

# Bridging Expert Knowledge and AI: A Semantic Architecture for Manufacturing Knowledge Management using Knowledge Graphs and Large Language Models

Camilla Hemmer<sup>1,\*</sup>, Annariina Komljenovic<sup>1</sup> and Jaroslaw Warzecha<sup>1</sup>

<sup>1</sup>valantic GmbH, Ainmillerstrasse 22, 80801 München, Germany

## Abstract

Manufacturing companies face a critical knowledge management crisis where decades of operational expertise remains trapped in expert minds, leading to extended onboarding periods, inefficient decision-making and knowledge loss through retirement. This paper presents the development of a semantic AI assistant system that combines Knowledge Graphs with Large Language Models to transform how manufacturing engineers access and utilize institutional knowledge. The approach addresses the integration of scattered data across incompatible systems through a three-stage pipeline: multi-format document ingestion, knowledge graph construction, and LLM-enhanced natural language querying. Our implementation demonstrates how semantic technologies can address real-world industrial challenges, reducing information retrieval time and creating pathways for knowledge preservation at enterprise scale.

## Keywords

Knowledge Graphs, Large Language Models, Manufacturing Knowledge Management, Semantic Web Technologies, Industrial AI, Enterprise Knowledge Systems, Natural Language Processing, Manufacturing Automation

## 1. Introduction

Modern manufacturing organizations face an unprecedented knowledge management crisis. As industrial processes become increasingly complex and specialized, critical operational knowledge becomes concentrated in the minds of expert engineers and technicians[1].

The scope of this challenge is exemplified by our case study organization: a manufacturing company where new engineering hires require over two years of training to become productive contributors often due to the fragmented and inaccessible nature of institutional knowledge. Critical information resides across incompatible systems including Customer Relationship Management (CRM) platforms, legacy Lotus Notes databases, scattered document repositories, and informal email communications.

This paper presents a novel approach that bridges this semantic gap through the integration of Knowledge Graphs (KGs) and Large Language Models (LLMs)[2]. Our architecture transforms manufacturing documentation into a semantic knowledge base, enabling natural language access to technical expertise and institutional knowledge. The contribution of this work extends beyond technical implementation to demonstrate practical deployment challenges and business impact measurement in real-world manufacturing environments. This industrial focus aligns with recent efforts in the ISWC community to bridge the gap between academic research and practical enterprise deployment[3, 4].

## 2. Industrial Automation Knowledge Graph: Beyond Simple Taxonomies

Our system is comprised of a three-stage pipeline that transforms heterogeneous documentation into a unified, accessible knowledge resource. This approach builds upon established manufacturing knowl-

*ISWC 2025 Companion Volume, November 2–6, 2025, Nara, Japan*

\*Corresponding author.

✉ camilla.hemmer@valantic.com (C. Hemmer); annariina.komljenovic@valantic.com (A. Komljenovic); jaroslaw.warzecha@valantic.com (J. Warzecha)

ORCID 0009-0000-6317-7143 (C. Hemmer)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

edge graph methodologies while addressing the specific challenges of multi-format enterprise document processing.

## 2.1. Document Ingestion

The pipeline processes diverse content formats prevalent in manufacturing, including PDFs, HTML, PowerPoint, and structured data exports. The export formats are raw chunked text and accompanying image formats that are linked to the corresponding text:

- **PDFs:** Extracted while maintaining formatting context.
- **HTML:** Analyzed for structural hierarchies.
- **PowerPoint:** Extracts both text and embedded diagrams critical to specifications.

## 2.2. Knowledge Graph Construction

In the first iteration of the pipeline, a schema is inferred from example chunks of the text. The ingested documents are then extracted based on the schema inferred and uploaded into the KG. This approach leverages recent advances in LLM-powered knowledge graph construction[5, 6], demonstrating enterprise-grade accuracy in domain-specific extraction tasks.

## 2.3. Knowledge Extraction Results

The schema creation resulted in 11 entity types with 38 unique relationship types. The inferred schema exhibits a three-tier hierarchy: core entities (Equipment, Document), operational entities (Specification, Material, Variant, Position), and parameter entities (SpeedRange, LoadRequirement, VentilationRequirement, ControlMechanism, OperatingCondition). Manufacturing engineers validated both the extracted schema and entity relationships, confirming that the resulting ontology is a valid approach to structuring the available specification documents. However, additional analysis remains necessary to expand the schema and examine the ontological hierarchies for this specific domain.

The knowledge graph construction phase extracted 982 entities from the manufacturing documentation corpus. The extraction process revealed varying completion rates across different entity attributes within manufacturing knowledge domains:

**Table 1**  
Knowledge Graph Entity Extraction Results

Entity Type	Core Attributes	Completion Rate
Equipment	name/type/purpose/model	100%/100%/100%/74%
Specification	name/value/unit/constraint	100%/98%/96%/65%
Position	name/location/relative_location	100%/99%/80%
Material	name/type/properties	100%/100%/94%
Document	name/type/date/link/version	100%/100%/52%/41%/49%
Variant	name/type/priority/selection_criteria	100%/100%/34%/93%
ControlMechanism	name/type/automation_required	100%/100%/96%
OperatingCondition	name/environment/temperature/humidity/pressure	100%/100%/12%/6%/24%
LoadRequirement	name/load_scenario/load_value/load_unit	100%/100%/78%/78%
SpeedRange	name/unit/min_speed/max_speed	100%/100%/100%/86%

Table 1 reveals distinct patterns in attribute completion rates across entity types. Core identifiers (names and types) consistently achieved high extraction rates, while specialized attributes such as environmental conditions (temperature, humidity, pressure) and document metadata (dates, links, versions) showed significantly lower completion rates.

## 2.4. AI Assistant Chatbot with Knowledge Graph Querying

We provide a frontend for the user with a chat interface that communicates with a Claude-3.5 instance in the backend. The LLM instance can write Cypher queries to query the Knowledge Graph for specific information. The answer then verbalizes this information and refers to the section of the document that is attached to this entry in the KG. Refer to Appendix A for an architectural overview.

## 3. LLM-Enhanced Knowledge Retrieval: Making Experts Accessible

The integration of Large Language Models transforms knowledge graph access from a technical query interface into a natural conversation platform that mimics expert consultation[7, 8]. This transformation addresses the adoption barriers that limited previous knowledge management approaches.

Natural language query translation represents a critical component of the LLM integration[9]. The system processes queries expressed in technical terminology and conversational language, identifying relevant entities and relationships within the knowledge graph. Recent comparative studies demonstrate that GraphRAG approaches significantly outperform traditional vector-based RAG systems[10] in enterprise applications, with accuracy improvements from 0% to over 90% for complex knowledge retrieval tasks[11, 12].

Manufacturing operations are inherently fragmented, highly interdependent, and often involve latent knowledge, making traditional approaches inadequate. By extracting and structuring this complex information before embedding it into a Knowledge Management system, we achieved a far more intelligent and context-aware retrieval process. This approach aligns with recent advances in graph neural retrieval frameworks that demonstrate state-of-the-art performance on complex reasoning tasks[13].

We piloted the software with Sales Engineers in their specialization. The underlying document base consisted of Application Instructions and Technical Specifications. Initial feedback was enthusiastic and scaling the system should result in higher-quality decision-making, less turnaround time and higher employee satisfaction.

## 4. Lessons Learned: Semantic Technologies in Manufacturing

Technical challenges center primarily on data quality and schema evolution requirements. The provided documentation often repeats but with different terminology and contains often outdated information. Schema evolution proves particularly challenging as manufacturing processes and equipment evolve, requiring knowledge graph structures that can accommodate change without breaking existing functionality.

Document format diversity presents ongoing challenges for content ingestion. Manufacturing organizations typically maintain information in dozens of different formats, many of which require specialized processing approaches. This challenge suggests the value of standardized documentation approaches within manufacturing organizations, as advocated by recent Industry 4.0 ontology standardization efforts[14].

Organizational challenges focus on user adoption and knowledge contribution patterns. Engineers demonstrate enthusiasm for knowledge access capabilities but show resistance to knowledge contribution requirements. This finding echoes challenges reported in other industrial knowledge graph deployments[15, 16]. In future projects we will focus on being able to submit and interact with the Knowledge Graph from the user interface.

Best practices emerged from the deployment experience that can guide similar initiatives. Iterative schema development allows for continuous refinement based on user feedback and system performance. This approach aligns with established methodologies for industrial ontology engineering platforms[14] and semantic integration frameworks[17].

# Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

## References

- [1] G. Buchgeher, D. Gabauer, J. Martinez-Gil, L. Ehrlinger, Knowledge graphs in manufacturing and production: A systematic literature review, arXiv preprint arXiv:2012.09049 (2020).
- [2] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, X. Wu, Unifying large language models and knowledge graphs: A roadmap, *IEEE Transactions on Knowledge and Data Engineering* 36 (2024) 3580–3599.
- [3] Q. Cao, I. Grangel-Gonzalez, L. Du, A knowledge graph-based approach for the quality management of bosch, in: *Proceedings of ISWC 2023 Posters, Demos and Industry Tracks*, volume 3828 of *CEUR Workshop Proceedings*, 2023.
- [4] E. G. Kalaycı, I. Grangel González, F. Lösch, G. Xiao, A. ul Mehdi, E. Kharlamov, D. Calvanese, Semantic integration of bosch manufacturing data using virtual knowledge graphs, in: *International Semantic Web Conference*, Springer, 2020, pp. 464–481.
- [5] A. Mavridis, et al., Large language models for intelligent rdf knowledge graph construction: results from medical ontology mapping, *Frontiers in Artificial Intelligence* 8 (2025) 1546179.
- [6] X. Bai, et al., Construction of a knowledge graph for framework material enabled by large language models and its application, *npj Computational Materials* 11 (2025) 51.
- [7] K. Affolter, K. Stockinger, A. Bernstein, A comparative survey of recent natural language interfaces for databases, *VLDB Journal* 28 (2019) 793–819.
- [8] Y. Guo, et al., Talk2data: A natural language interface for exploratory visual analysis via question decomposition, *ACM Transactions on Interactive Intelligent Systems* 14 (2024) 1–24.
- [9] W. Yu, L. Wu, Y. Deng, R. Mahindru, Q. Zeng, S. Guven, M. Jiang, A technical question answering system with transfer learning, in: *Proceedings of EMNLP 2020 System Demonstrations*, 2020, pp. 92–99.
- [10] H. Han, et al., Retrieval-augmented generation with graphs (graphrag), arXiv preprint arXiv:2501.00309 (2025).
- [11] J. Sequeda, D. Allemang, B. Jacob, A benchmark to understand the role of knowledge graphs on large language model’s accuracy for question answering on enterprise sql databases, in: *Proceedings of the 7th Joint Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, 2024.
- [12] Microsoft Research, GraphRAG: Unlocking LLM discovery on narrative private data, Technical Report, Microsoft Corporation, 2024.
- [13] C. Mavromatis, G. Karypis, Gnn-rag: Graph neural retrieval for large language model reasoning, arXiv preprint arXiv:2405.20139 (2024).
- [14] M. H. Brynildsen, C. Jakobsen, N. Abildgaard, C. Woods, Building an industrial ontology engineering platform, in: *Proceedings of ISWC 2024 Posters, Demos and Industry Tracks*, volume 3828 of *CEUR Workshop Proceedings*, 2024.
- [15] J. Martinez-Gil, G. Buchgeher, D. Gabauer, et al., Root cause analysis in industrial domain using knowledge graphs, in: *Procedia Computer Science*, volume 200, 2022, pp. 944–953.
- [16] J. M. Rožanec, P. Zajec, K. Kenda, et al., Xai-kg: Knowledge graph to support explainable ai in manufacturing, in: *CAiSE 2021 Workshops*, 2021.
- [17] Z. Zheng, B. Zhou, Z. Tan, O. Savkovic, D. Rincon-Yanez, N. Nikolov, D. Roman, A. Soylu, E. Kharlamov, Semantic cloud system for scaling data science solutions for welding at bosch, in: *Proceedings of ISWC 2024 Posters, Demos and Industry Tracks*, volume 3828 of *CEUR Workshop Proceedings*, 2024.