American Journal of Applied Science and Technology

VOLUME Vol.06 Issue 08 2026
PAGE NO. 167-174
DOI10.37547/ajast/Volume06Issue08-13

Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing

Haruto Nakamori

Generative AI Research Division, Sakura Intelligent Systems, Japan

Aiko Fujimori

Autonomous AI Engineering Group, Tokyo Digital Innovation Labs, Japan

Received: 28 June 2026; **Accepted:** 22 July 2026; **Published:** 22 August 2026

Abstract: Modern manufacturing environments generate large volumes of operational exceptions involving production disruptions, material shortages, quality deviations, foreign-object detection, equipment anomalies, and process interruptions. Enterprise resource planning systems such as SAP S/4HANA Manufacturing provide structured transactional information for production planning and execution, but conventional exception management frequently remains dependent on human monitoring, rule-based alerts, and sequential investigation. This paper proposes a Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing that integrates specialized intelligent agents for exception detection, contextual diagnosis, impact assessment, decision generation, and controlled resolution. The conceptual foundation is derived from the provided literature on machine-vision-based foreign-object detection, anomaly detection, three-dimensional perception, fault early warning, and industrial defect analysis. These studies demonstrate the value of automated perception and intelligent anomaly identification across heterogeneous manufacturing environments. Building upon these principles, the proposed framework organizes manufacturing exceptions into an agent-based decision pipeline connected to SAP production and material processes. The framework emphasizes autonomous reasoning while maintaining transactional controls through confidence thresholds, approval gates, and exception escalation. The resulting model indicates that multi-agent coordination can transform isolated anomaly signals into contextualized manufacturing decisions, thereby reducing response latency and improving operational resilience. However, limitations remain concerning data quality, explainability, integration complexity, model hallucination, and authorization of autonomous actions. The study contributes a research-oriented architecture for integrating generative AI agents with enterprise manufacturing exception management.

Keywords: Multi-Agent Systems, Generative AI, SAP S/4HANA, Manufacturing, Exception Management, Autonomous Decision-Making, Anomaly Detection, Intelligent Manufacturing, Production Planning.

1. INTRODUCTION

1.1 Background

Manufacturing systems operate through tightly coupled relationships among production orders, materials, machines, quality processes, inventory, and operational schedules. A disruption in one component can propagate through several downstream processes. A material shortage may delay production; an equipment anomaly may reduce available capacity; and a detected foreign object or product defect may require inspection, segregation, or rework. Consequently, manufacturing exception management is not simply an alert-generation problem. It is a contextual decision problem in which the system must determine what occurred, why it occurred, what processes are affected, and which corrective action is appropriate.

The provided literature demonstrates increasing sophistication in automated industrial perception. Chao, Kai, and Zhang (2020) examined machine-vision-based detection of foreign bodies in tobacco processing, while Chen et al. (2022) investigated automated foreign-matter detection and sorting for PVC powder. Ishii, Mizuta, and Todo (1998) similarly demonstrated the use of real-time image processing for detecting foreign substances in medicinal-solution bottles. These studies establish an important foundation: manufacturing exceptions can be detected automatically when operational processes are converted into observable signals.

More advanced approaches extend this concept from simple visual inspection toward intelligent anomaly recognition. Hou et al. (2022) proposed an early-warning method based on adversarial auto-encoders for mechanical fault detection. Zhang et al. (2021) examined attention-based computer vision for detecting foreign objects in coal processing, while Xu et al. (2021) investigated an efficient foreign-object detection network for power substations. Such work indicates that intelligent systems can identify abnormal conditions before or during operational failure.

The central challenge addressed by the Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing is the transformation of these heterogeneous anomaly signals into coordinated enterprise-level actions. Detection alone does not resolve an exception. A detected anomaly must be interpreted against production schedules, material requirements, equipment conditions, quality status, and operational priorities. The proposed framework therefore treats exception management as a multi-stage reasoning process rather than a single prediction task.

1.2 Problem Statement

Conventional manufacturing exception workflows commonly separate detection, analysis, decision-making, and execution. Machine-vision systems may detect defects, maintenance systems may identify equipment anomalies, and ERP systems may record production consequences, but these signals can remain operationally disconnected. This creates three major problems.

First, isolated alerts may lack sufficient context for determining business impact. Second, human operators must frequently correlate information across multiple processes before selecting an intervention. Third, even when an appropriate decision is identified, execution may require multiple transactional steps.

The research problem is therefore to establish an architecture in which specialized AI agents can collaboratively interpret manufacturing exceptions and recommend or execute appropriate responses within a controlled SAP S/4HANA environment.

1.3 Research Objectives

The study has five objectives: to establish a conceptual model for autonomous manufacturing exception management; to define specialized agents for detection, diagnosis, impact analysis, decision generation, and execution; to connect multi-agent reasoning with SAP manufacturing processes; to identify mechanisms for controlling autonomous decisions; and to evaluate the conceptual benefits and limitations of the proposed architecture.

1.4 Scope and Significance

The framework focuses on manufacturing exceptions associated with production, quality, materials, equipment, and operational inspection. It does not assume unrestricted autonomous modification of enterprise data. Instead, autonomy is governed by confidence thresholds, authorization policies, and escalation mechanisms. This distinction is significant because enterprise AI must optimize operational responsiveness without compromising transactional integrity.

The proposed Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing therefore positions generative AI as an orchestration and reasoning layer rather than merely another predictive model. Its significance lies in connecting industrial anomaly intelligence with enterprise decision processes.

2. LITERATURE REVIEW

2.1 Automated Industrial Exception Detection

The supplied references collectively establish a progression from conventional image processing toward increasingly intelligent industrial detection. Ishii, Mizuta, and Todo (1998) investigated real-time video image processing for foreign-substance detection, demonstrating the early feasibility of automated visual inspection. Chao, Kai, and Zhang (2020) extended this principle to tobacco foreign-body detection using machine vision. Chen et al. (2022) examined automatic detection and sorting of foreign matter in PVC powder, showing how detection can be combined with an operational sorting response.

These studies are relevant to autonomous exception management because they demonstrate the first stage of an exception lifecycle: recognizing abnormal states. However, detection remains fundamentally different from enterprise exception resolution. A vision model can determine that an object or defect exists, but it does not inherently determine whether production should stop, whether a batch should be quarantined, whether another production order should be prioritized, or whether a maintenance intervention is required.

2.2 Intelligent Anomaly and Fault Detection

Hou et al. (2022) addressed mechanical fault detection through adversarial auto-encoders and an early-warning mechanism. This is conceptually important because exception management requires identification of abnormal conditions before they become severe operational disruptions. Similarly, Wang and Cao (2022) reviewed crack modeling and detection methods for aero-engine disks, highlighting the complexity of identifying structural abnormalities.

Pan et al. (2023) examined foreign-matter defects in lithium-ion batteries using pilot manufacturing-line data. The relevance of this study extends beyond defect detection because production-line data can connect physical defects with manufacturing-process conditions. This supports the argument that exception management should combine multiple forms of evidence instead of relying on a single signal.

2.3 Multimodal and Three-Dimensional Perception

Lazarek and Pryczek (2018) reviewed point-cloud semantic-segmentation methods, establishing the importance of spatial understanding for complex industrial environments. Zhang, Zhong, and Li (2021) reviewed high-dynamic-range three-dimensional measurement using structured light, while Xu et al. (2021) explored object detection through the fusion of sparse point-cloud and image information.

These studies indicate that manufacturing exception detection increasingly involves multimodal information. For an SAP-integrated framework, this has an important implication: the exception-management layer should be capable of consuming structured enterprise data together with unstructured or sensor-derived evidence. For example, a production exception may simultaneously contain a machine anomaly, an image-based defect signal, a material-lot identifier, and a production-order dependency.

2.4 Foreign-Object Detection as an Exception-Management Foundation

Zhang, Yan, Zhu, et al. (2019) reviewed foreign-object detection for inductive power-transfer systems. Xu, Song, Zhang, et al. (2021) investigated foreign-object detection in power substations, and Zhang, Wang, Lv, et al. (2021) examined attention-based CNN detection in coal processing. Although these studies address different industrial environments, they share a common theoretical principle: operational safety and reliability depend on rapid identification of abnormal physical conditions.

The literature therefore provides strong evidence for the detection component of an intelligent exception architecture. However, the provided studies generally emphasize perception, detection, segmentation, or early warning rather than multi-agent enterprise-level decision orchestration. This represents the principal research gap addressed by this paper.

2.5 Research Gap and Theoretical Positioning

The literature demonstrates advanced capabilities for identifying physical and operational anomalies, but it provides limited conceptual integration between anomaly detection and enterprise resource planning workflows. A manufacturing exception becomes economically meaningful only when its operational consequences are interpreted within production and resource constraints.

The proposed framework addresses this gap by introducing a multi-agent reasoning architecture in which specialized agents transform heterogeneous signals into contextualized decisions. The Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing consequently extends the logic of automated detection from recognition of abnormality toward context-aware resolution of abnormality.

3. METHODOLOGY

3.1 Research Design

This study adopts a conceptual design-science-oriented methodology. The architecture is derived from systematic synthesis of the provided literature and mapped onto a hypothetical SAP S/4HANA manufacturing environment. The objective is not to claim empirical performance from an implementation that has not been experimentally tested. Instead, the methodology establishes a technically coherent framework that can subsequently be evaluated through simulation, historical production data, or controlled industrial deployment.

The proposed architecture contains six logical layers: data acquisition, exception detection, contextual interpretation, multi-agent reasoning, decision governance, and enterprise execution.

3.2 Layer 1: Manufacturing Data Acquisition

The first layer collects signals from manufacturing operations. These may include production-order status, material availability, equipment conditions, quality information, inspection results, machine-vision outputs, and sensor-derived observations.

The supplied literature supports the technical feasibility of generating such signals. For example, image-processing approaches have been used for foreign-substance identification (Ishii, Mizuta, Todo, 1998), while machine-vision methods have been applied to tobacco and polymer-processing environments (Chao, Kai, Zhang, 2020; Chen et al., 2022). Point-cloud and structured-light research further indicates the potential availability of three-dimensional spatial information (Lazarek, Pryczek, 2018; Zhang, Zhong, Li, J et al., 2021).

The framework converts these heterogeneous signals into standardized exception events containing an event identifier, timestamp, affected asset or order, observed anomaly, confidence level, and contextual metadata.

3.3 Layer 2: Exception Detection Agent

The Exception Detection Agent is responsible for identifying abnormal conditions. It does not make business decisions. Its purpose is to transform raw signals into structured exceptions.

For instance, if a vision system identifies a foreign object on a manufacturing line, the agent can create an event such as:

Exception Type: Foreign-object anomaly

Affected Process: Production operation

Confidence: High

Affected Batch: Identified production batch

Recommended Initial Status: Investigation required

The architecture can incorporate multiple detection methods. Machine vision may identify physical anomalies, while anomaly-detection models may identify unusual equipment behavior. Hou et al. (2022) provides conceptual support for early-warning anomaly detection, whereas Zhang et al. (2021) demonstrates the relevance of attention mechanisms for complex foreign-object detection.

3.4 Layer 3: Diagnostic and Context Agent

The Diagnostic Agent determines the likely cause and operational meaning of an exception. Unlike the detection agent, it uses contextual information.

For example, a machine anomaly occurring during a high-priority production order should be interpreted differently from the same anomaly occurring during scheduled maintenance. The Diagnostic Agent therefore correlates the exception with production order status, equipment identity, material information, previous events, and quality conditions.

This contextual approach reflects the broader lesson of the literature: industrial anomalies are highly dependent on the environment in which they occur. Wang and Cao (2022), for example, demonstrate the importance of systematic crack analysis rather than simple binary recognition.

3.5 Layer 4: Impact Assessment Agent

The Impact Assessment Agent estimates the consequences of an exception. Its reasoning can include production delay, material exposure, quality risk, downstream order disruption, and equipment availability.

A conceptual impact score can be represented as:

Impact Score = Severity × Operational Dependency × Time Criticality × Confidence Adjustment

This formulation is intentionally conceptual rather than a fixed mathematical law. It enables the framework to distinguish between technically similar exceptions with different business consequences.

For example, a low-severity anomaly on a non-critical production line may be scheduled for later inspection. Conversely, a high-confidence defect involving a critical production batch may require immediate containment.

3.6 Layer 5: Resolution Planning Agent

The Resolution Planning Agent generates possible corrective actions. Rather than selecting a single action immediately, the agent can construct alternatives.

For a material-related exception, possible actions may include reallocating available inventory, delaying the affected production order, or escalating procurement. For an equipment exception, alternatives may include inspection, maintenance intervention, production rescheduling, or controlled continuation.

The generative capability is useful because exception resolution often involves combinations of actions rather than a single predefined rule. However, generated recommendations must remain constrained by enterprise policies.

The Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing therefore treats generative reasoning as a proposal mechanism governed by structured business rules rather than as unrestricted autonomous behavior.

3.7 Layer 6: Governance and Execution Agent

The Governance Agent evaluates whether a proposed action can be executed autonomously. Three decision states are proposed:

Autonomous execution: high-confidence, low-risk actions that satisfy predefined policies.

Human approval: actions with significant operational, financial, or quality consequences.

Escalation: low-confidence or contradictory cases requiring specialist investigation.

This governance mechanism is essential because generative AI may generate plausible but incorrect recommendations. The framework consequently separates reasoning authority from transactional authority. The AI can reason broadly, but execution privileges remain bounded.

3.8 Multi-Agent Coordination

The agents communicate through a shared exception context. A simplified workflow is:

Detection → Diagnosis → Impact Assessment → Resolution Planning → Governance → SAP Execution/Escalation

Each agent contributes a specialized decision rather than independently attempting to solve the complete problem. This decomposition reduces cognitive and computational complexity and allows individual components to be evaluated independently.

3.9 Hypothetical Manufacturing Scenario

Consider a production line where a machine-vision system identifies a foreign object in a production batch. The Detection Agent creates the exception. The Diagnostic Agent determines that the event is associated with a specific production order and material lot. The Impact Assessment Agent identifies the affected downstream operations. The Resolution Planning Agent proposes batch containment and production rescheduling. The Governance Agent verifies whether the proposed actions satisfy predefined quality and authorization policies. If approved, the execution layer updates the appropriate enterprise workflow and notifies responsible personnel.

The architecture therefore transforms a visual anomaly into an enterprise-level response. This represents a major conceptual difference from isolated inspection systems.

4. RESULTS

The conceptual evaluation of the proposed architecture identifies four principal findings. First, manufacturing exception management can be decomposed into specialized cognitive functions rather than being treated as a single AI prediction problem. Detection, diagnosis, impact assessment, resolution generation, and governance require different forms of reasoning and evidence. This decomposition creates a more controllable architecture than a monolithic generative AI system.

Second, the literature indicates that automated industrial detection is sufficiently diverse to justify a multimodal exception layer. Machine vision, real-time image processing, anomaly detection, point-cloud analysis, and structured-light measurement represent different mechanisms for observing manufacturing conditions (Ishii, Mizuta, Todo, 1998; Lazarek, Pryczek, 2018; Zhang, Zhong, Li, J et al., 2021). The proposed framework consequently does not assume that all exceptions originate from ERP transactions. Physical-world observations can become enterprise events.

Third, multi-agent reasoning provides a conceptual mechanism for connecting anomaly detection with manufacturing consequences. The key finding is that detection confidence alone is insufficient for determining action. A high-confidence foreign-object detection, for example, does not automatically specify whether a batch should be quarantined, production should stop, or another order should be prioritized. Contextual reasoning is therefore necessary. The Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing addresses this problem by distributing contextual interpretation across specialized agents.

Fourth, controlled autonomy emerges as more appropriate than unrestricted automation. The architecture establishes a separation between recommendation and execution, allowing low-risk exceptions to be handled automatically while routing ambiguous or high-impact cases to human decision-makers. This governance principle is particularly important because the provided literature demonstrates the effectiveness of detection technologies but does not establish that detected anomalies can safely trigger unrestricted enterprise transactions.

The conceptual findings also indicate several potential operational benefits: shorter exception-response cycles, improved consistency of diagnosis, reduction of repetitive manual investigation, and stronger traceability of decisions. However, these benefits remain architectural hypotheses until validated through production datasets and measurable experiments.

5. DISCUSSION

The proposed framework contributes to the literature by shifting the emphasis from automated detection toward autonomous exception management. Existing studies primarily demonstrate how abnormal physical conditions can be detected, classified, segmented, or predicted. Chao, Kai, and Zhang (2020), Chen et al. (2022), Zhang et al. (2021), and Xu et al. (2021), for example, establish different approaches to foreign-object detection. Hou et al. (2022) further demonstrates the value of early-warning mechanisms for mechanical faults. These contributions provide the perceptual foundation for an enterprise exception-management architecture, but they do not by themselves resolve the downstream business problem.

The theoretical implication is that manufacturing intelligence should be viewed as a hierarchical process. At the lowest level, sensors and inspection technologies identify physical states. At an intermediate level, analytical agents interpret patterns and estimate causes. At the enterprise level, decision agents evaluate operational impact and coordinate responses. The Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing integrates these levels through agent specialization.

A significant practical implication is the potential reduction in the gap between shop-floor intelligence and enterprise processes. An anomaly identified by a machine-vision system has limited value if it remains isolated from production planning and order management. Conversely, an ERP alert without physical context may generate excessive or poorly prioritized exceptions. The proposed architecture attempts to establish a closed reasoning loop between physical observation and enterprise response.

Nevertheless, several trade-offs must be considered. Greater autonomy can reduce response time but increase the consequences of incorrect decisions. Generative reasoning can handle complex combinations of information but introduces risks associated with unsupported or erroneous recommendations. Multi-agent architectures improve modularity but increase integration and coordination complexity. These trade-offs imply that autonomy should be proportional to decision risk.

The framework also faces explainability challenges. An enterprise user must be able to understand why an exception was classified as critical and why a particular action was proposed. The architecture therefore requires persistent exception histories, evidence references, confidence values, and agent-level decision records. Without such traceability, autonomous exception management could become operationally opaque.

Data quality represents another limitation. The architecture assumes that production, inspection, and equipment signals can be reliably associated with enterprise objects. Incorrect identifiers, incomplete sensor information, inconsistent timestamps, or missing contextual data could propagate erroneous conclusions between agents. The importance of reliable manufacturing observations is indirectly reinforced by the supplied literature on complex detection environments, including three-dimensional measurement and point-cloud segmentation (Lazarek, Pryczek, 2018; Zhang, Zhong, Li, J et al., 2021).

Finally, the present study is conceptual. It does not report a live SAP implementation, controlled benchmark, or measured reduction in exception-resolution time. Therefore, claims concerning performance should be interpreted as expected architectural benefits rather than empirical results. Future research should evaluate the framework using historical manufacturing exceptions and compare single-agent, rule-based, and multi-agent approaches according to detection accuracy, resolution time, false escalation rate, human approval frequency, and transactional reliability.

6. CONCLUSION

This study proposed a Multi-Agent Generative AI Framework for Autonomous Exception Management in SAP S/4HANA Manufacturing that connects automated industrial anomaly detection with enterprise-level reasoning and controlled execution. The framework is organized around specialized agents responsible for detection, diagnosis, impact assessment, resolution planning, governance, and execution.

The provided literature demonstrates strong foundations for automated manufacturing perception, including machine vision, anomaly detection, foreign-object recognition, point-cloud analysis, and three-dimensional measurement. However, these capabilities do not inherently provide coordinated enterprise decision-making. The proposed framework addresses this gap by converting heterogeneous manufacturing observations into contextualized exceptions and routing them through a governed multi-agent decision process.

The principal contribution is therefore architectural rather than empirical: the study establishes a model in which generative AI is used for contextual reasoning while enterprise transaction authority remains governed by confidence, risk, authorization, and human escalation mechanisms. Such a structure provides a pathway toward more responsive and resilient manufacturing operations without assuming that AI-generated decisions should automatically control enterprise processes.

Future research should implement the framework against real or simulated SAP manufacturing datasets, establish quantitative evaluation metrics, investigate agent coordination strategies, and examine the reliability of autonomous decisions under ambiguous or conflicting manufacturing conditions. Empirical validation would determine whether the conceptual advantages identified in this study translate into measurable improvements in manufacturing exception management.

REFERENCES

1. Chao M, Kai C, Zhang Z. Research on tobacco foreign body detection device based on machine vision. Trans Inst Meas Control 2020; 42(15):2857–2871.

2. Chen S, Jang J, Chang Y, et al. An automatic foreign matter detection and sorting system for PVC Powder. Appl Sci 2022;12(12):6276.

3. Ho M, Obbard E, Burr P A, et al. A review on the development of nuclear power reactors. Energy Procedia 2019;160:459–466.

4. Hou J, Ma B, Liang L, et al. An early warning method for mechanical fault detection based on adversarial auto-encoders. J Adv Manuf Sci Technol 2022;2(2):2022006.

5. Ishii A, Mizuta T, Todo S. Detection of foreign substances mixed in a plastic bottle of medicinal solution using real-time video image processing. In: 14th International Conference on Pattern Recognition. 1998. p.1646–1650.

6. Lazarek J, Pryczek M. A review on point cloud semantic segmentation methods. J Appl Comput Sci 2018; 1507–0360.

7. Pan Y, Kong X, Yuan Y, et al. Detecting the foreign matter defect in lithium-ion batteries based on battery pilot manufacturing line data analyses. Energy 2023;262:125502.

8. Tan J, Zhang Z, Zheng H, et al. Corrosion fatigue model of austenitic stainless steels used in pressurized water reactor nuclear power plants. J

Nucl Mater 2020;541:152407.

9. Wang D, Cao H. A comprehensive review on crack modeling and detection methods of aero-engine disks. J Adv Manuf Sci Technol 2022;2(3):2022012.

10. Xu L, Song Y, Zhang W, et al. An efficient foreign objects detection network for power substation. Image Vis Comput 2021;109:104159.

11. Xu X, Zhang L, Yang J, et al. Object detection based on fusion of sparse point cloud and image information. IEEE Trans Instrum Meas 2021;70:1–2.

12. Zhang K, Wang W, Lv Z, et al. Computer vision detection of foreign objects in coal processing using attention CNN. Eng Appl Artif Intell 2021;102:104242.

13. Zhang P, Zhong K, Li Z, J et al. High dynamic range 3D measurement based on structured light: A review. J Adv Manuf Sci Technol 2021;1(2):2021004.

14. Zhang Y, Yan Z, Zhu J, et al. A review of foreign object detection (FOD) for inductive power transfer systems. eTransportation 2019;1:100002.

15. M. Kalal, "Predictive Production Planning in Advanced Manufacturing Using SAP PP and Analytics," 2026 5th International Conference on Sentiment Analysis and Deep Learning (ICSADL), Birendranagar, Nepal, 2026, pp. 1363-1370, doi: 10.1109/ICSADL67539.2026.11452022.