processing, and pharmaceutical packaging. [21], [22], [23]

Within this domain, the YOLO (You Only Look Once) family of object detectors has emerged as one of the most influential architectures for real-time industrial inspection. Unlike traditional two-stage detectors, which first generate region proposals and subsequently classify them, YOLO performs localization and classification simultaneously in a single forward pass through the network. This design dramatically reduces computational latency while maintaining high detection accuracy. [24], [25]

Recent developments have culminated in the YOLOv8 architecture, which introduces several improvements relevant to industrial anomaly detection. The model adopts an anchor-free detection paradigm, eliminating the need for predefined anchor boxes and simplifying the detection process. This design improves flexibility when identifying small, irregularly shaped defects commonly encountered in manufacturing environments. Additionally, YOLOv8 incorporates advanced optimization mechanisms such as Distribution Focal Loss (DFL) and Complete Intersection over Union (CIoU), improving both classification confidence and bounding-box precision. [27]

Recent research has also investigated integrating transformer backbones into object detection architectures. These hybrid models improve feature representation while maintaining competitive inference speed, making them suitable for industrial anomaly detection involving highly complex textures and irregular defect geometries [52], [53].

The effectiveness of modern object-detection systems depends heavily on the availability of high-quality datasets. Several benchmark datasets have become widely adopted within the

industrial anomaly-detection community. The MVTec AD dataset remains one of the most extensively utilized benchmarks due to its diverse categories and precise defect annotations. Other datasets, including VisA, MPDD, Real-IAD, and ReinAD, provide increasingly realistic industrial scenarios involving complex assemblies, metallic surfaces, and large-scale production environments. These datasets enable researchers to evaluate model performance under diverse operational conditions and facilitate fair comparison among competing approaches. [20], [36]




Figure 5.1: Architecture for a Yolo based inspection system

The proposed pipeline presents a complete framework for developing a YOLO-based industrial anomaly detection system. It processes visual inputs to detect anomalies in real time and incorporates evaluation metrics such as precision, recall, mAP, and inference time to assess model performance. These metrics facilitate iterative refinement and optimization, ensuring that the system meets the accuracy, speed, and reliability requirements for industrial deployment.

Although YOLO-based inspection systems[Figure 5.1], demonstrate remarkable performance, practical deployment introduces additional considerations. Industrial environments often require continuous operation under strict latency constraints. Consequently, model compression, quantization, edge deployment, and hardware acceleration have become important areas of research. These developments aim to balance detection accuracy with computational efficiency, enabling real-time quality control in production settings.

6. MULTI-AGENT SYSTEMS IN MANUFACTURING AND INDUSTRIAL AUTOMATION

The increasing complexity of manufacturing systems has motivated researchers to explore decentralized approaches to process monitoring and decision-making. Multi-Agent Systems (MAS) represent one of the most promising solutions to this challenge. A Multi-Agent System consists of multiple autonomous software entities, referred to as agents, that interact with one another and their environment to achieve individual and collective objectives. [3], [4] Each agent possesses localized knowledge, decision-making capabilities, and communication mechanisms that allow collaborative problem solving.

The concept of agent-based manufacturing emerged from the recognition that centralized control architectures often struggle to manage large-scale industrial operations. [5], [7] As production facilities expand and incorporate thousands of interconnected devices, centralized systems become vulnerable to communication bottlenecks, computational overload, and single points of failure. Multi-Agent Systems address these limitations by distributing intelligence across multiple computational entities.

Within industrial environments, agents may assume specialized responsibilities depending on their operational roles. Sensor agents acquire data from physical devices and convert it into standardized formats suitable for analysis. Analysis agents process incoming data streams using anomaly-detection algorithms and machine-learning models. Coordination agents manage communication among local subsystems and facilitate collaborative decision-making. Meta-agents oversee system-wide performance, maintain historical records, and optimize operational strategies [5], [6], [38] based on accumulated knowledge. The hierarchical architecture described in the dissertation reflects this layered approach to industrial intelligence.

Research projects such as GRACE demonstrated the feasibility of agent-based manufacturing environments by integrating Product Agents, Resource Agents, and Quality-Control Agents into coordinated production systems. These implementations showed that decentralized architectures could improve flexibility, adaptability, and fault tolerance while reducing dependence on centralized supervisory mechanisms.

A major advantage of Multi-Agent Systems lies in their resilience. If a single agent fails, the remainder of the system can continue functioning with minimal disruption. This characteristic is particularly valuable in manufacturing environments where operational continuity is critical. [3], [5] Furthermore, agent-based architectures naturally support scalability because new agents can be introduced without requiring substantial redesign of the overall system.

The rise of Industry 4.0 has renewed interest in MAS research. Modern manufacturing systems increasingly demand autonomous adaptation, distributed intelligence, and real-time responsiveness. Multi-Agent Systems provide a conceptual and technological foundation capable of supporting these requirements while enabling seamless integration with artificial intelligence and machine-learning technologies.

7. AGENT COMMUNICATION, COORDINATION, AND DECISION FUSION

Effective communication is essential for the successful operation of Multi-Agent Systems. Since individual agents possess only partial knowledge of the environment, collaboration is necessary to achieve global objectives. Consequently, considerable research has focused on developing communication protocols, coordination mechanisms, and decision-fusion strategies suitable for industrial environments.

One of the most influential standards in this area is the Foundation for Intelligent Physical Agents (FIPA) framework. FIPA defines Agent Communication Languages (ACLs) that enable structured interactions among heterogeneous agents. [30], [31] Through standardized performatives such as *request*, *inform*, *propose*, and *confirm*, agents can exchange information, negotiate responsibilities, and coordinate actions. Traditional MAS frameworks such as JADE were built around FIPA standards and remain widely cited in academic literature.

Recent developments have shifted attention toward Python-based frameworks capable of integrating more naturally with modern artificial intelligence ecosystems. SPADE (Smart Python Agent Development Environment) has gained popularity because it combines asynchronous messaging capabilities with direct compatibility with machine-learning libraries such as PyTorch and TensorFlow. [32] SPADE utilizes the Extensible Messaging and Presence Protocol (XMPP) to support scalable communication among distributed agents while maintaining low implementation complexity.

Beyond communication, coordination mechanisms determine how agents collectively respond to detected anomalies. Common approaches include majority voting, weighted averaging, Bayesian decision fusion, contract-net protocols, and consensus algorithms. Decision fusion is particularly important in industrial environments because multiple sensors and detection models may produce conflicting observations. Combining outputs from diverse sources generally improves reliability and reduces false alarms. [36]

Research has demonstrated that fusion strategies can significantly enhance anomaly-detection performance when compared with individual detectors. By integrating information from multiple agents, manufacturing systems can develop a more comprehensive understanding of operational conditions and generate more robust decisions. The fusion architecture described in the dissertation illustrates how outputs from multiple analysis agents can be aggregated through majority voting, weighted averaging, and Bayesian reasoning to produce final quality-control decisions.

As industrial systems continue to evolve, communication and coordination mechanisms will play an increasingly important role in enabling autonomous factories. Future manufacturing environments are likely to rely heavily on collaborative intelligence, where networks of specialized agents work together to monitor processes, detect anomalies, optimize production schedules, and adapt dynamically to changing conditions [Figure 7.1].




Figure 7.1: Multi agent based industrial anomaly detection system

The proposed multi-agent industrial anomaly detection system enhances detection performance by enabling multiple specialized agents to collaborate during data analysis and decision-making. This architecture reduces human intervention, improves detection accuracy, and increases the reliability of predictions. By combining the outputs of individual agents through a coordinated decision-making process, the system minimizes the bias and limitations of any single agent, resulting in more robust and accurate anomaly detection.

8. EDGE COMPUTING, FOG COMPUTING, AND DISTRIBUTED INTELLIGENCE

The emergence of Industry 4.0 has significantly increased the volume of data generated by manufacturing systems. Modern production facilities continuously collect information from sensors, cameras, programmable logic controllers, robotic systems, and enterprise platforms. While centralized cloud computing provides substantial computational resources for data storage and analytics, transmitting all industrial data to remote servers introduces latency, bandwidth consumption, and reliability concerns. Consequently, researchers have increasingly explored edge and fog computing architectures as complementary solutions for real-time industrial intelligence.

Edge computing refers to the deployment of computational resources close to data sources. Instead of transmitting raw sensor information to distant cloud servers, edge devices perform localized processing and decision-making. This approach reduces communication delays, minimizes bandwidth requirements, and enables immediate responses to critical operational events. [42], [43] In manufacturing environments, edge devices may include industrial PCs, embedded systems, smart cameras, and specialized AI accelerators capable of executing machine-learning models directly on the production floor.

Fog computing extends this concept by introducing an intermediate processing layer between edge devices and cloud infrastructures. Fog nodes aggregate information from multiple edge systems, perform regional analysis, and coordinate localized decision-making. This hierarchical architecture enables organizations to balance responsiveness and computational

efficiency. [42], [43] Critical anomaly-detection tasks can be executed at the edge, while computationally intensive activities such as model retraining and historical analytics can be delegated to fog or cloud environments.

The integration of Multi-Agent Systems with edge computing has emerged as a particularly promising research direction. In such architectures, edge agents process local sensor streams and execute anomaly-detection algorithms independently. Coordination agents operating at the fog layer aggregate information from multiple production stations, while cloud-based meta-agents perform long-term optimization and strategic planning. The hierarchical architecture described in the dissertation follows this design philosophy, distributing computational responsibilities across sensor agents, analysis agents, coordination agents, and cloud-based supervisory agents. [Figure 8.1]




Figure 8.1 A cloud-based anomaly detection agent

This system enables deployment across multiple industrial sites through a centralized architecture. This approach facilitates efficient coordination among distributed systems, reduces downtime by enabling rapid detection and response to equipment malfunctions, and improves overall system reliability. By aggregating data from multiple locations, the cloud platform supports more informed and accurate decision-making, allowing the anomaly detection agents to continuously learn and adapt while maintaining consistent performance across all connected facilities.

Recent studies demonstrate that edge-based anomaly detection can significantly reduce response times while improving operational reliability. By eliminating dependence on centralized infrastructures, edge computing enhances system resilience and supports continued operation even during network disruptions. These characteristics make edge-enabled MAS architectures particularly suitable for industrial quality-control applications. [5], [38], [42]

9. COMPARATIVE ANALYSIS OF MULTI-AGENT FRAMEWORKS

The selection of an appropriate Multi-Agent System framework significantly influences the performance, scalability, and maintainability of industrial anomaly detection systems. Traditional frameworks such as **JADE** and **SPADE** have been extensively adopted in manufacturing research because they provide mature communication protocols, decentralized agent coordination, and compliance with established Multi-Agent System standards. More recently, **LLM-based agent frameworks** including **AutoGen**, **LangGraph**, **OpenAI Agent SDK**, and **Google Agent Development Kit (ADK)** have emerged as alternatives that introduce advanced reasoning and decision-making capabilities through Large Language Models. Although these frameworks were initially developed for conversational AI and autonomous task execution, researchers have begun investigating their applicability to industrial automation and intelligent manufacturing systems.

JADE (Java Agent Development Framework) remains one of the most widely used platforms for implementing distributed Multi-Agent Systems. Built upon the Foundation for Intelligent Physical Agents (FIPA) standards, JADE provides standardized Agent Communication Language (ACL) messaging, lifecycle management, directory services, and robust agent coordination. These characteristics make JADE highly suitable for distributed manufacturing control, resource allocation, and process coordination. However, its Java-based architecture introduces integration challenges with modern deep learning frameworks such as PyTorch and TensorFlow, which are predominantly developed in Python. Consequently, deploying advanced computer vision models within JADE-based systems often requires additional middleware or cross-language communication mechanisms. [3], [5], [30], [31]

SPADE (Smart Python Agent Development Environment) has become an increasingly attractive alternative because it is implemented entirely in Python. SPADE employs asynchronous communication through the Extensible Messaging and Presence Protocol (XMPP), allowing seamless interaction among distributed agents while supporting direct integration with machine learning libraries. Unlike JADE, SPADE enables developers to incorporate deep learning inference models such as YOLOv8 directly into agent behaviors without requiring complex interoperability layers. This compatibility significantly simplifies the implementation of decentralized industrial anomaly detection systems where each edge agent performs localized inference before communicating results to higher-level coordination agents. [32]

The rapid advancement of Large Language Models has introduced a new category of agentic frameworks capable of autonomous planning, tool usage, reasoning, and collaborative problem solving. Frameworks such as AutoGen, LangGraph, OpenAI Agent SDK, and Google ADK extend traditional agent architectures by enabling agents to dynamically generate plans, invoke external tools, maintain memory, and coordinate multi-step workflows. These capabilities are particularly valuable for high-level manufacturing tasks including production scheduling, maintenance planning, report generation, and decision support.

Despite these advantages, LLM-based frameworks remain less suitable for latency-sensitive industrial anomaly detection applications. Visual inspection systems require deterministic execution, predictable inference latency, and real-time responses measured in milliseconds. Current LLM agents depend on computationally intensive transformer models and, in many cases, cloud-based inference services that introduce additional latency and operational

uncertainty. Furthermore, hallucination, variable reasoning behavior, and increased computational costs present challenges for deployment in safety-critical manufacturing environments.

For these reasons, current industrial research generally favors hybrid architectures rather than complete replacement of conventional Multi-Agent Systems. In such architectures, lightweight frameworks such as SPADE or JADE perform real-time equipment monitoring and anomaly detection at the edge, while LLM-based agents operate at supervisory levels to analyze production trends, interpret historical quality reports, recommend maintenance strategies, and assist human operators in complex decision-making. This layered approach combines the deterministic execution required for industrial control with the advanced reasoning capabilities offered by modern foundation models.

Study	Method / Architecture	Dataset / Domain	Precision	mAP	Inference Time	Key Contribution
Peng et al. (2025)	YOLOv7 + ConSinGAN synthetic augmentation	DIP-switch board anomalies	—	95.50% (mAP@0.5)	≈285 ms	Synthetic augmentation improved YOLOv7 accuracy.
Pati et al. (2026)	YOLOv3, YOLOv5 & YOLOv8	Weld surface defects	Highest (NR)	≈0.59	Fastest (NR)	YOLOv8 achieved the best speed–accuracy trade-off.
Okano et al. (2025)	YOLOv8 Nano & Small	Polyamide tube inspection	0.932 / 0.951	0.938 / 0.941	470 / 1315 ms	Demonstrated efficient embedded deployment on Raspberry Pi.
Kalušev et al. (2026)	Edge MAS with YOLO + LLM	Generic object detection	—	—	Real-time	Validated multi-agent vision on low-cost hardware.
Hudnurkar et al. (2025)	YOLOv8 vs. VAE-GAN	Screw defect detection	95–97%	—	Faster than VAE-GAN	YOLOv8 outperformed hybrid anomaly detection.
Liu et al. (2025)	Task-Aligned YOLOv8	Spot weld defects	Recall ↑12.26%	+1.35%	2.8% faster	Lightweight head improved accuracy and speed.
Krassnig et al. (2025)	Synthetic + real defect training	ISP-AD	—	—	—	Real defects enhanced model generalization.
Zhang et al. (2025) – AutoIAD	Multi-agent automated pipeline	MVTec AD & custom datasets	—	Improved	Development time ↓40%	Automated end-to-end model development.

Table 9.1 Comparison of various studies made in multiagent framework

Despite their promise, current LLM-based frameworks face significant limitations in manufacturing contexts. Real-time industrial systems require predictable latency, deterministic behavior, and high reliability, characteristics that are not always guaranteed by

current generative AI architectures. Consequently, additional research is needed before such frameworks can be widely adopted in safety-critical manufacturing environments.

9.2. EMERGING TRENDS: AGENTIC AI, INDUSTRIAL LARGE LANGUAGE MODELS, AND TRANSFORMER-BASED ANOMALY DETECTION

The rapid advancement of Large Language Models (LLMs) has introduced a new paradigm for intelligent industrial automation. Unlike conventional Multi-Agent Systems, which primarily execute predefined behaviors and communication protocols, modern agentic AI frameworks enable autonomous reasoning, task planning, memory management, and dynamic tool utilization. These capabilities allow software agents to perform complex multi-step workflows, interact with external applications, and collaborate with both human operators and other intelligent agents.[46]

Recent frameworks such as **OpenAI Agent SDK** [49], **AutoGen** [47], **LangGraph** [50], **Google Agent Development Kit (ADK)** [51], and **CrewAI** have demonstrated the feasibility of orchestrating multiple specialized agents capable of autonomous collaboration. Although originally designed for knowledge-intensive applications, researchers have begun exploring their use in predictive maintenance, manufacturing analytics, intelligent scheduling, industrial report generation, and adaptive production planning. These frameworks extend traditional Multi-Agent Systems by integrating Large Language Models with external software tools, databases, and sensor interfaces, thereby creating more flexible and context-aware industrial decision-support systems.

Another significant development is the growing adoption of transformer-based architectures for industrial anomaly detection [52]. While Convolutional Neural Networks remain the dominant approach for image-based quality inspection, Vision Transformers (ViTs) and hybrid CNN-Transformer models have demonstrated superior capability in learning long-range spatial dependencies. Unlike CNNs, which primarily focus on local receptive fields, transformers utilize self-attention mechanisms that enable the model to capture global contextual information across the entire image [52], [53]. This characteristic has proven particularly valuable for detecting subtle surface defects, structural inconsistencies, and complex industrial anomalies that may span multiple regions of a product.

Recent industrial benchmarks indicate that transformer-based anomaly detection models consistently outperform conventional convolutional architectures on challenging datasets such as MVTec AD, VisA, BTAD, MPDD, and Real-IAD. Furthermore, multimodal transformer architectures capable of simultaneously processing visual, thermal, acoustic, and vibration data have demonstrated improved robustness under complex manufacturing conditions [53],[54].

Despite these promising developments, current LLM-based agentic frameworks are not yet optimized for low-latency industrial control. Manufacturing environments require deterministic execution, predictable response times, and high operational reliability, whereas modern Large Language Models often introduce variable inference latency and increased computational requirements. Consequently, researchers increasingly advocate hybrid architectures in which lightweight edge agents perform real-time anomaly detection using optimized computer vision models, while higher-level LLM agents provide strategic reasoning, maintenance recommendations, production optimization, and natural-language interaction with engineers.

Future industrial automation is therefore expected to integrate three complementary technologies: distributed Multi-Agent Systems for decentralized coordination, transformer-based deep learning models for high-accuracy anomaly detection, and Large Language Models for autonomous planning and intelligent decision support. Such hybrid systems represent one of the most promising research directions for next-generation smart manufacturing [55].

Framework	Language	Strengths	Limitations	Suitability
JADE	Java	Mature, FIPA compliant, reliable communication	Difficult ML integration, Java ecosystem	High
SPADE	Python	Native Python support, asynchronous messaging, easy AI integration	Smaller ecosystem than JADE	Very High
ROS/ROS-Industrial	C++ / Python	Excellent robotics communication	Not a true MAS framework	Medium
AutoGen	Python	Multi-agent reasoning, collaborative workflows	High latency, LLM dependency	Medium
LangGraph	Python	Structured AI workflows and memory	Limited real-time capability	Medium
OpenAI Agent SDK	Python/JavaScript	Advanced planning and tool orchestration	Requires LLM inference, not deterministic	Medium-Low
Google ADK	Python	Easy AI service integration	Early-stage industrial adoption	Medium-Low

Table 9.2 Comparison of JADE, SPADE and LLM based frameworks

10. RESEARCH GAPS AND CRITICAL ANALYSIS

Although substantial progress has been achieved in industrial anomaly detection and Multi-Agent Systems, several critical research challenges remain unresolved.

The first major gap concerns the separation between anomaly-detection research and industrial control architectures. Most studies focus primarily on improving detection accuracy through increasingly sophisticated machine-learning models. Conversely, MAS research often concentrates on communication, coordination, and distributed control without fully integrating state-of-the-art anomaly-detection techniques. Consequently, relatively few studies investigate end-to-end systems that combine modern deep-learning detectors with decentralized agent-based infrastructures. [5], [6], [36], [38]

A second challenge involves scalability. While deep-learning models achieve remarkable accuracy under controlled laboratory conditions, deploying these systems across large manufacturing facilities introduces substantial computational demands. Edge devices frequently possess limited memory, storage, and processing capabilities, restricting the complexity of models that can be executed in real time. [5], [6], [36], [38]

Third, industrial datasets remain highly imbalanced. Most manufacturing processes generate vast quantities of normal operational data while producing relatively few defective samples. This imbalance complicates model training and can reduce anomaly-detection performance.

Furthermore, many publicly available datasets fail to fully capture the diversity and complexity of real-world production environments. [5], [6], [36], [38]

Another important limitation concerns explainability. Industrial operators often require transparent decision-making processes, particularly when anomaly-detection systems influence production schedules or product acceptance decisions. Deep neural networks generally provide limited interpretability, creating barriers to industrial adoption. [5], [6], [36], [38]

Although recent LLM-based agent frameworks have demonstrated impressive reasoning capabilities, their application to real-time industrial anomaly detection remains limited because of latency constraints, computational requirements, and the need for deterministic execution [47], [49], [50], [51].

Cybersecurity has received insufficient attention within many MAS-based quality-control architectures. As industrial systems become increasingly interconnected, protecting agent communications and preventing malicious interference will become essential for maintaining operational integrity.

Most existing studies investigate Multi-Agent Systems, deep learning models, and Large Language Models independently. Very limited research has explored unified architectures capable of combining real-time computer vision, decentralized agent communication, transformer-based anomaly detection, and autonomous reasoning within a single industrial framework. This integration represents one of the most promising research directions for Industry 5.0 [47], [52], [55].

11. FUTURE RESEARCH DIRECTIONS

Future industrial quality control systems are expected to evolve toward intelligent autonomous ecosystems that integrate edge computing, transformer-based computer vision, Multi-Agent Systems, digital twins, and Large Language Models. Such hybrid systems will not only detect manufacturing defects in real time but will also explain their causes, recommend corrective actions, generate maintenance reports, and coordinate production schedules through collaborative intelligent agents [46], [47], [55].

Another promising direction is the integration of multimodal sensor fusion. Future quality-control systems are likely to combine visual imagery, thermal measurements, acoustic emissions, vibration signals, and process parameters within unified anomaly-detection architectures. Multimodal learning can improve robustness and provide a more comprehensive understanding of manufacturing processes.

The advancement of edge AI hardware will further expand opportunities for real-time industrial intelligence. Specialized accelerators capable of executing complex deep-learning models directly on production equipment will reduce dependence on centralized infrastructures and enable faster decision-making. [42], [43]

Research into explainable artificial intelligence (XAI) is also expected to play a critical role. Developing interpretable anomaly-detection models will increase operator trust and facilitate adoption within highly regulated industries.

The convergence of MAS, digital twins, and agentic AI represents another transformative opportunity. Digital twins provide virtual representations of physical manufacturing systems, while intelligent agents can continuously interact with these models to simulate, predict, and optimize operational outcomes. Combining these technologies could create highly adaptive manufacturing ecosystems capable of self-monitoring and self-correction [8], [10], [38].

Finally, Future industrial automation is expected to integrate conventional Multi-Agent Systems, transformer-based anomaly detection, and agentic AI frameworks to create intelligent manufacturing environments capable of autonomous monitoring, adaptive planning, and continuous learning [46], [52], [55].

12. CONCLUSION

The transition toward Industry 4.0 has fundamentally transformed the requirements of industrial quality control. Traditional inspection methods and statistical monitoring techniques are increasingly inadequate for managing the complexity, scale, and speed of modern manufacturing systems. As a result, anomaly detection has emerged as a critical technological capability for identifying defects, process deviations, and equipment failures before they generate significant operational losses.

This review examined the evolution of anomaly-detection methodologies, ranging from classical statistical approaches to advanced deep-learning architectures. Particular attention was given to computer-vision systems and YOLO-based object detectors, which have demonstrated exceptional effectiveness for real-time industrial inspection. Simultaneously, the review explored the role of Multi-Agent Systems as a decentralized architectural framework capable of supporting intelligent monitoring, distributed decision-making, and adaptive quality control.

The literature indicates that anomaly detection and Multi-Agent Systems are highly complementary technologies. Deep-learning models provide powerful analytical capabilities, while agent-based architectures offer scalability, fault tolerance, and operational flexibility. However, the integration of these domains remains relatively underexplored, creating substantial opportunities for future research.

Emerging technologies including Vision Transformers, foundation models, and agentic AI significantly broaden the scope of intelligent manufacturing by extending anomaly detection beyond visual inspection toward autonomous industrial reasoning. Their integration with decentralized Multi-Agent Systems represents a promising step toward fully autonomous smart factories capable of continuous adaptation and self-optimization.

By combining intelligent anomaly detection, edge computing, multimodal sensor fusion, and decentralized agent coordination, next-generation manufacturing systems can achieve unprecedented levels of autonomy, reliability, and operational efficiency. Such developments will play a central role in realizing the vision of fully intelligent and self-optimizing smart factories.

ACKNOWLEDGMENTS

The authors declare that no financial or institutional support was received for this research.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

REFERENCES:

- [1] International Organization for Standardization, *ISO 9001:2015 Quality Management Systems—Requirements*. Geneva, Switzerland: ISO, 2015.
- [2] International Automotive Task Force, *IATF 16949:2016 Quality Management System Standard for Automotive Production*. Detroit, MI, USA, 2016.
- [3] M. Wooldridge, *An Introduction to MultiAgent Systems*, 2nd ed. Chichester, U.K.: John Wiley & Sons, 2009.
- [4] N. R. Jennings and M. Wooldridge, “Applications of intelligent agents,” in *Agent Technology: Foundations, Applications and Markets*. Berlin, Germany: Springer, 1998, pp. 3–28.
- [5] P. Leitão, “Agent-based distributed manufacturing control: A state-of-the-art survey,” *Engineering Applications of Artificial Intelligence*, vol. 22, no. 7, pp. 979–991, 2009.
- [6] P. Leitão, S. Karnouskos, L. Ribeiro, J. Lee, T. Strasser, and A. W. Colombo, “Smart agents in industrial cyber–physical systems,” *Proceedings of the IEEE*, vol. 104, no. 5, pp. 1086–1101, 2016.
- [7] V. Marík and P. Vrba, “Manufacturing intelligence: An emergent property of cooperative industrial agents,” *IEEE Transactions on Systems, Man, and Cybernetics Part C*, vol. 34, no. 1, pp. 59–68, 2004.
- [8] F. Tao, Q. Qi, A. Liu, and A. Kusiak, “Data-driven smart manufacturing,” *Journal of Manufacturing Systems*, vol. 48, pp. 157–169, 2018.
- [9] A. Kusiak, “Smart manufacturing,” *International Journal of Production Research*, vol. 56, nos. 1–2, pp. 508–517, 2018.
- [10] J. Lee, B. Bagheri, and H. A. Kao, “A cyber-physical systems architecture for Industry 4.0-based manufacturing systems,” *Manufacturing Letters*, vol. 3, pp. 18–23, 2015.
- [11] S. Yin, X. Li, H. Gao, and O. Kaynak, “Data-based techniques focused on modern industry: An overview,” *IEEE Transactions on Industrial Electronics*, vol. 62, no. 1, pp. 657–667, 2015.
- [12] C. C. Aggarwal, *Outlier Analysis*, 2nd ed. Cham, Switzerland: Springer, 2017.
- [13] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009.

-
- [14] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, Pisa, Italy, 2008, pp. 413–422.
- [15] B. Schölkopf, J. Platt, J. Shawe-Taylor, A. Smola, and R. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
- [16] D. M. J. Tax and R. P. W. Duin, "Support vector data description," *Machine Learning*, vol. 54, no. 1, pp. 45–66, 2004.
- [17] I. Goodfellow et al., "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, Montreal, Canada, 2014, pp. 2672–2680.
- [18] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *Proc. Int. Conf. Learning Representations (ICLR)*, 2014.
- [19] T. Schlegl, P. Seeböck, S. Waldstein, G. Langs, and U. Schmidt-Erfurth, "Unsupervised anomaly detection with generative adversarial networks," *Information Processing in Medical Imaging*, pp. 146–157, 2017.
- [20] P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, "MVTec AD—A comprehensive real-world dataset for unsupervised anomaly detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9592–9600.
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [22] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015.
- [23] C. Szegedy et al., "Going deeper with convolutions," in *Proc. CVPR*, 2015, pp. 1–9.
- [24] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, real-time object detection," in *Proc. CVPR*, 2016, pp. 779–788.
- [25] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *Proc. CVPR*, 2017, pp. 7263–7271.
- [26] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," *arXiv preprint arXiv:2004.10934*, 2020.
- [27] G. Jocher et al., "YOLO by Ultralytics," GitHub Repository, 2024.
- [28] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [29] A. Vaswani et al., "Attention Is All You Need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [30] FIPA, *Foundation for Intelligent Physical Agents Agent Communication Language Specification*, Geneva, Switzerland, 2002.
- [31] F. Bellifemine, G. Caire, and D. Greenwood, *Developing Multi-Agent Systems with JADE*. Chichester, U.K.: Wiley, 2007.

-
- [32] M. Palanca, J. Del Val, and V. Julian, "SPADE 3: Supporting the new generation of multi-agent systems," *IEEE Access*, vol. 8, pp. 182537–182549, 2020.
- [33] M. Quigley et al., "ROS: An open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, 2009.
- [34] J. Dean et al., "Large scale distributed deep networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1223–1231.
- [35] S. Karnouskos and A. W. Colombo, "Architecting the next generation of service-based SCADA/DCS system of systems," in *Proc. IECON*, 2011, pp. 359–364.
- [36] Y. Liu, Y. Yang, X. Lv, and L. Wang, "Deep-learning-based industrial anomaly detection: A survey," *IEEE Access*, vol. 10, pp. 122456–122482, 2022.
- [37] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for image recognition at scale," in *Proc. ICLR*, 2021.
- [38] P. Leitão, J. Barbosa, and A. W. Colombo, "Industrial agents in Industry 4.0," *IEEE Industrial Electronics Magazine*, vol. 15, no. 2, pp. 18–29, 2021.
- [39] H. Lasi, P. Fettke, H.-G. Kemper, T. Feld, and M. Hoffmann, "Industry 4.0," *Business & Information Systems Engineering*, vol. 6, no. 4, pp. 239–242, 2014.
- [40] K. Zhou, T. Liu, and L. Zhou, "Industry 4.0: Towards future industrial opportunities and challenges," in *Proc. Fuzzy Systems and Knowledge Discovery*, 2015, pp. 2147–2152.
- [41] T. Erl, W. Khattak, and P. Buhler, *Big Data Fundamentals*. Upper Saddle River, NJ, USA: Prentice Hall, 2016.
- [42] W. Shi and S. Dustdar, "The promise of edge computing," *Computer*, vol. 49, no. 5, pp. 78–81, 2016.
- [43] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, 2017.
- [44] H. Boyes, B. Hallaq, J. Cunningham, and T. Watson, "The industrial internet of things (IIoT): An analysis framework," *Computers in Industry*, vol. 101, pp. 1–12, 2018.
- [45] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of Things (IoT): A vision, architectural elements, and future directions," *Future Generation Computer Systems*, vol. 29, no. 7, pp. 1645–1660, 2013.
- [46] T. Wu et al., "A Survey of Large Language Models," *arXiv preprint arXiv:2303.18223*, 2023.
- [47] Q. Chen et al., "AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation," *Microsoft Research*, 2024.
- [48] H. Touvron et al., "Llama 3: Open Foundation and Instruction Models," *Meta AI*, 2024.
- [49] OpenAI, "OpenAI Agents SDK Documentation," 2025.
- [50] LangChain Inc., "LangGraph: Stateful Multi-Agent Applications with LLMs," 2025.
- [51] Google DeepMind, "Google Agent Development Kit (ADK)," 2025.

-
- [52] C. Jia *et al.*, "Vision Transformer for Industrial Anomaly Detection: A Survey," *Engineering Applications of Artificial Intelligence*, 2024.
- [53] Y. Cao *et al.*, "Transformer-Based Industrial Visual Inspection: Recent Advances and Challenges," *IEEE Access*, vol. 12, pp. 110234–110258, 2024.
- [54] Y. Huang *et al.*, "Multimodal Industrial Anomaly Detection Using Vision Transformers," *Pattern Recognition*, vol. 150, 2025.
- [55] Z. Li *et al.*, "Foundation Models for Smart Manufacturing: Opportunities and Challenges," *Journal of Manufacturing Systems*, vol. 75, pp. 25–42, 2025.