American Journal of Intelligent Systems 2026, 15(1): 1-6
DOI: 10.5923/j.ajis.20261501.01

AI-Driven Discovery Bot for Manufacturing: Empowering Product and Plant Engineering through Historical Issue Resolution Guidance

Shuvdeep Bhattacharya

IT Director - Seating Engineering, Lear Corporation, 21577 Telegraph Road, Southfield MI 48033

Abstract This paper presents an artificial intelligence–driven Discovery Bot designed to improve manufacturing issue resolution, engineering knowledge reuse, and risk management. The system consolidates historical issue records, engineering change information, quality reports, production data, and collaboration records to support faster root cause analysis and reuse of proven corrective actions. The proposed framework combines natural language processing, machine learning, retrieval-augmented generation, and knowledge graphs to retrieve contextually relevant prior resolutions, identify recurring failure patterns, and support dynamic Failure Modes and Effects Analysis (FMEA). The architecture integrates with product lifecycle management, manufacturing execution, enterprise resource planning, quality, and collaboration systems through modular services and secure APIs. To address practical deployment requirements, the paper distinguishes currently deployed capabilities from planned extensions, defines evaluation metrics, and describes validation against expert-defined baselines. Case-study evidence across manufacturing contexts indicates measurable benefits, including shorter issue-resolution cycles and reduced downtime in selected use cases. The results suggest that an AI-enabled Discovery Bot can convert fragmented engineering knowledge into a scalable decision-support capability, enabling organizations to move from reactive troubleshooting toward predictive quality, reliability, and continuous improvement.

Keywords Artificial intelligence, Manufacturing knowledge management, FMEA, DFMEA, PFMEA, Knowledge graphs, Retrieval-augmented generation, Natural language processing, Root cause analysis, PLM, MES, ERP

Managerial Relevance Statement

Manufacturing organizations often lose time because relevant engineering knowledge is distributed across PLM, MES, ERP, quality systems, email, and local expert memory. This paper provides a practical framework for converting that fragmented knowledge into a governed decision-support capability. The Discovery Bot helps engineers retrieve similar historical issues, understand previous corrective actions, and update FMEA knowledge as new production and field data emerge.

For managers, the contribution is threefold. First, the framework reduces dependence on undocumented tribal knowledge by making historical lessons searchable and reusable. Second, AI-enhanced FMEA supports a shift from static risk documentation to continuously updated risk assessment. Third, the modular API-based architecture allows organizations to adopt the capability incrementally without replacing existing enterprise systems.

* Corresponding author:
shuvdeep@umich.edu (Shuvdeep Bhattacharya)
Received: Jun. 14, 2026; Accepted: Jul. 10, 2026; Published: Jul. 15, 2026
Published online at http://journal.sapub.org/ajis

1. Introduction

Manufacturing organizations are increasingly leveraging artificial intelligence to tackle complex engineering challenges and improve operational performance. Traditional issue resolution methods—reliant on tribal knowledge and manual searches—are often inefficient and prone to error, leading to delayed decision-making and obscured critical insights.

The Discovery Bot transforms this process by automating the retrieval and interpretation of historical data, enabling faster and more informed decisions. Its AI-driven approach to Failure Modes and Effects Analysis (FMEA) enhances risk assessment by identifying potential failure points early in the lifecycle. The system dynamically learns from past issues and corrective actions to generate predictive insights.

By integrating natural language processing (NLP), machine learning (ML), and domain-specific knowledge graphs, the bot streamlines root cause analysis and supports continuous improvement. It empowers cross-functional teams with contextual guidance, shifting engineering from reactive to predictive operations, thereby reducing downtime and improving product quality.

2. Literature Review

A. Manufacturing Knowledge Fragmentation

Manufacturing knowledge is often embedded in structured records, unstructured reports, engineering documents, and informal collaboration channels. Prior research on smart manufacturing knowledge management emphasizes that disconnected technical information limits reuse and slows operational decision-making [1], [2]. In issue-resolution contexts, the value of historical knowledge depends not only on storage but also on retrieval, contextualization, and traceability to affected components, processes, and corrective actions.

B. AI, NLP, Knowledge Graphs, and RAG in Manufacturing

Natural language processing enables systems to interpret engineering descriptions, maintenance notes, quality reports, and operator queries. Large language models have recently been explored for factory knowledge sharing, where retrieval from manuals and issue reports can support faster access to information while still requiring human review [8]. Retrieval-augmented generation (RAG) improves factual grounding by retrieving relevant evidence before generating a response [10]. Knowledge graphs further strengthen retrieval by representing relationships among parts, symptoms, failure modes, processes, suppliers, and corrective actions [7], [12], [13].

C. AI-Enhanced FMEA and Novelty Relative to Prior Work

FMEA is a foundational method for identifying failure modes, causes, effects, controls, and risk priorities. However, conventional FMEA can be manual, time-consuming, and difficult to keep current. Prior studies have proposed data-driven FMEA for maintenance planning [3], ontology-based reasoning over FMEA tables [4], and LLM-supported FMEA generation [5], [6].

Reference [5] is particularly relevant because it integrates LLMs into FMEA using data collection, preprocessing, risk identification, and human-in-the-loop validation. The present work differs in scope and contribution. Rather than focusing primarily on generating FMEA content from selected datasets, this paper embeds AI-enhanced FMEA inside an enterprise Discovery Bot that supports historical issue-resolution guidance. The proposed system links FMEA elements to PLM parts, MES processes, ERP supplier data, quality records, and prior corrective actions. Its novelty lies in combining enterprise retrieval, knowledge-graph traceability, AI-assisted root cause guidance, and continuous FMEA update in one deployment-oriented framework.

Manufacturing organizations operate in data-rich but knowledge-fragmented environments. Engineering changes, quality defects, maintenance logs, supplier issues, process deviations, and corrective actions are often stored in different systems and described using inconsistent terminology. As a result, engineers may spend substantial time searching for prior resolutions or relying on experienced personnel who may not be available.

Traditional issue-resolution workflows are therefore vulnerable to delay, duplicated investigation, repeated failure modes, and loss of institutional knowledge. These challenges are especially visible in complex product and plant engineering environments, where product design, tooling, process parameters, suppliers, and production conditions interact.

The Discovery Bot addresses this problem by using AI to retrieve, interpret, and connect historical issue data. It supports engineers through natural-language queries, contextual retrieval, knowledge-graph relationships, and AI-enhanced FMEA updates. Instead of treating FMEA as a static document, the framework treats it as a living risk model that evolves with operational evidence.

The paper makes four contributions:

1. It proposes a modular AI architecture for historical issue-resolution guidance across enterprise manufacturing systems.
2. It defines a five-stage methodology for data acquisition, preprocessing, retrieval, validation, and continuous learning.
3. It adds an evaluation approach with explicit baselines, metrics, and validation against expert-defined outputs.
4. It clarifies the novelty of the Discovery Bot relative to prior LLM-FMEA work by extending the focus from automated FMEA generation to enterprise-scale issue-resolution guidance and cross-system traceability.

3. Methodology

A. Study Design and Evaluation Scope

The study uses a multi-case, pre/post evaluation design. Each deployment context compares the Discovery Bot-assisted workflow against the incumbent baseline workflow, defined as manual search across enterprise systems, consultation of prior reports, and escalation to experienced subject-matter experts. The evaluation focuses on three outcome categories: issue-resolution time, downtime, and yield-related quality indicators.

Because the deployment contexts involve proprietary manufacturing operations, plant names, product identifiers, and absolute production volumes are anonymized. Results are therefore reported as normalized percentage improvements were supported by available operational records.

B. Data Sources

The Discovery Bot ingests structured and unstructured data from the following sources:

- **PLM systems:** Part metadata, engineering changes, design specifications, and product history.
- **MES systems:** Production events, machine logs, process deviations, and downtime records.
- **ERP systems:** Supplier, procurement, inventory, and part-availability information.
- **Quality systems:** FMEA records, 8D reports, audit findings, control plans, defect logs, and corrective actions.

- **Collaboration records:** Emails, issue summaries, meeting notes, and informal problem-solving discussions where access is authorized.
- **Data are normalized** through part-number mapping, metadata enrichment, duplicate detection, terminology standardization, and access-control tagging.

C. System Architecture and Implementation

The Discovery Bot follows a modular architecture consisting of five layers:

- **User interaction layer:** natural-language interface for engineers, quality teams, and plant users.
- **NLP and intent layer:** query parsing, entity extraction, intent classification, and terminology normalization.
- **Retrieval and knowledge layer:** hybrid retrieval across indexed documents, structured records, and knowledge-graph relationships.
- **Recommendation and FMEA layer:** retrieval of similar issues, corrective-action ranking, failure-mode mapping, and risk-score update support.
- **Integration and governance layer:** REST APIs, role-based access control, audit logging, and enterprise-system connectors.

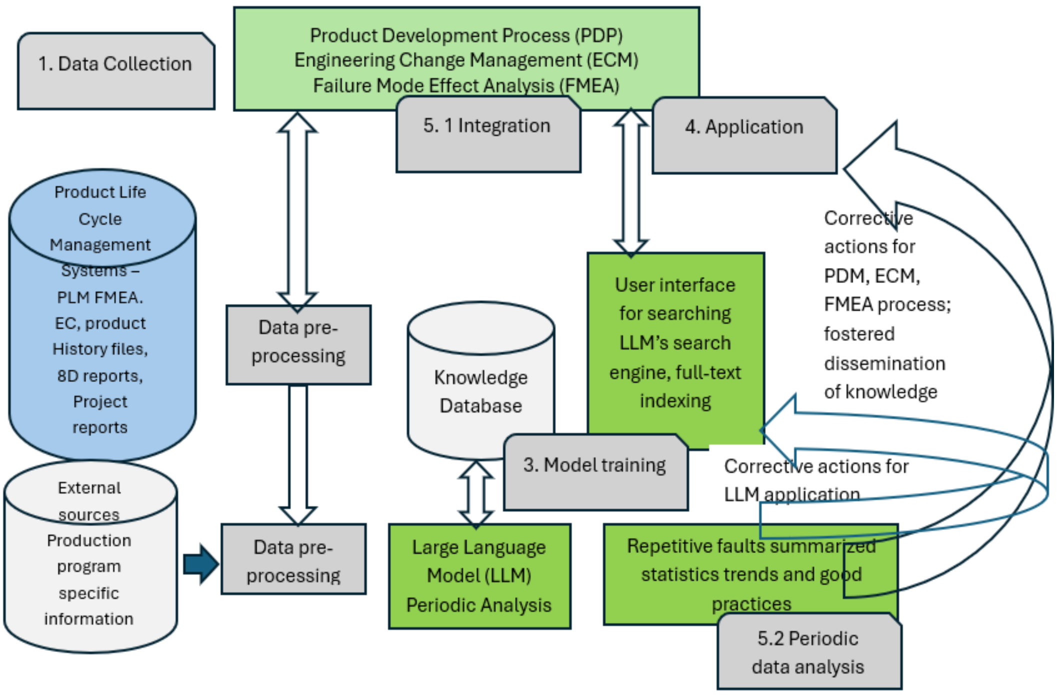

Figure 1. Information system model for LLM application in FMEA

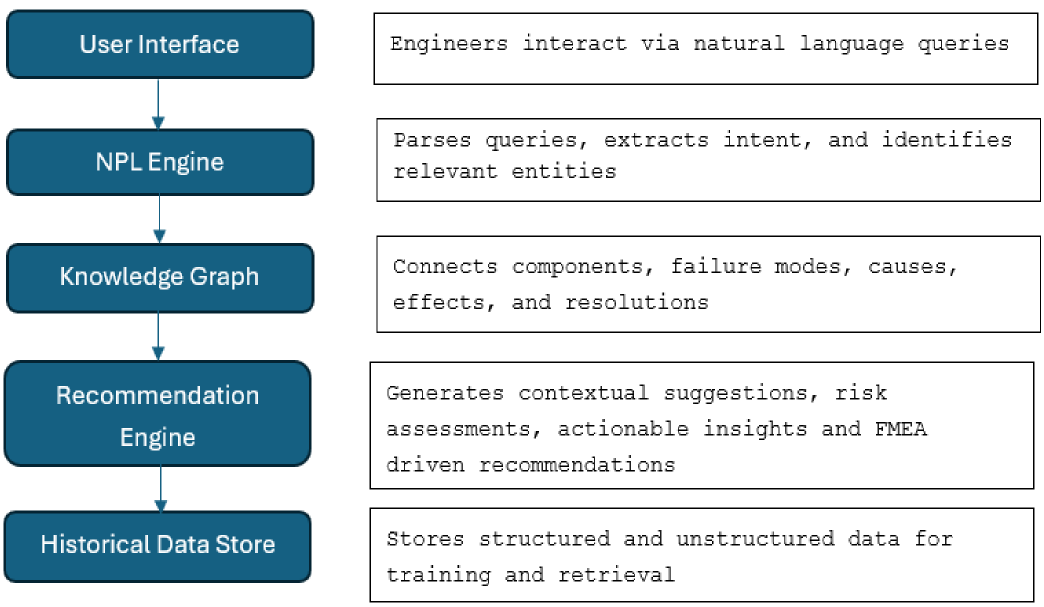

Figure 2. Discovery Bot reference architecture

The figure should show data acquisition from PLM, MES, ERP, quality systems, and collaboration sources; preprocessing and indexing; knowledge-graph construction; LLM/RAG retrieval; expert validation; and continuous FMEA update.

The figure should show the user interface, NLP engine, retrieval services, vector index, knowledge graph, recommendation engine, FMEA module, API gateway, security services, and enterprise-system connectors.

D. Five-Stage Framework

The methodology operates through five stages:

- **Data acquisition and preprocessing:** collect authorized records, clean text, map part/process metadata, and prepare vector and graph indexes.
- **Contextual retrieval:** retrieve similar historical issues using semantic search, metadata filters, and graph relationships.
- **AI-assisted reasoning:** summarize prior symptoms, suspected causes, implemented corrective actions, and linked FMEA entries.
- **Expert validation:** engineers review AI outputs before corrective actions or FMEA changes are accepted.
- **Continuous learning:** validated resolutions are added back to the knowledge base, improving future retrieval and risk assessment.

E. Deployed Versus Planned Components

Table 1. Illustrate stages of deployment

Component	Current status	Notes
Natural-language issue search	Deployed	Used to query prior issue records and corrective actions.
PLM/MES/ERP/quality connectors	Deployed in modular form	Integration depends on local system availability and permissions.
Knowledge graph for parts, failures, causes, and actions	Deployed	Used for traceability and contextual retrieval.
AI-assisted FMEA update recommendations	Deployed with human review	Recommendations require SME approval.
Role-based access and audit logging	Deployed	Used to govern sensitive engineering and quality data.
Real-time IoT sensor streaming	Planned	Future extension for early warning and live risk updates.
Multilingual and multimodal input	Planned	Future support for voice, images, and CAD annotations.
Closed-loop model retraining from action outcomes	Planned	Requires longitudinal validation.
Digital twin integration	Planned	Future capability for simulation-driven FMEA.

F. Metrics and Measurement

Issue-resolution time is measured as elapsed time from issue intake or ticket creation to documented containment or corrective-action approval. The baseline is the historical manual workflow for comparable issue types.

Downtime is measured using production or maintenance records that capture unavailable equipment or interrupted production time. Downtime reduction is calculated as the percentage decrease relative to the baseline period or comparable historical cases.

Yield is measured using first-pass yield, scrap, rework, or defect escape indicators where such records are available. Because yield definitions differ by process and industry, yield is reported only when the underlying quality system provides consistent measurement.

G. LLM Extraction Validation

To validate the AI extraction layer, LLM outputs are compared against expert-defined baselines. Experts identify the expected component, symptom, failure mode, likely cause, prior corrective action, and affected FMEA entry. The Discovery Bot output is then evaluated using:

- **Exact or semantic agreement:** whether extracted failure modes match expert baselines in meaning.
- **Precision:** proportion of extracted items judged correct by experts.
- **Recall:** proportion of expert-identified items retrieved by the system.
- **F1 score:** harmonic mean of precision and recall.
- **Hallucination rate:** proportion of generated claims not supported by retrieved evidence.
- **Actionability rating:** expert judgment of whether the recommendation is usable for investigation.

Table 2. LLM Extraction Validation Protocol

Validation element	Definition	Reviewer
Component match	Correct part/process identified	Product or process SME
Failure-mode match	Extracted failure mode matches expert baseline	Quality/FMEA SME
Cause mapping	Suggested cause is supported by retrieved evidence	Engineering SME
Corrective-action relevance	Suggested action matches prior successful resolution	Plant/quality SME
FMEA linkage	Correct FMEA row or risk category identified	FMEA owner
Unsupported claim check	Any generated claim lacking retrieved evidence is flagged	Reviewer panel

H. End-to-End Walkthrough: Weld-Failure Issue

A plant engineer enters a natural-language issue: "Recurring weld failure at seat-frame bracket after changeover." The Discovery Bot extracts key entities: weld failure, seat-frame bracket, changeover, affected process step, and potential part/process identifiers. It retrieves prior quality reports, related engineering changes, similar 8D records, and FMEA entries associated with weld strength, fixture alignment, and process-parameter drift.

The recommendation engine ranks prior cases by semantic similarity, shared component, process step, and corrective-action outcome. The engineer receives a summary showing previous containment actions, suspected causes, inspection points, and corrective actions. The system also identifies whether the failure mode already exists in PFMEA and whether severity, occurrence, or detection assumptions may need review.

After the engineer validates the recommendation and records the final corrective action, the outcome is added to the knowledge base. If the issue indicates a recurring or newly emerging risk, the FMEA owner receives a proposed update for review. In this way, the issue-resolution workflow and FMEA process remain connected.

4. Results

A. Cross-Industry Evaluation Contexts

Table 3. Impact of Discovery Bot Across Evaluation Contexts

Context	Primary issue type	Main data sources	Measured outcome
Automotive	Weld failure	Quality reports, MES records, FMEA, engineering changes	Downtime reduction
Aerospace	Assembly or process deviation	Quality records, audit findings, engineering history	Resolution-time reduction
Electronics	Defect recurrence	Test logs, quality reports, corrective actions	Yield or rework indicator where available
Heavy machinery	Hydraulic diagnostics	Maintenance logs, quality reports, part history	Downtime reduction
Consumer goods	Packaging optimization	Line reports, production data, issue logs	Downtime reduction

B. Quantitative Performance Outcomes

Table 4. Downtime Reduction by Use Case

Use case	Downtime reduction (%)
Weld failure	40
Hydraulic diagnostics	50
Packaging optimization	30

Across evaluated deployments, the Discovery Bot reduced issue-resolution time by approximately 30–50% where comparable baseline records were available. These improvements are attributed to faster retrieval of similar historical issues, reuse of validated corrective actions, and improved traceability between symptoms, components, causes, and FMEA entries. Yield improvements were observed qualitatively in selected contexts, but yield is not reported as a uniform percentage because measurement definitions varied across systems.

C. Validation Findings

The validation process showed that expert review is essential for safe use of AI-generated engineering recommendations. The Discovery Bot is therefore positioned as decision support rather than autonomous decision-making. Outputs are accepted only when retrieved evidence supports the recommended failure mode, cause, or corrective action. Unsupported claims are rejected or returned for further investigation.

5. Discussion

The results support the practical value of combining historical issue retrieval with AI-enhanced FMEA. The primary benefit is not simply faster search; it is the transformation of fragmented records into contextual engineering guidance. By linking symptoms, components, root causes, corrective actions, and FMEA entries, the Discovery Bot helps engineers identify whether an issue is new, recurring, or related to a previously mitigated risk.

The framework also addresses a limitation of static FMEA practice. Traditional FMEA documents may become outdated when new field or production evidence is not systematically incorporated. The proposed approach creates a feedback loop in which validated issue-resolution outcomes can trigger FMEA review. This improves organizational learning while preserving human accountability.

The system's value depends on data quality, terminology consistency, access governance, and expert validation. AI recommendations should not be treated as final engineering decisions. Instead, they should accelerate evidence gathering, highlight relevant prior experience, and support structured review by qualified personnel.

6. Conclusions

The AI-driven Discovery Bot provides a practical framework for improving manufacturing issue resolution and FMEA currency. The key conclusions are:

- Historical knowledge can be operationalized. The system converts fragmented issue records, quality reports, engineering changes, and production data into searchable guidance.
- AI-enhanced FMEA improves risk visibility. Failure modes, causes, corrective actions, and risk assumptions can be reviewed as new evidence emerges.
- Measured outcomes are encouraging. Reported cases show 30–50% reduction in issue-resolution time and downtime reductions of 40%, 50%, and 30% in selected use cases.
- Human validation remains essential. The Discovery Bot supports engineers but does not replace expert judgment, especially for safety, quality, and design decisions.
- Enterprise integration is central to
- scalability. Secure APIs, access control, knowledge graphs, and modular services allow deployment across plants and product lines.

7. Limitations and Future Directions

The study has several limitations. First, deployment contexts are anonymized, and absolute counts are not disclosed. Second, yield metrics are not uniformly reported because definitions differ across systems and industries. Third, the validation approach relies on expert-defined baselines, which improves safety but requires SME availability. Fourth, real-time IoT integration, multilingual and multimodal input, closed-loop retraining, and digital twin linkage remain planned extensions.

Future work should expand validation sample sizes, report precision/recall/F1 values for LLM extraction, test the system longitudinally in live production settings, and evaluate whether closed-loop FMEA updates reduce recurrence over time. Additional research should also examine governance, privacy, and explainability requirements for AI-enabled manufacturing decision support.

ACKNOWLEDGEMENT

The author thanks colleagues and industry partners who supported the development, evaluation, and refinement of the Discovery Bot across manufacturing environments.

REFERENCES

[1] S. C. Feng, W. Z. Bernstein, J. T. Hedberg, and A. B. Feeney, “Toward knowledge management for smart manufacturing,” *Journal of Computing and Information Science in Engineering*, vol. 17, no. 3, 2017.

[2] C. Gröger, H. Schwarz, and B. Mitschang, “The manufacturing knowledge repository,” *ICEIS*, pp. 39–51, 2014.

[3] M.-A. Filz, J. E. B. Langner, C. Herrmann, and S. Thiede, “Data-driven failure mode and effect analysis to enhance maintenance planning,” *Computers in Industry*, vol. 129, 103451, 2021.

[4] M. Hodkiewicz, J. W. Klüwer, C. Woods, T. Smoker, and E. Low, “An ontology for reasoning over engineering textual data stored in FMEA spreadsheet tables,” *Computers in Industry*, vol. 131, 103496, 2021.

[5] I. El Hassani, T. Masrour, N. Kourouma, D. Motte, and J. Tavčar, “Integrating large language models for improved failure mode and effects analysis (FMEA): a framework and case study,” *Proceedings of the Design Society*, vol. 4, pp. 2019–2028, 2024.

[6] I. El Hassani, T. Masrour, N. Kourouma, and J. Tavčar, “AI-driven FMEA: integration of large language models for faster and more accurate risk analysis,” *Design Science*, vol. 11, e10, 2025.

[7] L. Bahr et al., “Knowledge graph enhanced retrieval-augmented generation for failure mode and effects analysis,” *Journal of Industrial Information Integration*, vol. 45, 100807, 2025.

[8] S. Kernan Freire et al., “Knowledge sharing in manufacturing using LLM-powered tools: user study and model benchmarking,” *Frontiers in Artificial Intelligence*, vol. 7, 1293084, 2024.

[9] C. Ouerghemmi and M. Ertz, “Integrating large language models into digital manufacturing: a systematic review and research agenda,” *Computers*, vol. 14, no. 8, 318, 2025.

[10] P. Lewis et al., “Retrieval-augmented generation for knowledge -intensive NLP tasks,” *NeurIPS*, vol. 33, pp. 9459–9474, 2020.

[11] Z. Wu, W. Liu, and W. Nie, “Literature review and prospect of the development and application of FMEA in manufacturing industry,” *International Journal of Advanced Manufacturing Technology*, vol. 112, pp. 1409–1436, 2021.

[12] Y. Wan et al., “Making knowledge graphs work for smart manufacturing: research topics, applications and prospects,” *Journal of Manufacturing Systems*, vol. 76, pp. 103–132, 2024.

[13] Y. Li and B. Starly, “Building a knowledge graph to enrich ChatGPT responses in manufacturing service discovery,” *Journal of Industrial Information Integration*, vol. 40, 100612, 2024.

[14] C. Zhang et al., “A survey on potentials, pathways and challenges of large language models in new-generation intelligent manufacturing,” *Robotics and Computer-Integrated Manufacturing*, vol. 92, 102883, 2025.

[15] AIAG and VDA, *Failure Mode and Effects Analysis—FMEA Handbook*, Automotive Industry Action Group and Verband der Automobilindustrie, 2019.