ISBN: 978-981-94-4198-3
2025 the 15th International Workshop on Computer Science and Engineering (WCSE 2025)
doi: 10.18178/wcse.2025.06.003

An Aviation Process Knowledge Question-Answering System Based on Knowledge Graph and Large Language Model

Qian Zhao[1+], Wenhao Xu[2], Rui Zuo[1], Ruiqiang Lv[1], Tingting Du[1]

[1] AVIC Manufacturing Technology Institute, Beijing, China

[2] School of Software, Beihang University, Beijing, China

Abstract. Discrete manufacturing enterprises like aviation face significant challenges in managing fragmented process knowledge. This paper presents an innovative question-answering system that synergizes knowledge graphs with large language models to overcome these limitations. Our approach systematically processes multi-source aviation data through dual pathways: constructing a structured aviation process knowledge graph (AeroKG) while simultaneously building a vector database for unstructured documents using LangChain and FAISS. The integrated system leverages retrieval-augmented generation to enhance the DeepSeek-v3 LLM, achieving a 0.91 F1-score - a 15% improvement over conventional methods. Experimental validation confirms the system's robust performance in supporting critical manufacturing operations, from process design to parameter optimization. The implemented web interface demonstrates practical viability for real-world industrial knowledge management applications.

Keywords: Aviation Process QA, Knowledge Graph, LLM

1. Introduction

In recent years, the informatization and intelligence levels of aviation manufacturing enterprises have significantly improved. Process knowledge is one of the critical knowledge assets for aviation manufacturing enterprises, possessing high utilization value. Despite the continuous iterative upgrades of technological methods, the management of aviation process knowledge still faces several challenges, such as knowledge fragmentation and barriers to sharing, difficulties in transferring tacit knowledge, insufficient dynamic updates and compliance management, challenges in integrating multi-source data, and inconvenient application of knowledge. Therefore, effectively integrating, managing, and applying multi-source heterogeneous aviation process knowledge has become an urgent problem that enterprises must address in advancing smart manufacturing.

Process knowledge management has evolved from paper-based manual data handling to Knowledge Management Systems (KMS), adopting informatization and digitalization for easier retrieval and utilization. Ontology and knowledge graph-based models now enable efficient integration, retrieval, and even recommendation and reasoning of process knowledge. Yujue Wang et al. [1] created an ontology repository for qualitative component processing techniques to integrate mechanical component data. Shixin Peng et al. [2] built a knowledge graph from database-extracted process knowledge ontologies. Jinzhou Zhu[3] applied knowledge graphs to UAV wiring harness assembly. However, aviation process knowledge bases remain underdeveloped, lacking comprehensive production process coverage.

Traditional QA systems mainly use rule bases, information retrieval, and shallow ML models. Though effective in specific domains with good interpretability, they struggle with complex semantics and multi-turn dialogues. Advances in deep learning and NLP have greatly enhanced intelligent QA performance. Zhang et al. [4] combined KG with LLM for TCM medical QA. Sun et al. [5] used KG prompts and aerospace data to enhance LLM QA. Jia et al. [6] applied LLMs to oil/gas enterprise knowledge management. LLMs overcome KG limitations in semantic understanding and text generation through pre-training/fine-tuning. Current aviation process QA still relies mainly on rules/retrieval methods.

To solve these challenges, we construct an aviation process knowledge QA system based on knowledge graphs and large language models. The method involves preprocessing multi-source heterogeneous aviation process data by constructing an aviation process knowledge graph and a process vector knowledge base. It then uses a retrieval-augmented large language model technique to develop the APK-DEEPSEEK-v3 QA

[+] Corresponding author.
E-mail address: zhaoqian0912@126.com

This research was funded by the Aero Science Foundation of China（Grant No. 2024M050025001）

model. Finally, the aviation process knowledge QA system is built using the aforementioned technologies, and its feasibility is validated through experiments.

2. Overall Framework

This paper aims to construct an aviation process knowledge QA system based on KG and LLM. It focuses on building a process KG and a process vector knowledge base in the aviation process domain to provide reliable contextual support for the LLM, enabling intelligent QA for aviation process knowledge. The overall framework consists of four main modules, as shown in Figure 1.

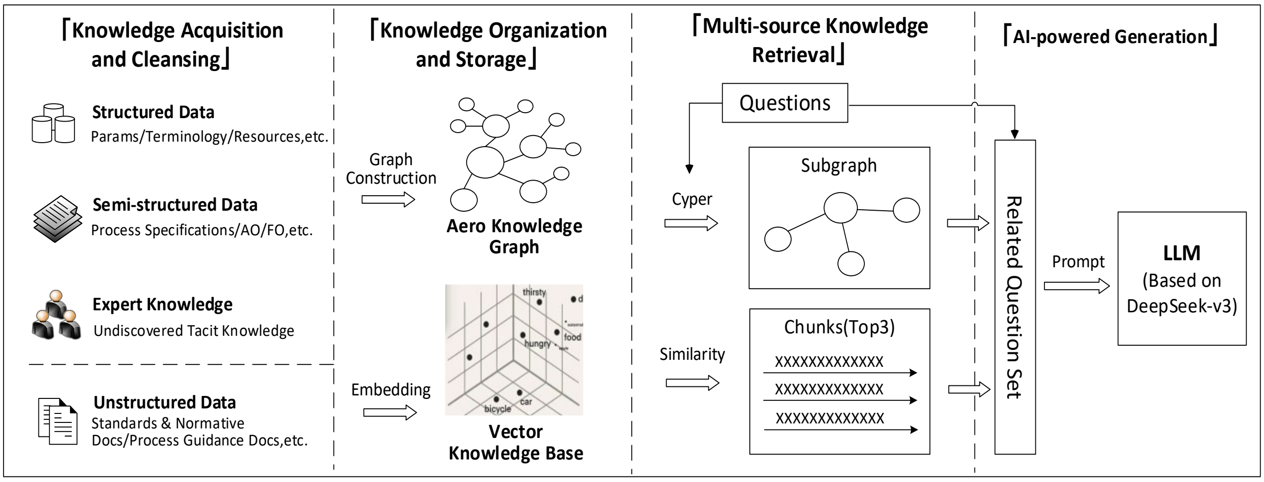

Fig. 1: Overall Framework

① Knowledge Acquisition and Cleaning: Multi-source heterogeneous aviation process data are collected, organized, and analysed to remove noise, handle missing values, and eliminate duplicate data, ensuring knowledge acquisition and cleaning.

② Knowledge Organization and Storage: The collected process data comprises two primary types: structured/semi-structured data, which is organized into a multi-granularity aviation process knowledge graph (AeroKG) for semantic storage and efficient knowledge retrieval, and unstructured document data, whose text content is sliced and vectorized to establish a process vector knowledge base (VectorKB) enabling efficient retrieval.

③ Multi-source Knowledge Retrieval: A combined retrieval strategy is proposed for answering questions. This strategy involves subgraph retrieval from AeroKG and knowledge matching from VectorKB to acquire domain-relevant knowledge.

④ AI-powered generation: Use prompt learning techniques to manually construct prompt word templates, combining the user question with the retrieved relevant domain knowledge as context information. This context is then provided to the large language model in the form of prompt words, helping the model better understand and generate answers.

3. Construction of Aviation Process Knowledge Graph

This study develops a structured approach to process knowledge management, handling both structured and semi-structured data through acquisition, extraction, organization, and storage. The methodology establishes a unified process knowledge system, ultimately constructing an aviation domain knowledge graph.

Knowledge graphs are semantic networks that represent entity relationships as structured knowledge bases. Unlike traditional approaches, they emphasize knowledge interconnectivity and logical structure through node-entity and edge-relationship representations. The construction of the process knowledge graph primarily consists of two layers: the schema layer and the data layer. Using ontology modeling, we developed an aviation process model focusing on four core domains. Through R2RML-based mapping and rule-based methods, we extracted structured/semi-structured data from aviation manufacturing knowledge, enabling the knowledge graph construction.

In this study, the construction of the aviation process knowledge graph consists of four main components. The following sections provide a detailed explanation of each component.

3.1. Process Knowledge Collation

Process knowledge encompasses the technical expertise, operational standards, process parameters, and quality control methods developed throughout the manufacturing lifecycle from design and production to inspection and maintenance, with effective process knowledge management requiring defining knowledge scope, classification, and sources while systematically organizing scattered multi-source knowledge, where in aviation manufacturing this knowledge primarily consists of four key categories as illustrated in Figure 2.

Fundamental Process Knowledge: This includes terminology, machining principles, material properties process standards, equipment operation guidelines, etc.

Specialized Process Technology: Summarizes the unique knowledge of each process, covering machining, forming, assembly, composites, welding, surface treatment, etc.

Procedural Knowledge: Knowledge covering the entire workflow from process design to process manufacture, including process routes, process, steps, features, NC programs, etc.

Manufacturing Resource Knowledge: Knowledge related to manufacturing resources used in process design, tooling design, and machining programming.

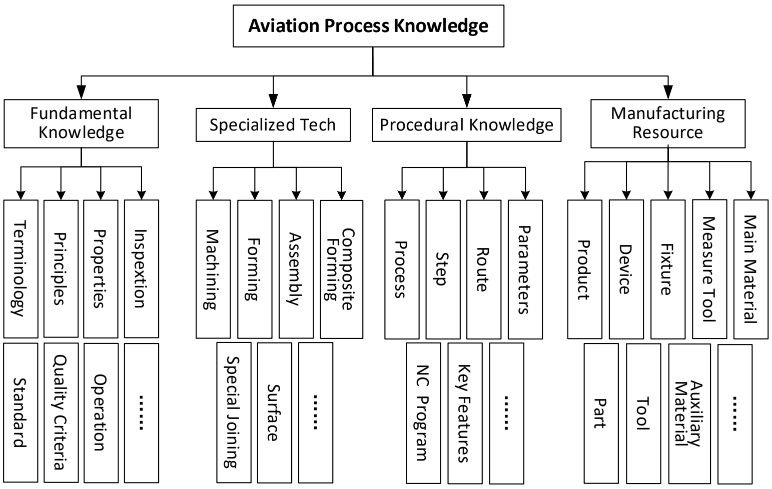

Fig.2: Aviation Process Knowledge

3.2. Process Ontology Modeling

Ontology modeling technology serves as an effective knowledge representation method that enables unified process data expression with inherent scalability and modularity. This study formalizes the aviation process ontology model as O=<C,R,A,E,F>, where:C (Concepts) denotes entity class collections, R (Relations) defines inter-concept logical connections (composition/inheritance/etc.), A (Attributes) specifies concept/entity characteristics and values, E (Entities) represents instantiated concept members, F (Functions/Axioms) establishes relationship constraints.

The framework systematically organizes aviation process knowledge into a networked structure where: Concepts encompass multiple entities inheriting all parental attributes, Cross-concept entity associations exist, Knowledge is ultimately stored as either <entity(concept)-relation-entity(concept)> or <entity-attribute-value> triples. This paper constructs an aviation process ontology model centred on fundamental process knowledge, specialized process technology, process knowledge and etc, as shown in Figure 3.

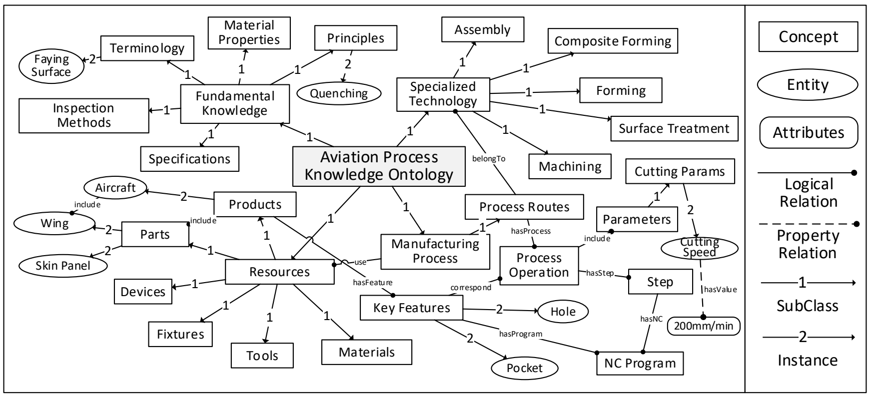

Fig. 3: Aviation Process Ontology Model

3.3. Process Knowledge Extraction

Knowledge extraction refers to the automatic identification of entities, relations, attributes, etc. Since 60% of the process data resulting is structured data stored in relational databases, and 10% is semi-structured data stored in Word and Excel documents, the process knowledge extraction discussed in this study primarily focuses on structured and semi-structured data.

For structured data, a pattern-based data mapping mechanism is used to convert the structured data into a networked knowledge structure, which is the RDB2RDF data conversion process. We use R2RML mapping to design multiple retrieval data logic tables for different process data tables. Each logic table is converted into RDF by a triple mapping, that is, each row of instance data in the logic table is mapped to several RDF triples. The R2RML mapping mechanism expression is:

$$triplesMap = \{logicalMap, subMap, preobjMap\} \tag{1}$$

Where, $logicalMap$ id used to describe the database table name; $subMap$ is the common subject of all RDF triple; $preobjMap$ is that each mapping consists of a predicate mapping and an object or a value mapping. RDF triples are generated based on the above mapping mechanism. For semi-structured data stored in Excel spreadsheets, we employ file parsing and rule-based extraction techniques. The Apache POI library in Java is used for it.

Finally, based on the aviation process ontology model and the extracted knowledge, the triples are converted into nodes, edges, and attribute values and stored in the Neo4j graph database, enabling information retrieval via Cypher queries.

4. Construction of Process Vector Knowledge Base

The process vector knowledge base is a specialized type of database that differs fundamentally from traditional relational databases in how it stores and processes information. In a process vector knowledge base, data is represented as vectors in a high-dimensional space, capturing semantic features and relationships. As a result, for large-scale unstructured data, this approach enables fast and efficient retrieval and query capabilities. To enable the retrieval of unstructured data such as documents, this paper proposes constructing a process vector knowledge base to vectorize useful text from documents. This approach transforms textual data into a vector format that is more easily understood and processed by computers while preserving key features and semantic information.

In the preliminary analysis of process knowledge sources, approximately 30% of the data originates from process documents, including national, industry, and enterprise standards, as well as process operation guidelines. This paper transforms these documents into a process vector knowledge base, utilizes the Lang Chain framework to construct a process vector knowledge base, with the main workflow divided into the following five steps:

(1) Data Collection and Preprocessing: A total of 210 standard specifications and 108 process guidance manuals in PDF or Word format were collected. After text extraction, data cleaning was performed to remove noise and irrelevant information.

(2) Text Chunking: Long texts were segmented into smaller chunks suitable for embedding. Two key parameters, `chunk_size` and `chunk_overlap`, were used to define it.

(3) Embedding Generation: Embeddings were generated using the open-source Sentence-BERT model, which converts segment into a low-dimensional vector.

(4) Vector Storage and Index Optimization: The FAISS open-source lightweight vector database was chosen for storing vectorized data. The indexing process employed the Hierarchical Navigable Small World (HNSW) index structure to facilitate approximate nearest-neighbour search using a graph-based approach, optimizing retrieval efficiency.

(5) Vector Database Persistence: Store vector data and original text data in association to ensure that the answer can be traced back to the source of the text.

5. Construction of Question Answering Model Based on Retrieval-enhanced Generation

We develop the APK-DEEPSEEK-v3 QA model using retrieval-augmented generation (RAG). The system first queries both the process knowledge graph and vector database to retrieve relevant information, then employs prompt engineering to guide the LLM in generating accurate answers.

Modern open-source LLMs span diverse applications, from general-purpose (LLaMA, Falcon) to specialized (BioGPT, FinGPT) and multilingual (ChatGLM) variants. For this study, we selected DeepSeek-v3 for its efficient design—employing pruning, quantization, and distillation to balance performance with resource efficiency—making it ideal for private deployment.

Models like DeepSeek-v3 have demonstrated significant capabilities in language understanding and generation, but they still have limitations, such as insufficient generation ability in vertical domains, lower accuracy, or the occurrence of "hallucinations." Compared to fine-tuning large language models, using retrieval-augmented generation (RAG) techniques combined with domain-specific knowledge bases can improve the QA capabilities of large language models, even in the absence of extensive training data. Therefore, in this study, aviation process knowledge graph and process vector knowledge base are used as sources of information for retrieval-augmented access, further enhancing the QA performance of the DeepSeek large language model in the field of aviation processes.

The process of building an RAG-based QA model is divided into the following four steps:

(1) Graph Retrieval. The relevant entities and relationships related to the question are retrieved from KG to provide structured knowledge support.

(2) Vector Retrieval. The system retrieves top-3 semantically similar text fragments from unstructured documents using cosine similarity (Formula 2) between question and document vectors.

$$similarity = \frac{A \cdot B}{||A||\,||B||} \tag{2}$$

(3) Prompt Word Construction. The system integrates retrieved subgraphs and vector-matched content into structured prompts, converting them into natural language context. It combines this knowledge with user queries using predefined templates (e.g., "As an aviation process expert, answer based on: {retrieved results}. Question: {user question}") to guide LLM response generation.

(4) Answer Generation. The locally deployed APK-DEEPSEEK-v3 LLM generates an answer based on the constructed prompt word and marks the knowledge source in the answer.

6. Question Answering System Implementation and Evaluation

We developed and validated a QA system through: (1) processing 20K+ multi-source data points into a Neo4j knowledge graph (20,345 entities, 31,850 facts); (2) creating a FAISS vector database from 318 documents using LangChain; and (3) integrating Cypher-based graph retrieval with vector search to power DeepSeek-v3 responses. The deployed web interface demonstrates effective knowledge access for aviation process queries.

6.1. Ablation Study

To evaluate the individual contributions of each component in system, we conducted an ablation study comparing four configurations: KG, VectorKB, DeepSeek-v3 and APK-DEEPSEEK-v3. Table 1 shows the performance indicators on the test set.

Table 1: Ablation Study Results

Model	Accuracy	Recall	F1 Score	Response time
KG	0.82	0.78	0.80	1.2
VectorKB	0.84	0.81	0.82	1.3
DeepSeek-v3 (no RAG)	0.85	0.80	0.82	2.0
APK-DEEPSEEK-v3	0.92	0.90	0.91	1.5

6.2. Subjective Evaluation

Taking the question “In the manufacture of aircraft engine blades, what is the heat treatment temperature range of Ti-6Al-4V titanium alloy? If the heat treatment temperature is too high or too low, what effect will it have on the mechanical properties of the blades?” as an example, the generated results of different models for the problem are compared (Table 2).

Table 2: Generated Results of Different Models for the Target Questions

Model	Answer	Evaluation
KG or VectorKB	Ti-6Al-4V titanium alloy is typically heat-treated at 900°C.	Is brief and lacks contextual details.
DeepSeek-v3	The heat treatment for Ti-6Al-4V titanium alloy generally ranges between 850°C and 950°C. If the temperature is too high, it may lead to grain coarsening, reducing the material's strength and toughness; if too low, residual stresses may not be fully relieved, affecting the blades's fatigue performance.	More comprehensive but lacks specific data references or cited sources.
APK-DEEPSEEKv3	The standard heat treatment temperature for Ti-6Al-4V titanium alloy is 900°C, with an acceptable range of 850°C to 950°C(Source: Process Manual, Chapter 5). Excessive temperatures (e.g., above 950°C) may cause grain coarsening, impairing strength and toughness, while insufficient temperatures (e.g., below 850°C) may fail to fully eliminate residual stresses, compromising blade fatigue resistance (Source: Journal of Materials Science, 2022). In production, heat treatment parameters should be adjusted based on blade geometry and service conditions.	Is detailed, accurate and supported by authoritative references.

6.3. Scalability and Dynamic Updates

The system employs a distributed architecture with incremental processing for scalable knowledge management. The KG utilizes Neo4j for horizontal scaling, enhanced by read replicas and caching to optimize query performance. Updates are efficiently managed through batch processing, real-time streaming for critical changes, and automated ontology propagation. For vector operations, FAISS provides high-performance indexing with incremental embedding capabilities, enabling seamless document integration. Dynamic updates are maintained through automated monitoring, database triggers, and periodic audits.

7. Summary and Prospect

This paper presents an aviation process knowledge QA system based on KG and LLM. This approach improves the knowledge generation accuracy, interpretability, and domain adaptability of the QA system to some extent. Future work will focus on three key areas: (1) Automated knowledge maintenance using ML to detect stale data, propose updates from literature, and validate changes with experts; (2) Scalability enhancements through graph neural networks for efficient retrieval, alternative vector databases for large document sets; (3) A continuous evaluation framework featuring automated accuracy testing, user feedback systems, and aviation-specific QA benchmarks. These developments will overcome current limitations and enable enterprise-scale deployment across manufacturing facilities.

8. References

[1] Wang Yujue, Liang Yuhao, Wang Suqin, et al. Construction of an Ontology Library for Machining Process of Mechanical Parts [J]. Computer Science, 2022, 49(S1): 661–666.

[2] Peng Shixin, Xiao Biao, Zhao Zhengcai, et al. Construction and Application of a Knowledge Graph for the Processing Technology of Aerospace Complex Thin-walled Parts [J]. Mechatronic Engineering, 2024, 41(4): 709–710.

[3] Zhu Jinzhuo. A Knowledge Service Method for UAV Wiring Harness Assembly Process Based on Knowledge Graph [D]. East China University of Science and Technology, 2024.

[4] Zhang Heyi, Wang Xin, Han Lifan, et al. Research on a Question-answering System Integrating Large Language Models and Knowledge Graph [J]. Journal of Computer Science and Exploration, 2023, 17(10): 2377–2388.

[5] Yan SUN, Lixin ZHOU, Lianjun SUN, Bin YE, Guolin WANG. Fusion Application of Knowledge Graph and Big Language Model to Aerospace TT&C Question-answering System[J]. AEROSPACE SHANGHAI, 2024, 41(5):

[6] TANG J,PANG D W,LIU S M , et al.A Comparative Study on the Effects of Fine-tuning Large Language Models and Retrieval Enhanced Generation Technology in Oil and Gas Enterprise Regulation Question Answering Application[J].Digital Communication World, 2024, (11): 104-106.