



FROM DOWNTIME TO UPTIME

AI-DRIVEN FAULT DIAGNOSIS WITH GRAPHRAG

Executive Summary

Manufacturers face increasing pressure to reduce downtime and boost operational efficiency, yet diagnosing machine faults remains a slow and error-prone process. A new AI-powered approach, Graph-based Retrieval-Augmented Generation (GraphRAG), is changing that. Developed through project activities in collaboration with the Fraunhofer Innovation Platform at the University of Twente and Dutch manufacturer Mintres BV, this system transforms years of scattered maintenance logs and manuals into structured knowledge that helps engineers troubleshoot more quickly and accurately. By turning unstructured data into intelligent support, GraphRAG enables teams to make smarter decisions, reduce downtime, and retain critical expertise on the factory floor.

Why Maintenance Is Still a Struggle

Despite advances in manufacturing technology, fault diagnosis remains a major bottleneck, largely due to how maintenance data is recorded and stored. Logs are often written in free text, with inconsistent levels of detail, and captured in a variety of formats such as handwritten notes, spreadsheets, PDFs, and emails. These inconsistencies make it difficult to search, analyse, or extract meaningful insights from the data.

Language barriers further complicate matters. Maintenance logs are typically written in local languages for example in Dutch, while equipment manuals and troubleshooting guides are often in

English. This makes cross-referencing difficult and error-prone. Key information ends up siloed, buried, or overlooked.

On top of this, the ageing of both machinery and skilled personnel compounds the challenge. As experienced engineers retire or move on, their knowledge frequently leaves with them. Without a structured way to retain and share this expertise, fault diagnosis becomes slower, less reliable, and more expensive.

GraphRAG: Smarter AI for Maintenance

GraphRAG is a next-generation artificial intelligence approach designed to make maintenance work faster, more accurate, and more consistent. It builds

on a method known as Retrieval-Augmented Generation (RAG), which enhances large language models (LLMs) like GPT by enabling them to draw from a company’s internal documents and historical records when answering specific technical questions.

While standard RAG systems retrieve relevant documents to support answers, GraphRAG takes things further by converting information into a structured knowledge graph. This graph links fault locations, symptoms, causes, and corrective actions, allowing the system to recognise patterns and offer intelligent, context-aware recommendations. The result is a solution that not only understands language, but also comprehends how your machines operate, where they fail, and how to resolve those failures effectively.

Use Case: Fault Diagnosis in Ion Beam Machine

Mintres BV applied GraphRAG to an Ion Beam Machine (IBM)—a complex piece of equipment involving multiple subsystems such as vacuum pumps, RF generators, and beam controllers. Historical maintenance logs, mostly in Dutch, alongside technical manuals in English, contained a wealth of information that had been difficult to access or utilise effectively.

Using GraphRAG, this unstructured data was extracted and organised into a single, multilingual knowledge graph. Engineers could now interact with the system using a conversational interface, asking questions and receiving answers drawn from years of recorded experience. The AI could connect symptoms to likely causes and proven solutions, regardless of language or document format. What was once a time-consuming search through siloed records became a seamless, knowledge-driven troubleshooting process.

Business Impact: Why This Matters

With a clear example of GraphRAG in action, it becomes clear how this technology delivers tangible business value. Converting unstructured, multilingual maintenance data into structured knowledge unlocks a competitive advantage that goes beyond individual fixes. Engineers are better equipped to detect recurring faults and root causes that would otherwise remain hidden. The system captures the insights of experienced

technicians and shares them across the organisation, so even newer team members can act with confidence.

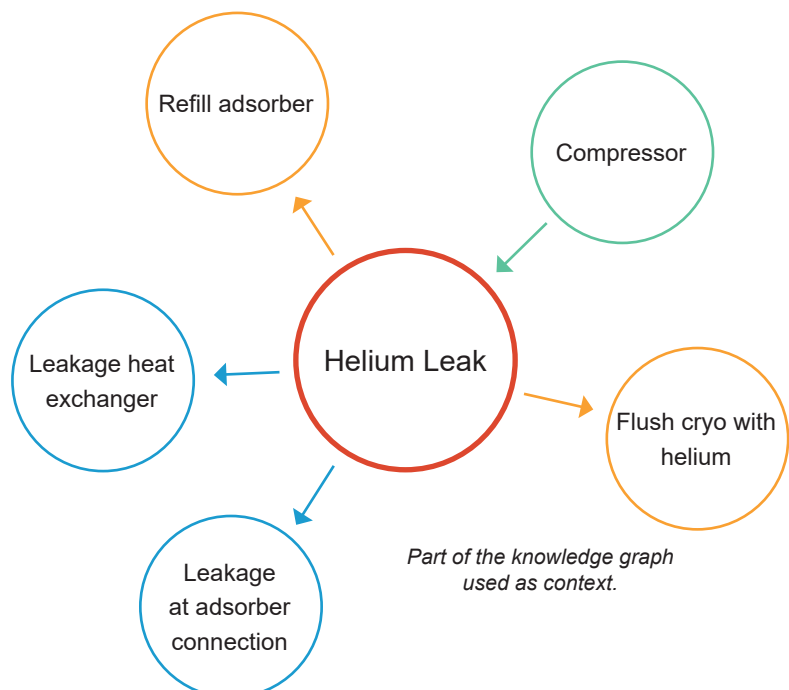
The result is more accurate diagnostics, fewer repeat failures, and a smoother maintenance process. This translates into shorter downtimes, lower costs, and greater consistency. Perhaps most importantly, it protects organisations from knowledge loss by ensuring that what senior engineers know is retained, shared, and built upon. GraphRAG doesn’t just help you respond more quickly; it also supports the development of long-term resilience.

Conversational Fault Diagnosis with GraphRAG



I found a problem in the compressor: a helium leak in the heat exchanger. What is the cause and the solution?

The helium leak in the heat exchanger may be due to leakage at the exchanger or adsorber connection. Recommended actions include replacing the heat exchanger and adsorber, flushing the cryo system and hoses with helium, and refilling the adsorber with activated carbon.



“With a clear example of GraphRAG in action, it becomes clear how this technology delivers tangible business value. Converting unstructured, multilingual maintenance data into structured knowledge unlocks a competitive advantage that goes beyond individual fixes.”

How to Get Started

You do not need to start from scratch to take advantage of GraphRAG. Begin with what you already have, such as maintenance logs, service reports, manuals, and shift records. Choose a single machine or system as a pilot to prove the concept. This focused approach helps manage scope and build confidence before scaling further.

Engage domain experts early. Their experience is essential in validating data accuracy and ensuring the output is both useful and trustworthy. Collaboration between engineering, data, and IT teams is vital to success. Together, they can align goals and develop an effective integration strategy.

By taking an incremental approach, your organisation can reduce risk and begin seeing benefits quickly.

Each additional system or machine strengthens the knowledge base and increases the system's overall value.

Augmenting, Not Replacing

GraphRAG is not about replacing skilled workers; it is designed to augment their abilities. The system takes on the time-consuming task of digging through logs and manuals, freeing engineers to concentrate on problem-solving and continuous improvement. Their insights become part of a shared knowledge system, allowing the entire team to benefit from their experience.

As team composition changes over time, the knowledge remains. New engineers can ramp up faster, and maintenance practices become more consistent. In this way, GraphRAG

strengthens both individual performance and organisational memory.

AI solutions like GraphRAG are transforming maintenance by making it faster, smarter, and more resilient. Instead of reacting to problems as they arise, manufacturers can anticipate them, diagnose them accurately, and resolve them efficiently. By turning fragmented data into usable knowledge, organisations protect themselves from disruption, reduce operational risk, and unlock real value from their existing information.

The future of maintenance is already within reach. It starts with your data and the decision to put it to work. ■

Author:



Wan Barmawi Muhammad Razaq
MSc Business Information Technology
Student, University of Twente