Optimizing Aerospace Product Maintenance

A Novel Multi-Modal Knowledge Graph and LLM Approach for Enhanced Decision Support

Raed Awill[1,2], Wajahat Ali Khan[1], Maqbool Hussain[1], and Ben Anderson[2]

[1] University of Derby, UK
[2] Addqual Ltd, Derby, UK

Abstract. Siloed and inaccessible repair knowledge hinders the efficient maintenance of critical Turbine Engine components in the aerospace industry. This research introduces a novel multi-modal knowledge graph, leveraging Natural Language Processing (NLP) and Large Language Models (LLMs) to extract and structure repair rules from unstructured documents into a 131-node, 148-relationship graph. This advancement enables immediate access to essential information and facilitates data-driven decision-making, enhancing repair accuracy and efficiency. Implemented at AddQual Ltd., the knowledge graph reduced information retrieval times by 70%, increased repair speed by 20%, and is projected to yield 20% annual cost savings. These results highlight the transformative potential of integrating AI with knowledge graphs in aerospace maintenance. Future work will focus on advancing robust data validation frameworks and developing adaptive AI algorithms, extending the benefits across the aerospace sector and beyond.

Keywords: Knowledge Graph, Multi-Modal, Turbine Repair, Decision Support, Aerospace, Text2Cypher

1 Introduction

In the fast-paced aerospace industry, maintaining turbine engine components such as blades is essential for safety and optimal performance [1]. Building on our previous work [2] that highlighted the challenges in knowledge extraction from structured sources, our current research extends these capabilities to unstructured PDF documents [7], enhancing knowledge management in aerospace maintenance. Unstructured documents, traditionally difficult to navigate and integrate into decision-making processes, have impeded the effectiveness of repair strategies, impacting component reliability. Recognizing this issue, our research introduces an effective approach by developing a multi-modal knowledge graph. This graph is designed to systematically capture, structure, and make accessible the vast and previously hidden repair knowledge contained within unstructured documents. By leveraging NLP with LLMs capabilities, we have extracted, categorized, and linked engine components, their images, repair rules, measurements, and actions related to engine component maintenance into a coherent, navigable structure. Our proposed knowledge graph addresses three core challenges in aerospace maintenance: (1) enhancing the accessibility of repair knowledge, thus overcoming the limitations posed by traditional document formats; (2) enabling the utilization of existing data for improved decision-making, by structuring and linking measurement data within the graph; and (3) ensuring consistent and informed repair decisions through a standardized representation of knowledge that supports data-driven analysis. The adoption of our knowledge graph at AddQual Ltd. demonstrates its practical value, reducing information retrieval times and enhancing the consistency and efficiency of the repair processes.

2 Approach

Our system employs Neo4j, a graph database ideal for handling complex, multi-modal data (text, visuals, and numerical data) in aerospace maintenance. Its adaptability is essential for seamless data integration, continuous system development, and sustained effectiveness in dynamic operational environments [4]. Furthermore, Cypher, Neo4j's query language, enables precise and efficient data retrieval, preserving semantic integrity as the data landscape evolves [3].

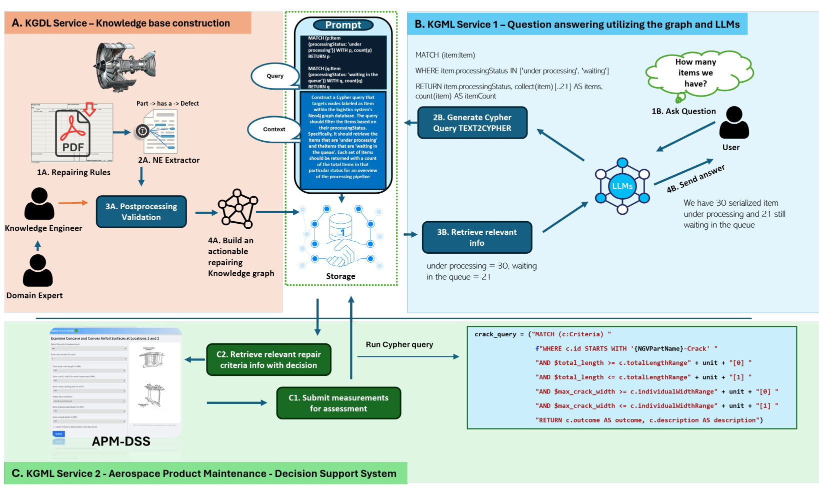

Fig. 1. Workflow Diagram of the Integrated Multi-Modal KG System

2.1 A: Knowledge Graph Definition Language Service (KGDL)

The system integrates multi-modal data from unstructured documents into a structured graph, extracting key repair rules, measurements, and images of engine components for maintenance. The LLM parses and organizes these data types per the input schema, enhancing the semantic depth of the graph for better usability. It starts with a domain expert converting a PDF with KG schema into a Neo4j graph using an LLM. The PDF is segmented, and processed by the LLM to extract data, and Cypher queries are generated and executed in Neo4j. This forms the basis of a DSS system [6], detailed in Figure 1, which shows the steps to build a validated knowledge base. Additionally, our semantic technologies not only store data but also enable advanced reasoning. By extracting entities like 'compressor' from PDFs and using iterative prompting with GPT-4, which improved based on the aerospace domain experts' feedback, precise Cypher queries are generated. For instance, it identifies necessary inspections and retrieves relevant maintenance actions for components like 'Gas Turbine blades', ensuring the data meets domain-specific needs for aerospace maintenance decision-making.

2.2 B: Knowledge Graph Manipulation Language Services (KGML)

KGML, illustrated in Figure 1, comprises two interconnected subsystems designed to maximize the knowledge graph's utility. KGML Service 1 facilitates user interaction through an NLP feature. This service allows users to submit queries about engine components using natural language, and KGML automatically translates these queries into Cypher queries suitable for the underlying graph database. This seamless integration enables users to receive clear and accurate information promptly. The KGML Service 2 serves as a dedicated Decision Support System (DSS) for aerospace product maintenance, streamlining the maintenance decision-making process. It automates the retrieval of relevant repair rules and efficiently stores and analyzes user-entered measurements. The system then executes complex Cypher queries to align the outcomes with established maintenance standards. By providing a deep analysis of technical data, this subsystem empowers users to make decisions. To address trust and reliability concerns, the KGML serves as a decision-support tool, complementing rather than replacing human judgment. Technicians use the outputs from the LLM to efficiently repair engine defects, benefiting from the knowledge graph's explainable insights for transparent decision-making. Our approach includes documented procedures for technician verification and managerial oversight, ensuring responsible integration of LLM outputs into the critical maintenance workflows.

3 Results, Insights, and Future Work

Implementing our multi-modal knowledge graph at AddQual Ltd. has substantially improved the efficiency and accuracy of aerospace engine maintenance. The knowledge graph displayed high entity recognition accuracy at 96.5% and relationship extraction accuracy at 95.2%, evaluated using a validation set of 500 entities, encompassing engine components, components images, and repair procedures. This thorough validation approach included automated data extraction by

LLM algorithms, supplemented by critical evaluations from domain experts, ensuring the results are practically applicable. Drawing on the Holistic Evaluation of Language Models (HELM) [5], GPT-4 exhibits a Robustness score of 88% in Cypher Query Construction from the extracted entities, demonstrating reliable performance in handling the complex linguistic challenges found in aerospace maintenance texts. The validation set was specifically designed to challenge the system with complex scenarios typical in aerospace maintenance, thereby proving the knowledge graph's capability to manage, categorize, and retrieve complex information accurately. This led to a marked improvement in operational efficiency: repair technicians experienced a 70% faster retrieval of necessary maintenance data, and the system handled over 2,000 user queries with an impressive average response time of just 4 seconds. Furthermore, leveraging historical user-entered measurement data has optimized repair decision-making, boosting the consistency of repair actions across different tasks by 35% and enhancing overall operational efficiency by 20%. These improvements have significantly reduced maintenance time, saving 240 hours over six months. Additionally, end-user feedback has been overwhelmingly positive, with a 90% satisfaction rate among repair personnel, underscoring the system's usability and effectiveness in improving decision-making processes. This research has illuminated the critical role of structured knowledge integration in improving decision-making within aerospace maintenance, offering a profound lesson on the interplay between domain-specific knowledge and AI technologies. We learned that the efficacy of AI-driven systems heavily relies on the depth and accuracy of the underlying data models, as evidenced by our knowledge graph's significant impact on operational efficiency and decision-making precision. Moving forward, it is essential to prioritize robust data validation frameworks and adaptive AI algorithms to further enhance the reliability and applicability of such systems across varied industrial landscapes.

References

1. Addqual | non-destructive testing inspection (2024), https://www.addqual.com/
2. Awill, R., Khan, W.A., Hussain, M., Zada, S., Anderson, B.: Aerospace qualification services knowledge graph: A leap towards enhanced data management. In: Proceedings of the International Semantic Web Conference 2023 (Nov 2023)
3. Holzschuher, F., Peinl, R.: Performance of graph query languages: comparison of cypher, gremlin and native access in neo4j. In: International Conference on Extending Database Technology (2013)
4. Kaur, K., Rani, R.: Modeling and querying data in nosql databases. In: 2013 IEEE International Conference on Big Data. pp. 1–7 (2013). https://doi.org/10.1109/BigData.2013.6691765
5. Liang, P., Bommasani, R., et al., T.L.: Holistic evaluation of language models (2023)
6. Peng, C., Xia, F., Naseriparsa, M., Osborne, F.: Knowledge graphs: Opportunities and challenges. Artificial Intelligence Review **56**(11), 13071–13102 (2023)
7. Zhao, H., Pan, Y., Yang, F.: Research on information extraction of technical documents and construction of domain knowledge graph. IEEE Access **8**, 168087–168098 (01 2020)